

IBM FileNet P8
Version 4.5.1

*Plan and Prepare Your Environment for
IBM FileNet P8 for installation on AIX
with IBM DB2 for Linux, UNIX, and
Windows, IBM WebSphere Application
Server, and IBM Tivoli Directory Server*



IBM FileNet P8
Version 4.5.1

*Plan and Prepare Your Environment for
IBM FileNet P8 for installation on AIX
with IBM DB2 for Linux, UNIX, and
Windows, IBM WebSphere Application
Server, and IBM Tivoli Directory Server*



Note

Before using this information and the product it supports, read the information in "Notices" on page 71.

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 18

Configuring UNIX for IBM FileNet P8 servers (all components) 18

Configuring Content Engine servers (all UNIX) 19

Configuring Content Search Engine servers (all UNIX) 19

Configuring Process Engine servers (all UNIX) 19

Configuring Process Engine servers (AIX) 21

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. 32

Configuring IBM Tivoli Directory Server. 32

Specifying IBM FileNet P8 accounts 32

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 40

Preparing DB2 for Linux, UNIX and Windows servers 42

Verifying that DB2 for Linux, UNIX and Windows is installed for IBM FileNet P8 44

Configuring DB2 client 49

Installing DB2 client software (Process Engine). 49

Creating DB2 client instances for Process Engine 49

Cataloging the Process Engine DB2 server node. 49

Creating the Process Engine DB2 database alias 49

Verifying the ability to connect to the database 50

Verifying the Process Engine database connection (DB2) 50

Application Server administrator installation tasks. 50

Creating Content Engine application server accounts 51

Configuring WebSphere for Content Engine 53

Creating the WebSphere profile for Content Engine 54

Specifying the WebSphere environment variables 54

Setting the primary administrative user name 56

Setting host aliases for deployment on multiple servers 56

Setting permissions for the Configuration Manager user. 56

Configuring WebSphere Application Server for Application Engine or Workplace XT 57

Configuring Process Engine clients for ORB 57

Configuring the documentation server 57

Appendix A. Preparing non-English environments for installing IBM FileNet P8 Platform. 59

Application Server administrator 59

Configuring character encoding on WebSphere

Application Server 59

Security administrator. 59

Extended characters and user names 60

Database administrator 60

Installing the DB2 server	60
IT administrator	60
Operating system considerations	61
UNIX	61
Configuring support for other languages in X Windows UNIX system	62
Setting the LANG, LC_TIME, and LC_MESSAGES environment variables	62
IBM FileNet P8 administrator	62
Process Engine services	63
Configuring Process Task Manager for Application Engine and Workplace XT	63

Limitations on installing in a non-English environment	63
---	----

Appendix B. IBM FileNet P8 ports . . . 65

Notices	71
Trademarks	73

Index 75

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 a WebSphere® Application Server managed environment, you must install and configure Content Engine on the network deployment 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.

Non-managed deployment

When you deploy Content Engine in a WebSphere Application Server non-managed environment, you install and configure Content Engine on a single server in the environment. After you deploy the bootstrapped Content Engine EAR file on the initial server, you copy the bootstrapped EAR file to the other servers and complete the configuration tasks for each server. Finally, you deploy the bootstrapped EAR file that was copied from the initial server.

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 50

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 50

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 18

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>
Database user name (DB2® for Linux®, UNIX® and Windows®): <i>ce_db_user</i>	<p>Operating system user account on the database server. This user is granted database permissions for Content Engine access to the DB2 database. The user needs access to each database created for an object store and the GCD database. Separate accounts can be used for each object store, but are not required.</p> <p>A user can be configured to access multiple object store databases, depending on your installation requirements. Permissions granted by the database administrator determine which object store databases any user can access.</p>
<p>Accounts required for DB2 for Linux, UNIX and Windows database installation and administration:</p> <ul style="list-style-type: none"> • Groups: instance owner primary group <i>ce_db_db2_group</i> • Users: instance owner <i>ce_db_db2_instanceowner</i> 	<p>Operating system users and groups that must exist on the database server. The user will create databases and set a number of configuration parameters.</p>

User or group name	Description
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.
Content Engine Operating System user: <i>ce_os_user</i>	<p>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.</p> <p>UNIX For UNIX-based Content Engine and file storage areas, configuring security requires the use of NFS.</p>

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. This includes DB2 database users if you are using DB2 as your database. In a configuration with workload partitioning (WPAR), create the users and groups from within the WPAR. In a Solaris zone configuration, create the users and groups in the local zone. Requirements for database permissions for the DB2 users are documented with tasks related to the DB Administrator. Process Engine on a Windows server can optionally be installed by a domain user and can optionally be configured to run by a domain user. Whether Process Engine is installed or configured to run by a domain user or local user, local users and groups must be configured.

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>fnsww</i>	<p>This user is an operating system user that is used to start and stop Process Engine software.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>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>fnsww</i> name to the installation program.</p> <p>Assign <i>fnswwr</i> as the primary group for the <i>fnsww</i> user.</p> <p>Assign the <i>fnsww</i> user to the following secondary groups:</p> <ul style="list-style-type: none"> • <i>fnswwadmin</i> • <i>fnswwnop</i>
The password for the <i>fnsww</i> user (UNIX): <i>fnsww_password</i>	<p>This password is for the <i>fnsww</i> user.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p>
An operating system group for Process Engine: <i>fnswwnop</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>fnsww</i> user to this group.</p> <p>If you will install Process Engine or run Process Engine Windows services as either the <i>fnsww</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>fnswwnop</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>

User or group name	Description
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
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> • root

User or group name	Description
The database runtime user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_sw</i>	<p>This user is the primary Process Engine DB2 for Linux, UNIX and Windows database run time user and is used only by the Process Engine software to access the DB2 database.</p> <p>For systems using SERVER or SERVER_ENCRYPT authentication, create this user on the database server. For systems using CLIENT authentication, create this user on the Process Engine server.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>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 <i>f_sw</i> is defined, indicate to the Process Engine installation program that aliases will be defined. When prompted, provide the <i>f_sw</i> name to the installation program.</p> <p>After creating the new users and setting their group memberships, log off as root user, log on as each of the new users, and change the password to avoid connection problems the first time they are used.</p> <p>In a farmed configuration, each Process Engine must have the same run time user name.</p>
The password for the database runtime user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_sw_password</i>	<p>This is the password for the <i>f_sw</i> user.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>The Process Engine installation program prompts for this password. The password entered during installation will be used to create a file containing an encrypted version of this password. If the encrypted version of the password does not match the operating system version of the password, Process Engine will fail to connect to the database. If the <i>f_sw_password</i> field is left blank during installation, the value will be set to <i>filenet</i>, which is the default.</p>

User or group name	Description
The database maintenance user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_maint</i>	<p>This user is the primary database maintenance user for Process Engine with DB2 for Linux, UNIX and Windows and is used only by the Process Engine software to access the DB2 database.</p> <p>For systems using SERVER or SERVER_ENCRYPT authentication, create this user on the database server. For systems using CLIENT authentication, create this user on the Process Engine server.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>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 <i>f_maint</i> is defined, indicate to the Process Engine installation program that aliases will be defined. When prompted, provide the <i>f_maint</i> name to the installation program.</p> <p>After creating the new users and setting their group memberships, log off as root user, log on as each of the new users, and change the password to avoid connection problems the first time they are used.</p> <p>In a farmed configuration, ensure that each Process Engine has the same run time user name.</p>
The password for the database maintenance user for Process Engine(DB2 for Linux, UNIX and Windows): <i>f_maint_password</i>	<p>This is the password for the <i>f_maint</i> user.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>The Process Engine installation program prompts for this password.</p> <p>The password entered during installation will be used to create a file containing an encrypted version of this password. If the encrypted version of the password does not match the operating system version of the password, Process Engine will fail to connect to the database. If the <i>f_maint_password</i> field is left blank during installation, the value will be set to <code>change\$this_obnoxious_password</code>, which is the default.</p>

User or group name	Description
The database instance owner for Process Engine(DB2 for Linux, UNIX and Windows and DB2 for z/OS®): <i>instance_owner</i>	<p>This is the instance owner for the Process Engine DB2 database instance. The user will create databases and set a number of configuration parameters.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>The instance owner is identified when the DB2 database instance is created. It is also identified to the Process Engine installation program.</p> <p>If the database is remote, this is the <i>instance_owner</i> for the client instance on the Process Engine server.</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)”

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)” on page 19

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 19

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 19

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

“Configuring Process Engine servers (AIX)” on page 21

The following operating system prerequisites apply to AIX IBM FileNet P8 Process Engine servers.

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 21 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 21

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.

(AIX only): For configurations with workload partitioning (WPAR), create a system non-shared WPAR to host the Process Engine using the **mkwpar** command.

(AIX only): Permissions must be set correctly on the /fnsw and /fnsw/local mount points and the file systems before and after mounting the file systems. In a WPAR configuration, set these permissions from within the WPAR.

To configure UNIX servers:

- Ensure minimum /tmp size. The /tmp directory must have 510 MB free.
- (AIX only): In a WPAR configuration, the /wpars/wpar-name/tmp file system is automatically created when a system non-shared WPAR is created with the **mkwpar** command. The base directory in the global environment is /wpars. Verify that /wpars/wpar-name/tmp has 510 MB free. Increase the size, if necessary, in the global environment.
- 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

- (AIX only): In a WPAR, create and mount these partitions on the WPAR base directory in the global environment. It is recommended that these file systems be created by specifying a mount group with the name of the WPAR to make them permanent.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fnsw	/wpars/wpar-name/fnsw	2 GB	fnsw	fnusr	775
local	/wpars/wpar-name/fnsw/local	10 GB	fnsw	fnusr	775

The following syntax shows how to create the file systems where *p8pewpar* is the name of the WPAR. The *-u* flag indicates the mount group.

```
crfs -v jfs2 -g datavg -m /wpars/p8pewpar/fnsw -u p8pewpar -a
logname=INLINE -a size=2G
```

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 (AIX):

The following operating system prerequisites apply to AIX IBM FileNet P8 Process Engine servers.

“Configuring AIX servers for Process Engine”

Perform the following prerequisite tasks on the AIX server on which Process Engine will be installed. You can perform these tasks in any order.

“Modifying /etc/rc.dt for AIX 5.3 and 6.1”

Configure network ports by modifying /etc/rc.dt. In a configuration with workload partitioning (WPAR), make these configuration changes in the global environment.

Configuring AIX servers for Process Engine:

Perform the following prerequisite tasks on the AIX server on which Process Engine will be installed. You can perform these tasks in any order.

To configure AIX servers for Process Engine:

- Set the kernel to 64-bit mode.
- Set the swap space to 1.5 - 2 times RAM.
- Set the Maximum Number of Processes allowed per user to at least 400.
- Set the Maximum KB of real memory allowed for **MBUFS** to 0. Setting the **MBUFS** parameter to 0 causes the system to use the default amount of available memory. This default amount is approximately 1/8 to 1/4 the amount of real memory.
- Set the Maximum Number of FIXED licenses (Num) to a minimum of 16.
- Install and commit the following filesets:
 - bos.adt.libm
 - bos.adt.lib
 - bos.adt.base
 - bos.perf.perfstat
 - bos.perf.libperfstat
 - bos.adt.debug
- Review and change the time zone parameters if necessary. In SMIT, choose **System Environments** → **Change/Show Date and Time** → **Change Time Zone Using System Defined Values**. Choose the Daylight Savings Time option if applicable. At the CUT Time Zone menu, choose the option associated with your site. For example, in California, the time zone needs to be set to the **Pacific time zone (PST8PDT) Pacific U.S.; Yukon (cut -8)**.

Modifying /etc/rc.dt for AIX 5.3 and 6.1:

Configure network ports by modifying /etc/rc.dt. In a configuration with workload partitioning (WPAR), make these configuration changes in the global environment.

To modify the network ports:

1. As the root user, run the following commands:

```
/usr/sbin/no -p -o tcp_sendspace=16384
/usr/sbin/no -p -o tcp_recvspace=16384
/usr/sbin/no -p -o tcp_keepidle=80
/usr/sbin/no -p -o tcp_keepintvl=20
/usr/sbin/no -p -o tcp_ephemeral_high=65535
/usr/sbin/no -p -o tcp_ephemeral_low=42767
/usr/sbin/no -p -o udp_ephemeral_high=65535
/usr/sbin/no -p -o udp_ephemeral_low=42767
```

This should result in these entries in the /etc/tunables/nextboot file:

```
no:
tcp_ephemeral_high = "65535"
tcp_ephemeral_low = "42767"
tcp_keepidle = "80"
tcp_keepintvl = "20"
tcp_recvspace = "16384"
tcp_sendspace = "16384"
udp_ephemeral_high = "65535"
udp_ephemeral_low = "42767"
```

2. If the CDE bundle is installed, the dtlogin daemon may bind to a required port before the kernel parameters listed above take effect. In order to avoid this problem, add the following statements at the beginning of the /etc/rc.dt file:

```
/usr/sbin/no -o tcp_sendspace=16384
/usr/sbin/no -o tcp_recvspace=16384
/usr/sbin/no -o tcp_keepidle=80
/usr/sbin/no -o tcp_keepintvl=20
/usr/sbin/no -o tcp_ephemeral_high=65535
/usr/sbin/no -o tcp_ephemeral_low=42767
/usr/sbin/no -o udp_ephemeral_high=65535
/usr/sbin/no -o udp_ephemeral_low=42767
```

In some circumstances you may need to create an /etc/rc.dt file containing the commands above even if the CDE bundle is not installed. This is indicated by an error message when trying to start Process Engine similar to the following, where *myserver* is the name of the machine running Process Engine:

```
Retrying to connect to host 'myserver'...
initfns: connect() failed with error 'connection refused'!
Is TM_daemon running?
```

3. Restart the server (shutdown -Fr) for these settings to take effect. Executing the commands at the command line is not sufficient. The parameters must be in effect at boot time to avoid bind failures.
4. Check the values by executing the following commands:

```
no -a | grep ephemeral
no -a | grep tcp
```

5. Check that the ports required by the Process Engine software have not been bound. The following command, entered as a single command line, should not produce any output:

```
for port in 32768 32769 32770 32771 32772 32773 32774
do
    netstat -an | grep $port
done
```

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 32

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 32

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:

- IBM Tivoli 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.

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.

(WebSphere only) Choose Stand-alone or Federated repository type.

There is an option in Configuration Manager's Configure LDAP task to select whether the WebSphere Application Server repository type is Stand-alone LDAP registry or Federated repositories. In order to have Configuration Manager use your repository type setting, you must select the Configuration Manager option to **Set as current active user registry**.

If you choose Stand-alone LDAP registry

Configuration Manager changes the administrative console user login to the new account you enter as the **Administrative console user name** (*ce_appserver_console_name*). This new account must reside in the Stand-alone LDAP registry location. The existing administrative console user login, if any, becomes invalid.

In order to have Configuration Manager replace an existing Stand-alone LDAP registry configuration, you must enable the Configuration Manager option **Overwrite existing repository**.

If you choose Federated repositories

By choosing the Federated repositories option in Configuration Manager, you are adding a new LDAP realm to an existing Federated LDAP repository. The administrative console user name (*ce_appserver_console_name*) you provide must be a unique user across all federated realms.

Avoid overlapping realm definitions

In the Configuration Manager task Configure LDAP, if you set the Websphere Application Server LDAP repository type option to **Federated repositories**, do not enter repositories with overlapping suffixes as they are not supported. For example, the following two repositories with overlapping Base entry distinguished names are not supported:

- dc=ibm,dc=com

- dc=filenet,dc=ibm,dc=com

This restriction especially applies to Active Directory parent and child domains, since by definition parent/child domains in AD have overlapping suffixes.

The repositories in the next example are supported, because they are sibling repositories and do not overlap:

- dc=tivoli,dc=ibm,dc=com
- dc=filenet,dc=ibm,dc=com

Administrative security must be enabled

Configuration Manager does not change the state of WebSphere administrative security. If it was on before running Configuration Manager, then it stays on; if it was off before, then it stays off.

(Note, however, that Process Engine administrative security must be enabled before starting Content Engine. The *IBM FileNet P8 Platform Installation and Upgrade Guide* contains steps for enabling WebSphere's administrative security.)

For more information on federated LDAP repositories, consult the IBM WebSphere Application Server information center, and search for the keywords "federated repositories".

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 IBM Tivoli Directory Server”

You can configure IBM Tivoli® Directory Server to be the directory service for IBM FileNet P8 Platform.

Configuring IBM Tivoli Directory Server:

You can configure IBM Tivoli Directory Server to be the directory service for IBM FileNet P8 Platform.

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 IBM Tivoli Directory Server features, refer to the IBM FileNet P8 Help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory Service Providers** → **Tivoli 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.”

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>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>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>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>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>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: IBM Tivoli Directory Server (Referred to as Directory service bind user name in Configuration Manager): <i>ce_service_user</i>	<p>An IBM Tivoli Directory Server user account that Content Engine uses to connect to IBM Tivoli 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 Workplace object after a user has been authenticated. <p>Using IBM Tivoli 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 <i>ce_service_user</i> 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*.

WebSphere Application Server

*WAS_HOME/profiles/default/installedApps/node_name/
app_engine_war.ear/app_engine.war*

*WAS_HOME/profiles/default/config/cells/machine_name/Node01cell/
nodes/machine_name/Node01/serverindex.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 59 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. After the IT Administrator creates operating system users and groups for DB2 databases, you must grant database permissions.

“Creating Process Engine database accounts”

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

“Preparing DB2 for Linux, UNIX and Windows servers” on page 42

Plan and prepare your IBM DB2 for Linux, UNIX and Windows servers for IBM FileNet P8 installation.

“Configuring DB2 client” on page 49

Configure the DB2 client. Configuration requirements will be different for Content Engine and Process Engine and will vary between DB2 for Linux, UNIX and Windows and DB2 for z/OS. In a configuration with workload partitioning (WPAR), perform all of the DB2 client configuration from within the WPAR.

“Verifying the ability to connect to the database” on page 50

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. After the IT Administrator creates operating system users and groups for DB2 databases, you must grant database permissions.

Create the following users and groups:

User or group name	Description
Database user name (DB2 for Linux, UNIX and Windows): <i>ce_db_user</i>	Grant the following database permissions to this user: <ul style="list-style-type: none">• Connect to the database• Create tables (CREATETAB) in the tablespace• Use the tablespace (USE OF) for User and User Temp tablespaces• SELECT on SYSIBM.SYSTABLES• Usage on workload SYSDEFAULTUSERWORKLOAD• Implicit_schema on database• SELECT on SYSIBM.SYSCOLUMNS (GCD creation only)• SELECT on SYSIBM.SYSVERSIONS, SYSCAT.DATATYPES (Object Store creation only)

Creating Process Engine database accounts

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

After the IT Administrator creates operating system users and groups for DB2 databases, you must grant database permissions.

Create the following users and groups:

User or group name	Description
The database runtime user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_sw</i>	Grant the following database permissions to this user if you are using SERVER and SERVER_ENCRYPT authentication: <ul style="list-style-type: none"> • permission to connect • permission to create tables in the tablespaces (CREATETAB) • permission to use the tablespaces (USE OF) Grant SYSADM permission to this user if you are using CLIENT authentication.
The password for the database runtime user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_sw_password</i>	This is the password for the <i>f_sw</i> user. The Process Engine installation program prompts for this password. The password entered during installation will be used to create a file containing an encrypted version of this password. If the encrypted version of the password does not match the operating system version of the password, Process Engine will be unable to connect to the database. If the <i>f_sw_password</i> field is left blank during installation, the value will be set to <i>filenet</i> , which is the default.

User or group name	Description
The database maintenance user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_maint</i>	<p>This user is the primary Process Engine DB2 for Linux, UNIX and Windows database maintenance user and is used only by the Process Engine software to access the DB2 database.</p> <p>For a database using SERVER or SERVER_ENCRYPT authentication for the Process Engine database, this user:</p> <ul style="list-style-type: none"> • Is recommended to be member of group having SYSADM authority of the DB2 instance that will be used by the Process Engine software • Does not need to have separate file systems for a home directory • Must belong to the primary group of the instance owner <p>Grant SYSADM permission to this user if you are using CLIENT authentication.</p> <p>After creating the new users and setting their group memberships, log off as root user, log on as each of the new users, and change the password to avoid connection problems the first time they're used.</p> <p>In a farmed configuration, ensure that each Process Engine has the same run time user name.</p>
The password for the database maintenance user for Process Engine (DB2 for Linux, UNIX and Windows): <i>f_maint_password</i>	<p>This is the password for the <i>f_maint</i> user.</p> <p>The Process Engine installation program prompts for this password.</p> <p>The password entered during installation will be used to create a file containing an encrypted version of this password. If the encrypted version of the password does not match the operating system version of the password, Process Engine will be unable to connect to the database. If the <i>f_maint_password</i> field is left blank during installation, the value will be set to <code>change\$this_obnoxious_password</code>, which is the default.</p>
The database instance owner for Process Engine (DB2 for Linux, UNIX and Windows and DB2 for z/OS): <i>instance_owner</i>	<p>This is the instance owner for the Process Engine DB2 for Linux, UNIX and Windows database instance.</p> <p>The instance owner is identified when the DB2 database instance is created. It is also identified to the Process Engine installation program.</p>

Preparing DB2 for Linux, UNIX and Windows servers

Plan and prepare your IBM DB2 for Linux, UNIX and Windows servers for IBM FileNet P8 installation.

The following are planning considerations for DB2 for Linux, UNIX and Windows:

- Determine whether DB2 server instances will be 64-bit or 32-bit. Use the version that is appropriate to the operating system where the server instance is installed.
- IBM FileNet P8 does not support partitioned DB2 databases or databases created with the RESTRICTIVE clause (or, from the Control Center, with **Restrict access to system catalogs** selected in the Create Database Wizard).
- Determine whether you want to use a dedicated or shared database. In this regard:
 - Content Engine and Process Engine can share a database engine, or they can each have a dedicated (unique) database engine.
 - Content Engine and Process Engine can each have a dedicated DB2 instance, or they can share an instance with one another or with non-IBM FileNet P8 applications.
 - Assign unique databases to the Content Engine global configuration data (GCD) and object stores. Although you can configure multiple Content Engine object stores in a single DB2 database, you should configure each object store in a separate database. If you configure object stores in separate databases, you have more flexibility and control with security access, backup scheduling and execution, updates, and scheduled outages.
- Plan to use automatic storage for tablespaces. For performance reasons, IBM recommends that you create tablespaces using automatic storage, rather than database managed or system managed tablespaces for Process Engine and Content Engine.
- Set DB2 collation. For Content Engine, use UTF-8. Process Engine will support the UTF-8 code set, code page 1208. Process Engine will also support all other single-byte character sets, for example:

Code Set	Code Page #
ISO8859-15	923
ISO8859-1	819

- It is a best practice to use SERVER authentication. SERVER_ENCRYPT and CLIENT authentication are also supported.
- Determine the maximum size of the content elements your users store. This affects setting up database storage areas or file storage areas. When you create an object store, a database storage area is provided by default, allowing you to store content as database BLOBs. You can also create one or more file storage areas to store content on local or remote file systems. If your users store large individual documents or other content elements, use only file storage areas. Otherwise, users can encounter memory-related errors when retrieving or indexing the large content.

Important: Controlled tests with limited concurrency exhibited errors when run with files that were 300 MB or larger. Factors affecting this file-size limitation include driver and application server memory demands, other activity such as concurrent retrieval or indexing of large content, and JVM memory allocations.

“Verifying that DB2 for Linux, UNIX and Windows is installed for IBM FileNet P8” on page 44

You can install and configure a DB2 instance that is dedicated or shared by one or more IBM FileNet P8 components. You can also share the instance with other (non-IBM FileNet P8) applications.

Verifying that DB2 for Linux, UNIX and Windows is installed for IBM FileNet P8:

You can install and configure a DB2 instance that is dedicated or shared by one or more IBM FileNet P8 components. You can also share the instance with other (non-IBM FileNet P8) applications.

- Dedicated to Content Engine
- Dedicated to Process Engine
- Shared by two or more Content Engines or Process Engines
- In a shared configuration, the IBM FileNet P8 components use the same instance, but different databases.
- Content Engine supports multiple object stores within the same DB2 database provided the users are properly partitioned via tablespace authorization.
- Content Engine and Process Engine should not share tablespaces.

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

When this procedure has been completed, proceed to “Configuring DB2 client” on page 49.

Important: Record the values for the following settings as you work through the database installation. Enter this information in the appropriate sections of the Installation and Upgrade Worksheet. This information must be entered during subsequent installations of Process Engine and Content Engine. Be aware that the Process Engine installation program allows only alphanumeric characters and underscores.

- DB2 Server name
- DB2 server database instance names (for example, P8inst)
- Content Engine and Process Engine dedicated database names (for example, VWdb)
- Dedicated tablespace names (for example, vwdata)
- DB2 instance port numbers
- Process Engine runtime database user (f_sw or alias) password
- Process Engine maintenance database user (f_maint or alias) password
- User ID and password for Content Engine DB2 user

“Installing DB2 for Linux, UNIX and Windows and creating DB2 instances” on page 45

Create DB2 for Linux, UNIX and Windows instances for Content Engine and Process Engine and set several instance values.

“Determining page size and user fields” on page 45

When you create a DB2 database, you have a choice of several page sizes for your tablespace: 4 KB, 8 KB, 16 KB, and 32 KB. The page size you choose affects the number and size of the user-defined index fields, such as property templates, the tablespace can support and the maximum row length of the tables within that tablespace.

“Creating and updating the DB2 databases for Content Engine and Process Engine” on page 45

To create and update the DB2 databases, log on as the database instance owner and ensure the following steps are taken.

“Creating the DB2 tablespaces for Process Engine” on page 46
Create five tablespaces for Process Engine on DB2 for Linux, UNIX and Windows. The default data, index and BLOB tablespaces will be prompted for by the Process Engine installation program.

“Creating the DB2 tablespaces for Content Engine” on page 48
Create four tablespaces for Content Engine on DB2 for Linux, UNIX and Windows. Each additional object store will require an additional tablespace and a unique tablespace user.

Installing DB2 for Linux, UNIX and Windows and creating DB2 instances:

Create DB2 for Linux, UNIX and Windows instances for Content Engine and Process Engine and set several instance values.

To install DB2 for Linux, UNIX and Windows and create DB2 instances:

1. Install the IBM DB2 for Linux, UNIX and Windows software. Make note of the TCP/IP port number assigned to the instance or instances, as the port number will be needed during the DB2 client configuration steps. The port number assigned can be found in the `/etc/services` file, associated with the DB2 instance(s) just created. After a successful installation, the DB2 instance should be up and running.
2. Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs.
3. Content Engine and Process Engine can share an instance, or each engine can have its own instance. Create the appropriate instances if they do not exist.
4. Set TCP/IP as the default protocol.
5. (Content Engine) Set the `DB2_OPTPROFILE` variable to YES.
6. (DB2 version 9.7 only) Set the `DB2_WORKLOAD` variable to `FILENET_CM`.

Determining page size and user fields:

When you create a DB2 database, you have a choice of several page sizes for your tablespace: 4 KB, 8 KB, 16 KB, and 32 KB. The page size you choose affects the number and size of the user-defined index fields, such as property templates, the tablespace can support and the maximum row length of the tables within that tablespace.

The important things to remember are:

- The total row length of all the fields (including system and user) cannot be larger than the page size.
- Choose a minimal set of add-ons before creating an object store. Each add-on reserves space in a row whether the add-on is used or not.
- The DB2 page size you select when you create the database must be large enough to hold at least one complete record.
- DB2 cannot retrieve a partial record or spread a single record onto two pages.
- It is a best practice to assign 8 KB minimum page sizes for Process Engine databases.
- 32 KB minimum page sizes are required for Content Engine databases.

Creating and updating the DB2 databases for Content Engine and Process Engine:

To create and update the DB2 databases, log on as the database instance owner and ensure the following steps are taken.

To create and update the DB2 database:

1. Log on as the database instance owner as defined earlier. At least three databases must be created, one for Process Engine, one for the Content Engine GCD, and one for a single Content Engine object store.

The database name needs to be unique and from 1 to 8 characters long.

2. For a database to be used by Content Engine object stores, update the following configuration parameter. Set the value, minimally, to the value indicated here:

APPLHEAPSZ 2560

3. (DB2 version 9.7 only) Set the CUR_COMMIT variable to ON.
4. For a database to be used by Content Engine, drop the default user [regular] tablespace - USERSPACE1 after creating the database.

Creating the DB2 tablespaces for Process Engine:

Create five tablespaces for Process Engine on DB2 for Linux, UNIX and Windows. The default data, index and BLOB tablespaces will be prompted for by the Process Engine installation program.

If you will be using the Process Engine region recovery feature, each Process Engine isolated region configured for region recovery must reside in dedicated data, index and BLOB tablespaces. If you are creating a tablespace for a new object store on an existing system, define the new tablespace with the same tablespace type and storage method used for existing object store tablespaces.

Process Engine Tablespaces	Actual Assigned Name	Record the values in the worksheet.	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
user temporary ts			40		must match the pagesize of VWDATA_TS
system temporary			40		must match the pagesize of VWDATA_TS
VWDATA_TS (default data tablespace) Configure as LARGE type and automatic storage.			200		8
VWINDEX_TS (default index tablespace) Configure as LARGE type and automatic storage.			200		must match the pagesize of VWDATA_TS

Process Engine Tablespaces	Actual Assigned Name	Record the values in the worksheet.	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
	<i>VWBLOB_TS</i>		200		8
	<p>(default blob tablespace)</p> <p>Configure as LARGE type and automatic storage.</p>				
	<i>region X data</i>		200		8
	<p>Data tablespace to be used by an individual PE region configured for recovery.</p> <p>Cannot be shared by any other region.</p> <p>Configure as LARGE type and automatic storage.</p>				
	<i>region X index</i>		200		must match the pagesize of <i>VWDATA_TS</i>
	<p>Index tablespace to be used by an individual PE region configured for recovery.</p> <p>Cannot be shared by any other region and must be separate from the default index tablespace. Configure as LARGE type and automatic storage.</p>				

Process Engine Tablespaces	Actual Assigned Name	Record the values in the worksheet.	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
	<i>region X blob</i>		200		8
	<p>BLOB tablespace to be used by an individual PE region configured for recovery.</p> <p>Cannot be shared by any other region and must be separate from the default index tablespace.</p> <p>Configure as LARGE type and automatic storage.</p>				

Creating the DB2 tablespaces for Content Engine:

Create four tablespaces for Content Engine on DB2 for Linux, UNIX and Windows. Each additional object store will require an additional tablespace and a unique tablespace user.

If you are creating a tablespace for a new object store on an existing system, define the new tablespace with the same tablespace type and storage method used for existing object store tablespaces.

Content Engine Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
	GCD_ts	256		32 (required)
	<p>(for the GCD database)</p> <p>Configure as LARGE type and automatic storage.</p>			
	cedata_ts	512		32 (required)
	<p>(for a single CE object store)</p> <p>Configure as LARGE type and automatic storage.</p>			
	user temporary ts	40		32 (required)

Content Engine Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
system temporary	ts	40		32 (required)

Configuring DB2 client

Configure the DB2 client. Configuration requirements will be different for Content Engine and Process Engine and will vary between DB2 for Linux, UNIX and Windows and DB2 for z/OS. In a configuration with workload partitioning (WPAR), perform all of the DB2 client configuration from within the WPAR.

“Installing DB2 client software (Process Engine)”

Install the DB2 Administration Client or the DB2 Runtime Client.

“Creating DB2 client instances for Process Engine”

Depending upon your configuration, create at least one DB2 Client instance for Process Engine. The client instance is referred to as the local instance in the Process Engine installation program.

“Cataloging the Process Engine DB2 server node”

Catalog the Process Engine server node for the DB2 database.

“Creating the Process Engine DB2 database alias”

Create an alias for the Process Engine DB2 for Linux, UNIX and Windows or DB2 for z/OS database.

Installing DB2 client software (Process Engine):

Install the DB2 Administration Client or the DB2 Runtime Client.

Content Engine does not require DB2 client software.

Creating DB2 client instances for Process Engine:

Depending upon your configuration, create at least one DB2 Client instance for Process Engine. The client instance is referred to as the local instance in the Process Engine installation program.

On UNIX, create a DB2 client instance for Process Engine.

Cataloging the Process Engine DB2 server node:

Catalog the Process Engine server node for the DB2 database.

1. Reboot the server and log on as the instance owner on the Process Engine.
2. Use the db2ca tool, or catalog the DB2 server node as follows:

```
db2 catalog tcpip node server_alias remote server_name
server server_side_instance_tcpip_port_#
```

For example: db2 catalog tcpip node sampnod remote hqwais20 server 60004

Creating the Process Engine DB2 database alias:

Create an alias for the Process Engine DB2 for Linux, UNIX and Windows or DB2 for z/OS database.

Use the db2ca tool, or create an alias for the Process Engine DB2 database. Record the alias name in the Installation and Upgrade Worksheet.

For example: db2 catalog database PEDBAIX at node aix20nod [as *alias name*]

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 (DB2)”

Verify the connection to the DB2 Process Engine database by executing the following commands to start a command line processor.

Verifying the Process Engine database connection (DB2):

Verify the connection to the DB2 Process Engine database by executing the following commands to start a command line processor.

1. Log on to the DB2 Control Center tool on the Process Engine, as follows:

UNIX Log on as the client instance owner and run the following at a command prompt:

```
db2
```

2. At the DB2 prompt, enter the following command indicating either the database alias or the instance name:

```
connect to database_alias or instance_name user f_sw using  
f_sw password
```

where:

database_alias is the Process Engine DB2 for Linux, UNIX and Windows database alias for remote databases. For local databases, use the database name.

instance_name is the Process Engine DB2 for z/OS database name

f_sw is the Process Engine runtime user, either the default *f_sw* user or the assigned alias

f_sw password is the runtime user's password.

DB2 will display the database connection information.

The following example shows the database connection command and the information returned:

```
db2 => connect to pedbinst user f_sw using fswpassword
```

```
Database Connection Information
```

```
Database server = DB2/AIX64 9.5.0
```

```
SQL authorization ID = F_SW
```

```
Local database alias = PEDBINST
```

In this example, the database alias is pedbinst, the user is *f_sw*, and the *f_sw* user password is fswpassword.

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 59 before you begin your preparation tasks.
- Review the following planning considerations:
 - “Application server planning considerations” on page 4
 - “Creating Content Engine application server accounts”
Create new or designate existing application server accounts for Content Engine, as shown in the following table.
 - “Configuring WebSphere for Content Engine” on page 53
You must prepare IBM WebSphere Application Server before you install Content Engine. You must create a WebSphere profile for the Content Engine application and set the environment variables for the database connection.
 - “Configuring WebSphere Application Server for Application Engine or Workplace XT” on page 57
You must install WebSphere Application 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 57
Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).
 - “Configuring the documentation server” on page 57
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>	<p>WebSphere Application Server:</p> <p>There are several things to consider when creating the <i>ce_appserver_admin</i> account for WebSphere Application Server. When you use WebSphere Application Server tools to create the WebSphere Application Server profile that will contain Content Engine, at the same time you also create an administrative account (<i>ws_admin</i>) and decide whether to enable security.</p> <p>If you do not enable security, then you can login to WebSphere Application Server using any account. Later on when you run Configuration Manager, the Set Properties for WebSphere Application Server configuration step creates the <i>ce_appserver_admin</i> account.</p> <p>If you enable security, then the administrative account (<i>ws_admin</i>) you created while creating the profile must be the account you use later on as <i>ce_appserver_admin</i>. This account is a WebSphere Application Server Local account. However, if you configure LDAP (either manually using the WebSphere Application Server console or later on using the Configure LDAP task in Configuration Manager) in the application server using Federated Repository and if this local <i>ce_appserver_admin</i> account is also in an LDAP directory that is configured for Content Engine authentication, you will lock yourself out of WebSphere Application Server. Therefore, the <i>ws_admin</i> account must be unique and must not be used by any other user that is in an LDAP directory that will be configured for Content Engine authentication.</p> <p>If you enable security but plan on using WebSphere Application Server Stand Alone mode, the <i>ws_admin</i> account can also be in the configured directory server, as long as there is not a WebSphere Application Server local account with the same name.</p>

User/Group	Description
Application server administrator user name (continued) <i>ce_appserver_admin</i>	Description of the WebSphere Application Server administrative account: <ul style="list-style-type: none"> This administrative account can be the same as the Content Engine system user (<i>ce_bootstrap_admin</i>), if you enter the same credentials for both in Configuration Manager, and provided that you are using an LDAP account, and also that the account will be a P8 Domain Administrator (<i>gcd_admin</i>) specified during GCD creation. Minimum required permissions for <i>ce_appserver_admin</i>: Full administrative control over the WebSphere Application Server domain that contains Content Engine. If your site uses Federated Repository, <i>ce_appserver_console_admin</i> must be a unique user across all federated realms including the WebSphere Application Server Local repository. After Content Engine installation is complete, you can reduce the account's permissions to a lesser role within WebSphere Application Server, such as Configurator.
Administrative Console user (continued) <i>ce_appserver_console_admin</i>	
	You must use the <i>ce_appserver_admin</i> account to perform the following tasks: <ul style="list-style-type: none"> Create directory service providers within the application server. Create JDBC providers and data sources for database connectivity. Deploy the Content Engine application. Stop and restart servers and cluster members (WebSphere Application Server). When you enable WebSphere Application Server Global Security as a post-install step, you must use <i>ce_appserver_console_admin</i> to login to the WebSphere Application Server console and run Configuration Manager and any scripts.

Configuring WebSphere for Content Engine

You must prepare IBM WebSphere Application Server before you install Content Engine. You must create a WebSphere profile for the Content Engine application and set the environment variables for the database connection.

1. "Creating the WebSphere profile for Content Engine" on page 54
You must create an IBM WebSphere Application Server profile for Content Engine if you do not already have a profile.
2. "Specifying the WebSphere environment variables" on page 54
You must specify the IBM WebSphere Application Server environment variables so that Content Engine can access its databases.

3. "Setting the primary administrative user name" on page 56
If you are using IBM WebSphere Application Server federated repositories for LDAP authentication, you must ensure that the name you entered for the WebSphere Application Server primary administrative user name is unique across all realms.
4. "Setting host aliases for deployment on multiple servers" on page 56
If you are deploying Content Engine to multiple IBM WebSphere Application Server servers on the same WebSphere node, you must define the host alias and port numbers.
5. "Setting permissions for the Configuration Manager user" on page 56
You must create an application server account with certain directory permissions that you will use later on to start the Configuration Manager tool.

Creating the WebSphere profile for Content Engine:

You must create an IBM WebSphere Application Server profile for Content Engine if you do not already have a profile.

To create the WebSphere profile for Content Engine:

1. Run the command script at one of the following (default) locations to create a new profile.

Option	Description
AIX	/usr/IBM/WebSphere/AppServer/manageprofiles.sh

2. Record the path to your new profile in the **Application server installation directory** property in your installation worksheet, as you must specify the path later when you configure Content Engine.
3. Grant write permission to the group *ce_appserver_install_group* (the user who runs Configuration Manager belongs to this group) on the following files in the logs directory of the WebSphere profile for Content Engine:
 - wsadmin.traceout
 - wsadmin.valout

You can find these files in one of the following locations:

Option	Description
AIX	/usr/IBM/WebSphere/AppServer/profiles/ <i>profile_name</i> /logs

where *profile_name* is the name of the WebSphere profile (for example, AppServer01).

Specifying the WebSphere environment variables:

You must specify the IBM WebSphere Application Server environment variables so that Content Engine can access its databases.

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for information on the JDBC driver file for the database type that you need for the Global Configuration Data (GCD) or for an object store you will be creating later.

To specify the WebSphere environment variables:

1. Install JDBC drivers on the server where WebSphere Application Server is installed.
 - a. Obtain the JDBC drivers for your database type.

DB2 Find the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the IBM Web site by searching for "JDBC Type 4".
 - b. Copy the JDBC driver file to the following suggested location:

UNIX /opt/jars

Do not copy the file to ...WebSphere/AppServer/lib/ext.
2. Start the WebSphere Application Server administrative console and log on to your Content Engine profile as *ce_appserver_console_admin*, the Administrator Console User.
3. Navigate to **Environment** → **WebSphere Variables** and specify the JDBC driver path:
 - a. Select **Cell scope** from the **All scopes** list.
 - b. Click **New** to create a WebSphere Application Server variable whose name is one of the JDBC environment variables shown in the Database environment variables table below.
 - c. Set the value of the variable to the JDBC driver path that you specified when you installed the JDBC drivers on the IBM WebSphere Application Server machine.
 - d. Save your change to the master configuration.
 - e. Select **Node scope** from the **All scopes** list.
 - f. In the table of substitution variables, click the *item name* in the **Name** column that corresponds to the JDBC environment variable for your database type in the Database environment variables table below.

Table 1. Database environment variables

Database	JDBC Environment Variable
DB2	DB2UNIVERSAL_JDBC_DRIVER_PATH

- g. Set the value of the *name_JDBC_DRIVER_PATH* item to the JDBC driver path you specified (/opt/jars or C:\jars).
 - h. Save your changes to the master configuration.
4. Navigate to **Servers** → **Application servers** → *server1* → **Java and Process Management** → **Process Definition** → **Java Virtual Machine**, and set the initial and maximum heap sizes, where *server1* is the name of the server where you will deploy Content Engine.
 - a. Set the values for the initial and maximum heap sizes:

Parameter	Value (in MB)
Initial Heap Size	At least 512
Maximum Heap Size	1024 or a size consistent with available RAM on the machine where WebSphere Application Server is installed

- b. Save your changes to the master configuration.
5. Optional: Increase the maximum transaction timeout to prevent administrative processes from failing:

- a. Navigate to the screen containing the **Maximum transaction timeout** parameter:
 - (WebSphere 6.1) Click **Servers** → **Application servers** → **server1** → **[Container Settings] Container Services** → **Transaction Service**.
 - (WebSphere 7.0) Click **Servers** → **Server Types** → **WebSphere application servers** → **server1** → **[Container Settings] Container Services** → **Transaction Service**.
 - b. Click the **Configuration** tab, and set the **Maximum transaction timeout** parameter value to at least 600 (seconds).

Important: If the timeout value is not large enough, some administrative processes (such as adding an expansion product) might fail.
 - c. Click **Apply** and then click **Save**.
6. Repeat this procedure as needed for any object store that uses a different database type.

Setting the primary administrative user name:

If you are using IBM WebSphere Application Server federated repositories for LDAP authentication, you must ensure that the name you entered for the WebSphere Application Server primary administrative user name is unique across all realms.

Setting host aliases for deployment on multiple servers:

If you are deploying Content Engine to multiple IBM WebSphere Application Server servers on the same WebSphere node, you must define the host alias and port numbers.

To set the host alias:

1. Log in to the WebSphere administrative console.
2. Navigate to **Environment** → **Virtual Hosts** → **default host** → **Host Aliases**.
3. If you are using SSL, add an alias for the SSL port number, such as port 9081.
4. Add an alias for the non-SSL port number, such as port 9444.
5. Click **Apply**.

Setting permissions for the Configuration Manager user:

You must create an application server account with certain directory permissions that you will use later on to start the Configuration Manager tool.

To set permissions for the Configuration Manager user:

Set permissions for the Configuration Manager user (*config_mgr_user*) on the WebSphere Application Server profile directory and all its subdirectories where Content Engine will be deployed:

Option	Description
UNIX	Read, write, and execute permissions

Configuring WebSphere Application Server for Application Engine or Workplace XT

You must install WebSphere Application 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 the 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 WebSphere Application 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 to pass session information between Application Engine or Workplace XT and the client browser.
2. Determine the Initial and Maximum Heap Size. Refer to your application server vendor's recommendation for Initial and Maximum heap size values. You will use this information when you configure WebSphere Application Server after you install Application Engine or Workplace XT. For IBM specific recommendations, see the *IBM FileNet P8 Performance Tuning Guide*.
3. When WebSphere is running as a service and a UNC path is specified in `web.xml` for configuration, upload, and download directories, the account that is specified to run the WebSphere service must have permissions to the share of the UNC path.

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
IBM WebSphere Application Server with the IBM JVM	No changes are required.

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 2. 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”

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

“Database administrator” on page 60

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

“IT administrator” on page 60

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 62

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 63

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 on WebSphere Application Server”

IBM FileNet P8 requires the following character encoding settings.

Configuring character encoding on WebSphere Application Server

IBM FileNet P8 requires the following character encoding settings.

1. Set the com.ibm.CORBA.ORBCharEncoding property to 0x05010001.
2. Set the com.ibm.websphere.security.BasicAuthEncoding property to UTF-8.

Security administrator

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

“Extended characters and user names” on page 60

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 the DB2 server”

Single instances of Content Engine and Process Engine on a DB2 server can support multiple languages.

Installing the DB2 server

Single instances of Content Engine and Process Engine on a DB2 server can support multiple languages.

Content Engine

A single instance of Content Engine on a DB2 server can support multiple languages using Windows language packs.

Process Engine (UNIX)

To make a single instance of Process Engine in a UNIX environment support multiple languages, select the UTF-8 database character set when configuring the DB2 database. If support for more than one language is not required, select the appropriate language database character set instead.

The following example shows a script line that could be used when creating a Process Engine database on AIX for multilingual support. The example illustrates a Unicode database code set with Japanese territory and collation support.

```
db2 create db VWDB using codeset UTF-8 territory JA_JP collate using system;
```

For more information, see the planning section in the DB2 administration guide.

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” on page 61

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”

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” on page 62

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” on page 62

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” on page 63

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” on page 63

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.

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	
WebSphere EJB	2809	This port is on the WebSphere Application Server for Content Engine, for communication with Content Engine by clients through EJB.
WebSphere WSI	9080	This port is on the WebSphere Application Server for Content Engine, for communication with Content Engine by clients through WSI.
DB2	50000	This port is on the DB2 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.
DB2	50000	This port is on the database server, for connection between Process Engine and DB2.
TMS	32768 (TCP)	Task Manager service. TM_daemon listens for requests from initfnsw running on the same or a different system analogous to COR_Listen listening for RPCs.

Port name	Default number	Notes
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.
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.
SNMPD SMUX (AIX only)	199 (TCP)	This port is on the Process Engine server.
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
WebSphere	9080	This port is on the WebSphere Application Server for clients to connect to Application Engine.
WebSphere SSL	443	This port is on the WebSphere Application Server for clients to connect to Application Engine through SSL.
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.

Port name	Default number	Notes
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.
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 Lipient

Port name	Default number	Notes
Lipient 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.

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 33
 - Content Engine application server 52
 - Content Search Engine 38
 - Process Engine 36
 - specifying 32
- ASA
 - installation tasks 50
- authentication
 - IBM Tivoli Directory Server 32

C

- Configuration Manager
 - setting permissions for user 56
- configurations
 - Baseline with options 2
 - demo 3
 - developer 3
 - sample 1
- Content Engine
 - configuring remote access
 - protocol 28
 - deployment on multiple servers 56
 - primary administrative user name 56
 - remote file access protocols 26
 - WebSphere 54
 - WebSphere environment variables 54
- Content Search Engine
 - preparing storage areas to support 28

D

- database
 - creating (DB2) 46
 - dedicated (DB2) 43
 - GCD (DB2) 48
 - Process Engine 46
 - shared (DB2) 43
 - storage area 24
 - verifying connection 50
- database accounts
 - creating 41
- database connection
 - verifying for DB2 50
- DB2
 - cataloging Process Engine node 49
 - client 49
 - creating Process Engine database
 - alias 49
 - dedicated database 43
 - install for FileNet P8 platform 44
 - installing 45
 - JDBC drivers 53
 - page size 45
 - shared database 43

- DB2 client
 - installing 49
- DB2 client instances
 - creating 49
- DBA
 - installation tasks 39
- Decru 24
- deployment
 - planning 3
 - planning application server 4
 - planning Content Engine 3
 - planning Content Engine multi-server on WebSphere and WebLogic 5
- directory servers
 - IBM Tivoli Directory Server 32
- documentation server
 - installing 58

E

- encrypted NTFS devices 24
- Encryption products for storage 24
- extensible authentication ports 65

F

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

H

- hostname 21
- hosts file 21

I

- IBM Encryption expert 24
- IBM Tivoli Directory Server 32
- installation
 - ASA tasks 50
 - DBA tasks 39
 - ITA tasks 10
 - planning 1
 - planning and preparing 1
 - SA tasks 29
- installation and upgrade worksheet 7
- IPv4 21
- IPv6 21
- ITA
 - installation tasks 10

J

- JDBC drivers 53
 - WebSphere
 - install for DB2 53

K

- Kerberos 29

N

- network name 21

O

- Object Request Broker(ORB)
 - application server-based PE clients 57
- object store
 - DB2 database 48
- operating system accounts
 - creating 13

P

- ports 65
- Process Engine
 - configuring AIX 22
 - configuring UNIX 20
 - configuring UNIX servers 20
 - database accounts, creating 41
 - operating system accounts, creating 13
- Process Engine (AIX)
 - configuring 22
 - kernel limits 22
- Process Engine tablespaces
 - DB2 database 46
- profile 54

R

- realm 29
- roles
 - definition of installation 5

S

- SA
 - installation tasks 29
- security
 - install considerations 29
 - specify accounts 32
- Single Sign-On 29
- SSO
 - Kerberos 29
 - planning considerations 29
 - requirements 29
- Storage areas 24

U

UNIX

- configuring for FileNet P8
- components 18

V

- Vormetrics 24

W

WebSphere

- configure for Content Engine 53
- configuring for Application Engine 57
- configuring for Workplace XT 57
- primary administrative user name 56

WebSphere profile for Content

- Engine 54

worksheet

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



Program Number: 5724-R76
5724-R81

GC19-2819-01

