


Plan and Prepare Your Environment for IBM FileNet P8



Plan and Prepare Your Environment for IBM FileNet P8

Note

Before using this information and the product it supports, read the information in “Notices” on page 279.

This edition applies to version 4.5.0 of IBM FileNet Content Manager (product number 5724-R81), version 4.5.0 of IBM FileNet Business Process Manager (product number 5724-R76), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2001, 2008.**

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

Table of contents

Table of contents	5
Revision log	11
About this document	12
Intended audience	12
Typographical conventions	13
Acronyms	14
About IBM FileNet documentation	15
Copy Web documents into the help system and make them searchable	15
Links to additional information	15
Gather reference documentation	16
Installation reference documents	16
Other available documentation	18
Autonomy K2 software documentation for configuring the FileNet P8 Content Search Engine	19
Access IBM FileNet documentation, compatibility matrices, and fix packs	19
Customer support	19
Feedback	20
Plan and prepare for IBM FileNet P8 installation	21
Plan the installation	22
General planning considerations	23
FileNet P8 client considerations	23
Non-English environment considerations	23
Review sample configurations	24
Baseline Configuration	25
Baseline Configuration With Optional Components	26
Developer Configuration	27
Demo Configuration	28
Plan your Content Engine deployment	29
Standalone deployment	29
Multi-instance single-server deployment	29
Multi-instance multi-server deployment	29
Managed deployment	29
Non-managed deployment	29
Cluster deployment	30
Definition of installation roles	31
Installation administrator	31
Information technology administrator	31
Security administrator	31
Database administrator	32
Application server administrator	32
FileNet P8 administrator	32

Email Administrator	32
Use the Installation and Upgrade Worksheet	33
Perform the required installation preparation tasks	34
IT Administrator tasks	35
Configure UNIX	36
Configure UNIX for FileNet P8 servers (all components)	36
Configure Content Engine servers (all UNIX)	36
Configure Content Search Engine servers (UNIX)	36
Configure Process Engine Servers (all UNIX)	38
Configure Process Engine servers (AIX)	40
Configure Process Engine servers (HP-UX)	42
Configure Process Engine servers (Solaris)	45
Configure Process Engine clients for ORB (all UNIX)	48
Configure Application Engine or Workplace XT servers (Linux)	48
Configure Microsoft Windows	49
Configure Windows for FileNet P8 servers	49
Configure Windows for Enterprise Manager, .NET and COM Compatibility clients	49
Content Engine authenticating with Active Directory	49
Process Engine on Windows	49
Content Search Engine on Windows	51
Configure network	52
Prepare storage areas for object stores	53
Configure File Servers for File Storage Areas	54
Remote File Access Protocols	54
Users and Groups	55
Configure the Remote Access Protocol on the Client Machine	57
Security Administrator tasks	59
Security planning considerations	60
Configure directory server	62
Configure Windows Active Directory	62
Configure Active Directory Lightweight Directory Services (AD LDS)	63
Configure Sun Java System Directory Server	63
Configure Novell eDirectory	64
Configure IBM Tivoli Directory Server	64
Specify IBM FileNet P8 accounts	65
Accounts for Content Engine	66
Task to be performed by: Database Administrator	66
Task to be performed by: Application Server Administrator	67
Task to be performed by: IT Administrator	68
Task to be performed by: Security Administrator	74
Accounts for Process Engine (Windows)	79
Accounts for Process Engine (UNIX)	80
Accounts for Application Engine or Workplace XT	85
Task to be performed by: Application Server Administrator	85
Accounts for Content Search Engine	87
Task to be performed by: IT Administrator	87
Database Administrator tasks	88
Prepare Microsoft SQL Server	89

Microsoft SQL Server planning considerations	89
Verify that Microsoft SQL Server is installed for IBM FileNet P8	89
Get Microsoft SQL Server ready for IBM FileNet P8	91
Configure Microsoft SQL Server Client for Process Engine	94
Prepare Oracle Server	97
Oracle Planning Considerations	97
General	97
Process Engine	97
Verify that Oracle Server is installed for IBM FileNet P8	98
Configure Oracle Client for Process Engine	106
Prepare DB2 Server for z/OS	109
DB2 for z/OS planning considerations	109
Verify that DB2 for z/OS Server is installed for IBM FileNet P8	109
Prepare DB2 Server for Linux, UNIX and Windows	113
DB2 for Linux, UNIX and Windows planning considerations	113
Verify that DB2 Server is installed for IBM FileNet P8	114
Configure DB2 Client	121
Verify the ability to connect to the database	125
To verify the Process Engine database connection (Oracle)	125
To verify the Process Engine database connection (DB2)	126
To verify the Process Engine database connection (SQL Server)	127
Application Server Administrator tasks	128
Application Server Planning Considerations	129
All platforms	129
WebSphere and WebLogic - Content Engine Multi-Server Deployment	129
JBoss Content Engine Multi-Server Deployment	130
Configure WebSphere for Content Engine	131
Configure WebLogic for Content Engine	134
Configure JBoss for Content Engine	138
Assign directory permissions	139
Configure JBoss server clusters	139
Configure WebSphere for Application Engine or Workplace XT	141
Configure WebLogic for Application Engine or Workplace XT	142
Configure JBoss for Application Engine or Workplace XT	143
Configure Process Engine clients for ORB	144
Plan and prepare for IBM FileNet P8 upgrade	145
Plan the upgrade	146
Upgrade overview	147
Standard upgrade from version 3.5 or 4.0	147
Staged upgrade from version 4.0 only	147
Staged upgrade scenario 1: upgrade Application Engine or Workplace XT first	148
Staged upgrade scenario 2: upgrade Content Engine first	149
Staged upgrade scenario 3: upgrade Process Engine first	150
Upgrade planning considerations	151
General requirements for all IBM FileNet P8 systems	151
IBM FileNet P8 expansion product considerations	152
Content Federation Services	152

Records Manager	153
Vendor software considerations	153
Definition of upgrade roles	154
Installation administrator	154
Information technology administrator	154
Security administrator	154
Database administrator	155
Application server administrator	155
FileNet P8 administrator	155
Email Administrator	155
 Perform the required upgrade preparation tasks	 156
IT Administrator upgrade tasks	157
Configure UNIX	158
Configure UNIX for FileNet P8 servers (all components)	158
Configure Content Engine servers (all UNIX)	158
Configure Autonomy K2 Servers (UNIX)	159
Configure Process Engine Servers (all UNIX)	160
Configure Process Engine servers (AIX)	163
Configure Process Engine servers (HP-UX)	165
Configure Process Engine servers (Solaris)	169
Configure Process Engine clients for ORB (all UNIX)	171
Configure Application Engine or Workplace XT servers (Linux)	171
Configure Microsoft Windows	172
Configure Windows for FileNet P8 servers	172
Content Engine and Enterprise Manager	172
Content Engine authenticating with Active Directory	173
Process Engine on Windows	173
Content Search Engine on Windows	175
Configure network	176
Assign directory permissions for Content Engine upgrade for 4.0.x to 4.5 on UNIX	177
Process Engine	179
Security Administrator upgrade tasks	180
Security upgrade planning considerations	181
Specify IBM FileNet P8 Accounts for upgrades	183
Accounts for Content Engine upgrade	184
Task to be performed by: Database Administrator	184
Task to be performed by: Application Server Administrator	186
Task to be performed by: IT Administrator	189
Task to be performed by: Security Administrator	190
Database Administrator upgrade tasks	196
Database Administrator planning considerations	197
General	197
Oracle	197
Microsoft SQL Server	198
Create the GCD data source for Content Engine (upgrades from 3.5.x)	199
Prepare DB2	199
Prepare Oracle Server	199
Prepare SQL Server	200

Application Server Administrator upgrade tasks	202
Upgrades from 3.5.x	202
Application server planning considerations (upgrades from 3.5.x)	203
All platforms	203
WebSphere and WebLogic - Content Engine Multi-Server Deployment	203
JBoss Content Engine Multi-Server Deployment	204
Application server considerations	204
Configure the application server for Content Engine (upgrades from 3.5.x)	205
Configure WebSphere for Content Engine	205
Configure WebLogic for Content Engine	207
Configure JBoss for Content Engine	211
Assign directory permissions	212
Configure JBoss server clusters	213
FileNet P8 Administrator upgrade tasks	214
Appendixes.	215
Installation and upgrade worksheet	216
Process Engine SQL scripts	218
SQL Scripts for SQL Server	218
SQL Scripts for Oracle	219
Changes Made by pe_filenet_site.sql and pe_oracle_users_defaults.sql	220
Changes Made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql	221
Use Cases for Running Scripts and Setting Passwords	224
Installation and upgrade glossary	228
Installing P8 Platform in a non-English environment	248
Application Server Administrator	249
WebSphere and WebLogic	249
WebLogic	250
Security Administrator	251
Extended characters and user names	251
Database Administrator	252
Install Microsoft SQL Server	252
Install Oracle server	252
Content Engine	252
Process Engine	253
Install Oracle client	253
NLS_LANG environment variable - Charset	253
UNIX	254
Install the DB2 server	256
Content Engine	256
Process Engine (UNIX)	256
IT Administrator	257
Operating system considerations	257
Application Engine or Workplace XT	257

Content Engine	257
Process Engine	257
Microsoft Windows	258
UNIX	258
FileNet P8 Administrator	259
Process Engine services	259
Configure Process Task Manager for Application Engine and Workplace XT	259
Limitations	260
WebLogic	260
English and EnglishX locale considerations	260
GB18030 Support for Windows PE	260
AE, Workplace XT, or PE Task Manager for UNIX	260
Publishing	260
Process Engine Modeler	261
 IBM FileNet P8 database character sets	 262
IBM FileNet P8 Database components character sets	262
Additional Oracle character set information	262
 IBM FileNet P8 ports	 265
 Notices	 279
Trademarks	281

Revision log

The following table identifies changes made to this document since the IBM FileNet P8 Platform 4.5 release.

Date	Revision
11/08	Initial release of this document for release 4.5.

About this document

This document provides information about how to prepare a FileNet P8 environment so that it is ready for an installation of IBM FileNet P8, which includes Application Engine (AE), Workplace XT, Content Engine (CE), and Process Engine (PE).

Information about installing and configuring the FileNet P8 Platform environment is in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE Installations that use Workplace XT as the web-based client instead of AE must refer to the *IBM FileNet Workplace XT Installation and Upgrade Guide* for Workplace XT information when topics in this document are about AE.

To download these documents from the IBM FileNet support Web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Intended audience

This document is intended for software professionals who will install the FileNet P8 applications. Installation by an IBM FileNet Certified Professional (FCP) is recommended. For more information about the FCP program, contact your IBM service representative.

Typographical conventions

This document uses the following typographical conventions:

Convention	Usage	Example
Bold	Platform-specific headings	Start the application server. WebSphere Refer to IBM WebSphere documentation for more information. WebLogic Refer to BEA WebLogic documentation for more information.
Gray bold	Clickable items, such buttons, and tabs.	Click OK .
	Menu paths or breadcrumb trails.	Select Start > Settings > Control Panels > Display > Screen Saver .
<i>Italics</i>	Variables that require user-supplied values	The calculation is: <i>number of object stores</i> * 16 + <i>number of concurrent users</i> .
	Document titles	You are reading the <i>IBM FileNet P8 Platform Installation and Upgrade Guide</i> .
Monospace	Text that has to be typed by the user	Copy the file by entering the following command: COPY filename
	Code samples	Find the following text in the web.xml file: <context-param><param-name>uploadDir</param-name><param-value>/opt/FileNet/AE/Upload1</param-value> </context-param>
	Display text, such as prompts and error messages	Are you sure you want to delete this object? You do not have permission to delete this object.
	Elements such as filenames, properties, classes and so on, whose meaning might get confused in regular text.	Open the <code>filed</code> file. Enter a value for the <code>new</code> property. Select the <code>senior</code> class.

Convention	Usage	Example
“Text with quotation marks”	User-interface fields that do not use initial capitalization and document headings referenced within a document	See the “Part number” field for the part number. For more information, see “Typographical conventions” on page 13 .
UPPERCASE	Case-sensitive text, where uppercase text is required.	Copy the file by entering the following command: <code>COPY filename</code>

Acronyms

This document uses the following IBM FileNet product names and acronyms.

Product Name	Acronyms
Application Engine	AE
Content Engine	CE
Content Federation Services	CFS
Content Search Engine	CSE
Enterprise Manager	EM
Global Configuration Data	GCD
Image Services Resource Adapter	ISRA
Process Engine	PE
Rendition Engine	RE

About IBM FileNet documentation

By default, this document is distributed as part of the *IBM FileNet P8* help system, but it is also available as a downloadable document from the IBM support Web site. Newer versions of *IBM FileNet P8* documentation are sometimes re-released with other events, such as fix pack releases or documentation refreshes. To ensure that you have the latest revision of a document, compare the document part number of your document to the document part number of the document that is posted on the support Web site:

www.ibm.com/support/docview.wss?rs=3278&uid=swg27010422

For example, the last two digits of “GC31-5585-05” indicate that the specified document has been revised five times after the original publishing, which is designated by 00.

Copy Web documents into the help system and make them searchable

The IBM FileNet P8 help system is designed so that you can download updated copies of this document and copies of other IBM FileNet documents into the *IBM FileNet P8* help system and index them so that they can be retrieved by a search in the help system. However, to search the *IBM FileNet P8* help system, it must be installed as a Web site on a Web server that supports Java™ applications.

For more information, see:

- “[Gather reference documentation](#)” on page 16
- “Install IBM FileNet P8 Platform documentation” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

Links to additional information

To help you locate additional information about a topic, this document includes links to:

- Other locations in this document
- External Web sites
- Topics in the FileNet P8 help system

Because this document is included in the default FileNet P8 help system, the links to the FileNet P8 help topics work only when you view this document from within the help system. If you view this document from outside of the help system, the links to FileNet P8 help topics do not work.

Gather reference documentation

Following are two tables with information about the IBM FileNet P8 documents that are available as part of the FileNet P8 release. To download these documents from the IBM FileNet support Web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Installation reference documents

Consider having one or more of the following documents (or help topics) nearby for reference purposes during the installation of IBM FileNet P8 Platform.

Document or help topic name...	Refer to this document...
<i>IBM FileNet P8 Platform Installation and Upgrade Guide</i>	For instructions on how to install and upgrade IBM FileNet P8 Platform, which includes Application Engine (AE), Content Engine (CE), and Process Engine (PE).
<i>IBM FileNet Workplace XT Installation and Upgrade Guide</i>	For information about how to install and upgrade Workplace XT in a FileNet P8 environment.
<i>IBM FileNet P8 Hardware and Software Requirements</i>	To confirm that the target environment has at least the minimum supported levels of software from independent software providers.
<i>IBM FileNet P8 Compatibility Matrix</i>	To confirm that the version of FileNet P8 to be installed is compatible with at least the minimum supported levels of other FileNet products that are or will be used with FileNet P8.
<i>IBM FileNet P8 Release Notes</i>	To familiarize yourself with the new features and known issues associated with the FileNet P8 release.
<i>IBM FileNet P8 Non-English Support Guide</i>	For information about how FileNet P8 supports non-English environments.
<i>IBM FileNet P8 High Availability Technical Notice</i>	For information about how to set up FileNet P8 using clusters, farms and other high-availability software and hardware.
<i>IBM FileNet P8 Performance Tuning Guide</i>	For critical tuning information required to make deployments of all sizes and levels of complexity work efficiently before going into production. Attention This guide provides many specific recommendations for making performance-related choices that are either difficult or impossible to change once the system goes into production.
<i>IBM FileNet P8 Troubleshooting Guide</i>	For troubleshooting tips associated with a FileNet P8 Installation.

Document or help topic name...	Refer to this document...
<i>IBM FileNet P8 help</i> topic: Administration > Enterprise-wide Administration > FileNet P8 Security > Users and Groups	For a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain a FileNet P8 system.
<i>IBM FileNet P8 help</i> topic: Administration > Enterprise-wide Administration > Shutdown and Startup	For information about how to shut down and start up FileNet P8 and any expansion products.

Other available documentation

Review the list of remaining FileNet P8 documents that you can download from the FileNet support Web site.

Document name	Refer to this document...
<i>IBM FileNet P8 System Overview</i>	For a technical summary of the FileNet P8 architecture, including a overview of features and capabilities.
<i>IBM FileNet P8 Disaster Recovery Technical Notice</i>	For information about potential options and solutions involved in a FileNet P8 disaster recovery plan.
<i>IBM FileNet P8 Process Task Manager Advanced Usage Technical Notice</i>	For information about properties found under the Advanced tab in Process Task Manager.
<i>IBM FileNet P8 Version Tools Technical Notice</i>	For information about the set of versions tools that are installed with FileNet P8 Platform and that identify the levels of Application Engine, Content Engine, and Process Engine in a FileNet P8 environment.
<i>IBM FileNet P8 Security Help Extract</i>	For security-related information from the FileNet P8 help system in PDF format.
<i>IBM FileNet Rendition Engine Installation and Upgrade document</i>	For information about how to install and upgrade Rendition Engine in a FileNet P8 environment.
<i>IBM FileNet P8 eForms Installation and Upgrade Guide</i>	For information about how to install and upgrade eForms in a FileNet P8 environment.
<i>IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide</i>	For information about how to install and upgrade IBM FileNet Connectors for SharePoint Web Parts in a FileNet P8 environment.
<i>IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide</i>	For information about how to install and upgrade IBM FileNet Connectors for SharePoint Document Libraries in a FileNet P8 environment.
<i>IBM FileNet P8 Portlets for WebSphere Installation and Upgrade Guide</i>	For information about how to install and upgrade FileNet P8 Portlets for WebSphere in a FileNet P8 environment.
<i>IBM FileNet Process Analyzer Installation and Upgrade Guide</i>	For information about how to install and upgrade Process Analyzer in a FileNet P8 environment.
<i>IBM FileNet Process Simulator Installation and Upgrade Guide</i>	For information about how to install and upgrade Process Simulator in a FileNet P8 environment.
<i>IBM FileNet Records Manager Installation and Upgrade Guide</i>	For information about how to install and upgrade Records Manager in a FileNet P8 environment.
<i>IBM FileNet Business Activity Monitor Installation and Configuration Guide</i>	For information about how to install and configure Business Activity Monitor in a FileNet P8 environment.

Document name	Refer to this document...
<i>IBM FileNet Content Federation Services Installation and Upgrade Guide</i>	For information about how to install and upgrade Content Federation Services in a FileNet P8 environment.
<i>IBM FileNet Content Federation Services for Image Services Planning and Configuration Guide</i>	For information about how to configure Image Services for document federation.
<i>IBM FileNet Content Federation Services for Content Manager OnDemand Planning and Configuration Guide</i>	For information about how to configure Content Manager for OnDemand for document federation.
<i>IBM FileNet Content Management Widgets Installation Guide</i>	For information about how to install Content Management Widgets in a FileNet P8 environment.

Autonomy K2 software documentation for configuring the FileNet P8 Content Search Engine

Autonomy (formerly Verity) K2 software, which underlies the optional CSE component, installs with a large body of documentation that is not included in the general FileNet P8 documentation.

For details on how to access this documentation, see the "Configure Content Engine for Content-Based Retrieval" topic in the IBM FileNet P8 Platform Installation and Upgrade Guide.

Access IBM FileNet documentation, compatibility matrices, and fix packs

To access documentation, compatibility matrices, and fix packs for IBM FileNet products:

1. Navigate to the Product Documentation for FileNet P8 Platform support page.
(<http://www-1.ibm.com/support/docview.wss?rs=3247&uid=swg27010422>).
2. Select a PDF or a Doc Link, whichever is appropriate.

Customer support

For information about contacting customer support:

1. Navigate to the FileNet Product Family support page:
(<http://www-01.ibm.com/software/data/content-management/filenet-product-family/support.html>).
2. Click **IBM FileNet Support Communications**, or search for a particular support topic under "Enter search terms".

Feedback

Your feedback helps us to provide quality information. Send your comments about this publication or any other IBM FileNet documentation by e-mail to comments@us.ibm.com. Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a chapter and section title, a table number, a page number, or a help topic title).

Plan and prepare for IBM FileNet P8 installation

The topics in this section explain the process for planning for and performing the required prerequisite tasks for a FileNet P8 installation. Review all of the topics in the following sections carefully before you begin.

- [“Plan the installation” on page 22](#)
- [“Perform the required installation preparation tasks” on page 34](#)

Plan the installation

This section describes considerations and information you should use when planning your IBM FileNet P8 system installation. Review this section thoroughly before you start to set up IBM FileNet P8 components or required third-party software.

Before you begin to install the IBM FileNet P8 Platform, review the following information:

- [“General planning considerations” on page 23](#)
- [“Review sample configurations” on page 24](#)
- [“Plan your Content Engine deployment” on page 29](#)
- [“Definition of installation roles” on page 31](#)
- [“Use the Installation and Upgrade Worksheet” on page 33](#)

General planning considerations

FileNet P8 client considerations

Determine your IBM FileNet P8 user application requirements. Install Application Engine if you will be using Workplace or creating a custom user application. Alternatively, you can install and use Workplace XT. If you plan to use Workplace XT, installing Application Engine is not required.

Non-English environment considerations

If you are installing IBM FileNet P8 Platform in a non-English environments, there may be additional preparation and configuration tasks required. These tasks are provided in an appendix to this document, [“Installing P8 Platform in a non-English environment” on page 248](#). Review this appendix carefully as part of your planning process, before you begin your preparation tasks.

Review sample configurations

This topic shows you some simple examples of how to distribute FileNet P8 Platform components across a variety of machines. Each example represents a minimum recommended configuration. The configurations include the major IBM FileNet P8 Platform components, both those that are core required components and those that are expansion product add-ons.

This topic includes the following sample configurations:

- [“Baseline Configuration” on page 25](#)
- [“Baseline Configuration With Optional Components” on page 26](#)
- [“Developer Configuration” on page 27](#)
- [“Demo Configuration” on page 28](#)

In all the sample configurations, note that:

- Collocation issues are not addressed. For details on collocation, see *IBM FileNet P8 Hardware and Software Requirements*.
- None of the samples shows the required directory service provider.
- You can scale out the components, but the following graphics do not attempt to show this.
- To administer Application Engine, you can run IBM FileNet P8 Workplace clients either from the computers shown or from one or more browser clients.
- You can also run Workplace as a user client to create and access stored content and processes. Optionally, you can configure these clients to integrate with Microsoft Outlook and Office applications, or to work in conjunction with expansion products, for example, IBM FileNet P8 eForms, IBM FileNet P8 Portlets, or IBM FileNet Records Manager, as they come available on a given release. You can also use WebDAV clients.
- If you want to apply business rules to Process Designer workflows, you can install a rules engine of your choice, such as ILOG JRules, which is not shown in the samples.
- Optional components not shown in the samples include: IBM FileNet P8 Portlets and Image Services Resource Adapter (ISRA). Check with your service representative for availability of other expansion products.
- You must set up Content Engine, Application Engine, and the P8 Platform Documentation on application servers. You can collocate the documentation with Content Engine or Application Engine, or deploy it on a dedicated server, as shown.

Baseline Configuration

This configuration shows a typical setup where basic process and content capabilities are required. It includes only the core components, and not the other add-on components that are shipped with the IBM FileNet P8 Platform as expansion products.




Figure 1: Baseline Configuration

Baseline Configuration With Optional Components

This configuration is useful for environments that plan to use not only the core IBM FileNet P8 Platform components but also optional expansion product components.

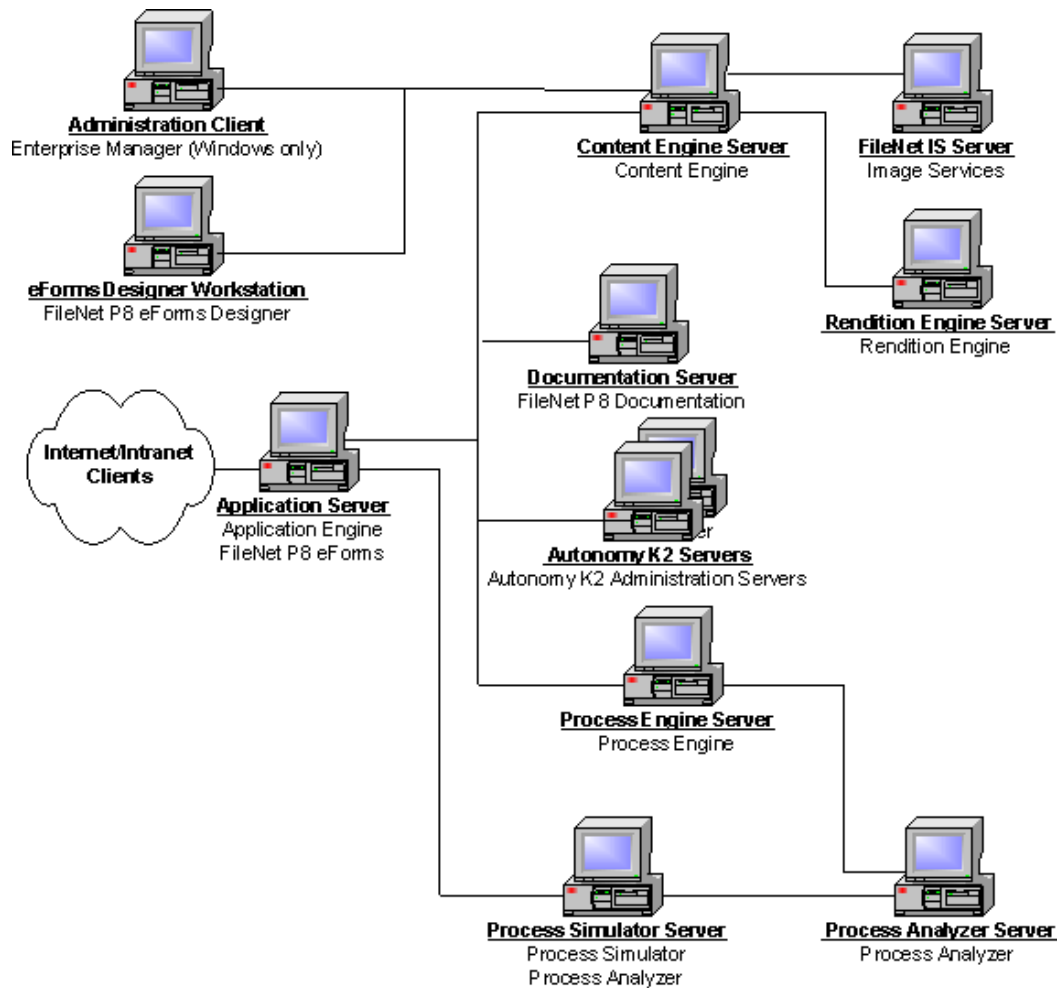


Figure 2: Baseline Configuration With Optional Components

Developer Configuration

This configuration illustrates how a development team might set up an environment for building an application that leverages the IBM FileNet P8 Platform.




Figure 3: Developer Configuration

NOTES

- It is useful to share all services, with the exception of Image Services, among the development workstations.
- You can use terminal services to run the Enterprise Manager from developer workstations, or install it directly on any Windows workstation.

- This configuration represents a single IBM FileNet P8 domain.
- 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.

Demo Configuration

This configuration supports demos, proof-of-concepts, and development on a single Windows server.



Figure 4: Demo Configuration

NOTE Avoid collocating IBM FileNet P8 components. See the *IBM FileNet P8 Hardware and Software Requirements* for details. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.](#)

Plan your Content Engine deployment

You can deploy Content Engine as a standalone (single instance) application, as multiple instances with different names on the same server, or in a cluster. Cluster deployments can be used for load balancing or to provide high availability.

Standalone deployment

When you deploy Content Engine as a standalone application, you configure a single application server. You will perform all the procedures in [“Configure Content Engine instances” on page 32](#) and in [“Deploy Content Engine instances” on page 46](#), both in the *IBM FileNet Platform Installation and Upgrade Guide*, 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 will perform all the procedures in [“Configure Content Engine instances” on page 32](#) and in [“Deploy Content Engine instances” on page 46](#) 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.

Multi-instance multi-server deployment

You can deploy Content Engine instances on multiple machines in a managed or non-managed application-server environment.

Managed deployment

When you deploy Content Engine in a WebSphere or WebLogic managed environment, you install and configure Content Engine on the Network Deployment node (WebSphere) or Administrative node (WebLogic). You do not install or configure Content Engine on any managed nodes.

NOTE JBoss has no managed-deployment capability.

Non-managed deployment

When you deploy Content Engine in a WebSphere or WebLogic non-managed environment, you install and configure Content Engine on each server in the environment, creating a distinct EAR file on each server. The contents of the EAR files have comparable configuration information (possibly with different directory paths).

When you deploy Content Engine in a JBoss non-managed environment, you install and configure Content Engine on an initial server in the environment, and then copy the Content Engine EAR file from the initial server to the other servers in the environment.

Cluster deployment

Cluster deployments can be used for load balancing or to provide high availability. Refer to the *IBM FileNet P8 Platform High Availability Technical Notice* for details on how to set up your IBM FileNet P8 system using clusters, farms, and other high availability software and hardware. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Definition of installation roles

The tasks in this *Plan and Prepare Your Environment for IBM FileNet P8* guide as well as the rows in the [“Installation and upgrade worksheet” on page 216](#) are organized by administrative roles, listed below.

Your organization may 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 installers during initial setup.
- Runs the Configuration Manager tool during initial setup, followed by launching Enterprise Manager.
- Runs IBM FileNet 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. See [“Installation and upgrade worksheet” on page 216](#).

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

Information technology administrator

- Responsible for the networking and operating systems issues 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.

For tasks assigned to the ITA, see [“IT Administrator tasks” on page 35](#).

Security administrator

- Responsible for configuring the directory servers required by IBM FileNet P8, including Content Engine, Application Engine, CFS Federation Administration application.
- Creates and maintains directory server user and group accounts.
- Decides on configuration parameters required to connect.
- 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.

For tasks assigned to the SA, see [“Security Administrator tasks” on page 59](#).

Database administrator

- Creates, configures, maintains database installations and database/tablespaces.
- Responsible for creating database accounts needed by FileNet P8.
- Might have responsibilities regarding the JDBC datasources.
- 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.

For tasks assigned to the DBA, see [“Database Administrator tasks” on page 88](#).

Application server administrator

- Responsible for providing the application servers required by 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.

For tasks assigned to the ASA, see [“Application Server Administrator tasks” on page 128](#).

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 using the *gcd_admin* account or an *object_store_admin* account is considered a 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.

Email Administrator

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

Use the Installation and Upgrade Worksheet

This planning and preparation guide includes a link to an associated *Installation and Upgrade Worksheet*, which is a spreadsheet file. You can enter information into the worksheet while you perform the preparation tasks for your FileNet P8 Platform environment. Information you collect in the worksheet will be used to complete various FileNet P8 installation, upgrade, and configuration tasks.

Here are some things to know about the worksheet:

- The rows in the worksheet correspond to field names, properties, and parameters that must be entered in the course of installing or upgrading FileNet P8 Platform components.
- Each row of the worksheet is designated for a particular administrator role. Review the rows in the worksheet before you begin your preparation tasks.
- Less obvious field names, properties, and parameter values have corresponding entries in the Installation and Upgrade Glossary, which appears as an appendix in this guide. Consult the Installation and Upgrade Glossary as part of your review of the worksheet rows to ensure that you can provide the correct values for the installation administrator.
- The *Installation and Upgrade Worksheet* is designed to be customized. You can sort the worksheet by different columns to get a specific view of the information that each administrator is expected to provide. Note that only some of the rows assigned to each administrator will apply for a particular installation. For example, some rows are specific to platform type, application server vendor type, database type, and so forth. Each row is also specific to a FileNet P8 component or add-on product. Work with your installation administrator or IBM services representative to determine the rows that are relevant for your installation. If these individuals provide you with variations of this table, use such variations instead to collect your data.

For more details on using the worksheet, and a link to the file, see [“Installation and upgrade worksheet” on page 216](#).

Perform the required installation preparation tasks

The tasks in this section are divided by administrator role. For information about assigning and defining these roles, see [“Definition of installation roles” on page 31](#).

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 [“Installation and upgrade worksheet” on page 216](#) for details.

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

- [“IT Administrator tasks” on page 35](#)
- [“Security Administrator tasks” on page 59](#)
- [“Database Administrator tasks” on page 88](#)
- [“Application Server Administrator tasks” on page 128](#)

IT Administrator tasks

As the Information Technology Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform:

- Review all rows assigned to the IT Administrator (ITA) in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

HINT 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 the Role you are interested in.
- 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).
- For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- If you are installing in a non-English environment, review the considerations and procedures in [“Installing P8 Platform in a non-English environment” on page 248](#) before you begin your preparation tasks.
- Configure the operating systems and network to prepare for component installation:
 - [“Configure UNIX” on page 36](#)
 - [“Configure Microsoft Windows” on page 49](#)
 - [“Configure network” on page 52](#)
- Prepare your storage areas for the creation of object stores:
 - [“Prepare storage areas for object stores” on page 53](#)
- Create the operating system accounts required by IBM FileNet P8, which are listed for the IT Administrator in the following:
 - [“Task to be performed by: IT Administrator” on page 68](#)

Configure UNIX

The topics in this task describe how to configure UNIX on the servers that will comprise your FileNet P8 system.

Configure UNIX for FileNet P8 servers (all components)

To configure UNIX FileNet P8 servers

1. **Ensure hosts file contents.** On each UNIX-based IBM 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 [“To configure the `/etc/hosts` file” on page 39](#) for additional information.

2. Consult with the application server, database, and P8 administrators to determine port requirements for all the servers in your install environment. For details, see [“IBM FileNet P8 ports” on page 265](#).

Configure Content Engine servers (all UNIX)

To configure UNIX FileNet CE servers

Content Engine running on a UNIX-based application server. 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/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 will need to specify when creating file storage areas on UNIX file servers.

NOTE This `umask` setting is required for the user running Content Engine Setup but need not be in the `.profile` file of the user.

Configure Content Search Engine servers (UNIX)

To configure HP-UX for CSE

To install Content Search Engine (Autonomy K2) on HP-UX, manually configure the kernel with following parameters before you begin the Autonomy K2 Master Administration Server installation:

Value	Setting
<code>maxdsiz</code>	1.9 Gbytes (0x7B033000)
<code>maxfiles</code>	2048 Kbytes

Value	Setting
maxfiles_lim	2048 Kbytes
maxssiz	160 Mbytes (0xA000000)
max_thread_proc	1024
maxswapchunks	8192
maxtsiz	1 Gbyte (0x40000000)
maxuprc	512
maxusers	128
nkthread	1024
nproc	517

To configure RedHat Enterprise Linux 5.1 for CSE

For RedHat Enterprise Linux 5.1 x86_64, install the following libraries from the original system media:

- *DVD-ROM_mount_point/Server/compat-libstdc++-33-3.2.3-61.i386.rpm*
- *DVD-ROM_mount_point/Server/compat-libstdc++-33-3.2.3-61.x86_64.rpm*

To set the Java_home variable for all UNIX types

Set the Java_Home variable for the supported Java SE Development Kit(JDK).

```
JAVA_HOME=/usr/jdk install path
export JAVA_HOME
```

NOTE Enter the Java_home variable in the .profile file to set this variable each time the user logs in.

To append the appropriate environment variables

Set the following variables, depending on your operating system:

AIX (default install path)

- `PATH=$PATH:/opt/verity/k2/_rs6k43/bin`
`export PATH`
- `LIBPATH=$LIBPATH:/opt/verity/k2/_rs6k43/bin`
`export LIBPATH`

HP-UX (default install path)

- `PATH=$PATH:/opt/verity/k2/_hpux/bin`
`export PATH`
- `SHLIB_PATH=$SHLIB_PATH:/opt/verity/k2/_hpux/bin`
`export SHLIB_PATH`

Linux (default install path)

- `PATH=$PATH:/opt/verity/k2/_ilnx21/bin`
`export PATH`
- `LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/verity/k2/_ilnx21/bin`
`export LD_LIBRARY_PATH`

Solaris (default install path)

- `PATH=$PATH:/opt/verity/k2/_ssol26/bin`
`export PATH`
- `LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/verity/k2/_ssol26/bin`
`export LD_LIBRARY_PATH`

Configure Process Engine Servers (all UNIX)

To configure UNIX servers for Process Engine

1. Perform the following prerequisite tasks in any order:

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

```
.cshrc
.Xdefaults
.Xresources
.dbxinit
.dtprofile
.env
.login
.mwmrc
.xinitrc
.profile
```

- Process Engine requires the presence of several partitions. 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.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fns	/fns file system	2GB	fns	fnusr	775
local	/fns/local file system	1GB	fns	fnusr	775
fn_sec_db0 (raw)	n/a logical volume	64MB	fns	fnusr	664
fn_sec_rl0 (raw)	n/a logical volume	64MB	fns	fnusr	664

NOTES

- (AIX 6.1 only) Permissions must be set correctly on both the /fns and /fns/local mount points and the file systems before mounting the file systems.
- Solaris volume management software might use port 32776. This is the default for the Process Engine Communication Port (IOR port).

To configure the /etc/hosts file

Information must be entered into either the server's DNS table or the hosts file related to Process Engine IP address, server name and NCH domain name. For non-farmed configurations this information can be in either the DNS table or the hosts file on the server. For farmed configurations this must be entered into the host file. In a farmed environment, entries must exist for every Process Engine server in the farm.

Entries must be 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 nch_domain-organization-nch-server load_balancer_name

where:

IP addr is the IP address of the Process Engine server.

hostname is the corresponding host name.

nch_domain-organization is the NCH domain and organization name, as provided to the Process Engine installation program.

load_balancer_name is the name of the load balancer in a farmed configuration.

When entering the domain-organization name, follow these rules:

- Eliminate all characters except ASCII alphanumeric characters and underscores.
- Convert all upper case characters to lower case.
- Insert a hyphen between the domain and organization names.
- Append “-nch-server” as a literal.

For example, a Process Engine has a domain “ace-1” and organization “FileNet”. Its IP address is 123.45.6.78. For this system, the hosts file entry is:

```
123.45.6.78 ace-1 ace1-filenet-nch-server
```

NOTE The hyphen in the *nch_domain-organization* name has been removed and the F and N in the “FileNet” organization name have been converted to lower case.

Configure Process Engine servers (AIX)

To configure AIX servers for Process Engine

1. Perform the following prerequisite tasks in any order:
 - The Kernel must be set to 64-bit mode.
 - The swap space must be set to 1.5 - 2 times RAM.
 - The Maximum Number of Processes allowed per user must be set to at least 400.
 - The Maximum Kbytes of real memory allowed for MBUFS must be set 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.
 - The Maximum Number of FIXED licenses (Num) must be set to a minimum of 16.
 - The following filesets must be installed and committed:
 - 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).

To modify /etc/rc.dt and /etc/tunables/nextboot for AIX 5.3 and 6.1

1. As the root user, execute the following commands or set them by editing the /etc/tunables/nextboot.

```
/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
```

2. Add the following statements at the beginning for /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
```

3. Restart the server (shutdown -Fr) for these settings to take effect. Executing the commands at the command line is not sufficient. The changes must be generated via the "nextboot" to avoid bind failures..

4. Check the values by executing:

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

To correct a required link in AIX 6.1

AIX 6.1 installs the file /usr/lib/libMrm.a in a directory that is different from the one required by the IS 4.1.0 mini-installer. As a result, the Process Engine installation will fail when running lic_admin.

To prevent this failure, use the following workaround after you install AIX 6.1, but before you install IS4.1.0:

1. Log in as a user with root privileges.
2. Remove any "filenet-*" entries in the \etc\services file.
3. Download and install either SP3 or APAR: IZ13179 on AIX 6.1.
4. Enter the following command:

```
ln -s /usr/lpp/x11/lib/R1/libMrm.a /usr/lib/libMrm.a
```

To install a required Oracle patch for AIX 6.1

AIX 6.1 requires the Oracle 10gR2 6613550 patch to fix a problem with rootpre.sh.

1. Download patch number 6613550 from the Oracle support web site.
2. As a user with root privileges, run the script.
3. As Oracle user, launch the Oracle Universal Installer (runInstaller).

Configure Process Engine servers (HP-UX)

The following operating system prerequisites apply to HP-UX FileNet PE servers.

To configure HP-UX servers for Process Engine

1. Perform the following prerequisite tasks in any order:
 - On each HP-UX 11 or HP-UX 11i machine where you will install a JVM-based IBM FileNet P8 component (such as Content Engine), or where an associated third-party JVM-based component (such as WebLogic or WebSphere) will run, increase the values of the kernel parameters `max_thread_proc` (maximum number of threads per process) and `nkthread` (maximum number of kernel threads in the system) beyond their default values, which are too small for IBM FileNet P8 applications.
 - Refer to the HP web page "Programmer's guide for Java 2 HP-UX configuration for Java support" for tools to determine values of these two kernel parameters that are sufficient for IBM FileNet P8.
 - The physical memory must be at least 512 MB.
 - The Kernel must be set to 64-bit mode.
 - The swap space must be set as follows:
 - Two times RAM if RAM < 1GB
 - 1.5 times RAM if RAM between 1GB and 2GB
 - Equal to RAM if between 2GB and 8GB
 - .75 times RAM if > 8GB
 - The `/etc/nsswitch.conf` should have the following entry:

```
hosts:  files [NOTFOUND=continue] dns
```

To perform symbolic links for X11 libraries

1. Log on as the root user.
2. At the prompt, execute the following:

```
cd /usr/lib
ln -s /usr/lib/libX11.3 libX11.sl
ln -s /usr/lib/libXIE.2 libXIE.sl
ln -s /usr/lib/libXext.3 libXext.sl
ln -s /usr/lib/libXhp11.3 libXhp11.sl
```

```
ln -s /usr/lib/libXi.3 libXi.sl
ln -s /usr/lib/libXm.4 libXm.sl
ln -s /usr/lib/libXp.2 libXp.sl
ln -s /usr/lib/libXt.3 libXt.sl
ln -s /usr/lib/libXtst.2 libXtst.sl
```

To configure kernel parameters

Ensure that the following parameters are set to at least the values shown unless otherwise noted. The values are appropriate for both HP PA-RISC and Integrity operating systems unless noted otherwise. These values are sufficient to install and initialize the software but system tuning will be required, specifically for the nfiles and maxfiles parameters.

Kernel Parameter	Minimum Setting PA-RISC	Minimum Setting Integrity V2 and V3 unless noted otherwise
maxdsiz	0x10000000 or 268435456 (256MB)	0x10000000 or 268435456 (256MB)
maxfiles	512	1024
nproc	1005	1005
maxuprc	400	400
nfile	1024	2048
ninode	1085	1085
semms	2000	2000
semmni	2000	2000
shmmax	0x10000000 or 268435456 (256MB)	0x20000000 or 536870912 (512MB)
shmseg	120	120
semmnu	1000	1000
semume	500	500
msgmni	2048	2048
msgseg	16384	16384 (obsolete in HP 11i V3)
msgtql	6640	6640
msgmap	msgtql + 2	msgtql + 2 (obsolete in HP 11i V3)

Kernel Parameter	Minimum Setting PA-RISC	Minimum Setting Integrity V2 and V3 unless noted otherwise
dbc_max_pct	1 to 10 (the value can not be greater than 30)	1 to 30 (the value can not be greater than 30) (obsolete in HP 11i V3)
dbc_min_pct	5	5 (obsolete in HP 11i V3)
timezone	Set appropriately	Set appropriately
msgmnb (DB2 only)	65535	65535
msgmax (DB2 only)	65535	65535

NOTE If you will be installing the 8.2 version of DB2 Client software, pay particular attention to the shmmax parameter. The minimum value documented for Process Engine might not be high enough to allow successful installation of the DB2 software. See the appropriate vendor documentation for the 8.2 release for recommended kernel parameter settings.

To check and optionally modify the timezone settings

HP-UX has two timezone settings: the kernel parameter timezone and the environment variable TZ. The value of both timezone settings must match. Review and, if necessary, change these settings on all servers.

1. As the root user, enter:

sam
2. Select the Kernel Configuration option.
3. Select the Configurable Parameters option, then check the Pending Value for the timezone parameter. The default value is 420 minutes west of Greenwich Mean Time (GMT), which is the U.S. Mountain timezone.

Determine the number of minutes east or west of GMT for your location by multiplying the number of hours east or west of GMT by 60 minutes per hour. For example, the U.S. Pacific timezone is 8 hours west of GMT. Multiply 8 x 60 to get 480 minutes. If your timezone location is east of GMT, you should use a negative number. For example, Middle European Time is one hour east of GMT. Multiply -1 x 60 to get -60 minutes for MET (Middle European Time).
4. If the Pending Value for the timezone parameter is correct, proceed to step 6. To change the value, continue with Step 5.
5. To change the timezone kernel parameter value:
 - a. Select the timezone parameter by pressing the spacebar and then press Tab to go to Actions menu.
 - b. Select the Modify Configurable Parameter option from the Actions menu and press Return.

- c. In the popup window that displays, the Specify New Formula/Value option should already be selected.
 - d. Tab to the Formula/Value field and type the new value.
 - e. Tab to OK and press Return. When the popup window disappears, you should see the new value in the Pending Value column.
 - f. Rebuild the kernel to make your change take effect:
 - i. Press the F4 key to access the menu bar.
 - ii. From the Action menu, select the Create a New Kernel option using the Arrow keys and press Return.
 - iii. Answer Yes when prompted about creating the kernel now.
 - iv. On the next screen, make sure the Move Kernel into Place and Shutdown/Reboot the System Now option is selected, tab to OK and press the Return key to reboot the system and make the new changes take affect.
6. As the root user, enter the following to check the current value of the TZ environment variable:

```
echo $TZ
```

7. If the current setting is not correct, enter the following to set the correct timezone:

```
/sbin/set_parms timezone
```

Choose the appropriate timezone from the menus displayed. Remember that the value must match that of the timezone kernel parameter.

If you change the current setting, you will be prompted to reboot the server.

NOTE If the HP-UX `set_parms` command is not available on your server, the timezone might be set via the SAM interface using the Kernel Parameters option in the same manner that other parameters are set. The System Administrator should consult the HP-UX operating system documentation to determine the appropriate way to set the TZ environment variable.

Configure Process Engine servers (Solaris)

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

To enable ports

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons. By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet uses several ports higher than 32768. See ["IBM FileNet P8 ports" on page 265](#) for details on which ports IBM FileNet uses.

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

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

1. To determine the current `tcp_smallest_anon_port` setting, enter the following at the command prompt:

```
ndd -get /dev/tcp tcp_smallest_anon_port
```

If the port is less than 32778, you must enable port 32777.

2. To enable port 32777 on Solaris9, use a text editor to edit the `/etc/rc2.d/S69inet` file.

Enter the following line:

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

3. To enable port 32777 on Solaris10, use a text editor to edit the `/lib/svc/method/net-init` file.

Enter the following line:

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

NOTE Put this entry in the file before the `exit 0` entry at the bottom of the file.

4. Reboot the Process Engine server to force the release of ports required by Process Engine that might be in use by the OS. Failure to reboot after these changes are made can result in port 32776 being unavailable, generating OpenSocket errors.

To verify national language character set and time settings

- The default time mask varies on UNIX depending on the `LANG` and `LC_TIME` environment settings. Verify the current `LC_TIME` settings by entering:

```
locale -k t_fmt
```

The result might appear similar to this:

```
t_fmt=%r
```

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

```
LC_TIME=C
export LC_TIME
```

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

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

```
LC_MESSAGES=en_US.ISO8859-1
LC_ALL=
```

To increase the operating system kernel limits

1. Make a copy of the system file (with a new name). Log on as root, and enter a command similar to the following:

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

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

```
vi /etc/system
```

3. Ensure that the following parameters are listed and are set to at least the values shown.

Solaris 9

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

Solaris 10

```
set semsys:seminfo_semmni=2000
set semsys:seminfo_semmsl=512
set semsys:seminfo_semopm=256
set shmsys:shminfo_shmmax=4294967295*
set shmsys:shminfo_shmmni=2000
set msgsys:msginfo_msgmni=2048
set max_nprocs=1000
set fnsod:sod_Debug=0
set rlim_fd_max=1024
set rlim_fd_cur=256
noexec_user_stack=1
```

* FileNet recommends this value be set to 4GB but do not set it higher than physical memory. It is recommended that this be set to less than 80% of physical memory.

4. Save your changes.

5. Reboot the server.

Configure Process Engine clients for ORB (all UNIX)

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 or custom applications

Therefore, if you have UNIX-based non-J2EE or custom applications, configure the Java installation on those servers with either the IBM or the Sun ORB.

Configure Application Engine or Workplace XT servers (Linux)

To configure Linux-based Application Engine or Workplace XT servers

Ensure that Linux® libraries are installed. To install Application Engine or Workplace XT on Linux, several legacy libraries are required. You must install the compat-libstdc++ packages on your RedHat system prior to beginning your install of Application Engine or Workplace XT.

Configure Microsoft Windows

Configure Windows for FileNet P8 servers

To configure Windows FileNet P8 servers

- Consult with the application server, database, and P8 administrators to determine port requirements for all the servers in your install environment. For details, see [“IBM FileNet P8 ports” on page 265](#).

Configure Windows for Enterprise Manager, .NET and COM Compatibility clients

To configure Windows for Enterprise Manager, .NET and COM Compatibility clients

1. On any Windows machine where you are going to install .NET API clients, COM Compatibility clients, or FileNet Enterprise Manager, you must first install the following:
 - Microsoft .NET Framework
 - Web Services Enhancements (WSE)

Check the IBM FileNet P8 Hardware and Software Requirements for the latest version requirements of these two components. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#). Enterprise Manager requires no other Content Engine services or files.

Content Engine authenticating with Active Directory

To configure Windows for Content Engine

- If Windows Active Directory is your directory service, set the primary DNS server IP address on your Content Engine machine to the IP address of the machine where DNS is installed.

Process Engine on Windows

To install Process Engine using a domain user

- If Process Engine will be installed by a domain user rather than a local user on the associated server, see [“Specify IBM FileNet P8 accounts” on page 65](#) for details on creating required users and groups.

To verify TCP/IP parameter settings (Windows)

1. Log on as the Administrator user and run regedit to verify the following registry key values.

NOTE These values are decimal. The default in regedit is hexadecimal.

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\MaxUserPort => 65534 (default = 5000)

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpTimedWaitDelay => 90 (default = 240, or 4 min)

2. If necessary, add or modify a new DWORD value with the values as described above and save the changes.

To add inbound rules to the Windows 2008 firewall

Configure inbound rules in the Windows 2008 firewall to allow the following Process Engine ports access.

Port	Protocol
32768	TCP
32769	TCP
32770	UDP

To configure the /etc/hosts file

Information must be entered into either the server's DNS table or the hosts file related to Process Engine IP address, server name and NCH domain name. For non-farmed configurations this information can be in either the DNS table or the hosts file on the server. For farmed configurations this must be entered into the host file. In a farmed environment, entries must exist for every Process Engine server in the farm.

Entries must be 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 nch_domain-organization-nch-server load_balancer_name

where:

IP addr is the IP address of the Process Engine server.

hostname is the corresponding host name.

nch_domain-organization is the NCH domain and organization name, as provided to the Process Engine installation program.

load_balancer_name is the name of the load balancer in a farmed configuration.

When entering the domain-organization name, follow these rules:

- Eliminate all characters except ASCII alphanumeric characters and underscores.
- Convert all upper case characters to lower case.
- Insert a hyphen between the domain and organization names.
- Append "-nch-server" as a literal.

For example, a Process Engine has a domain "ace-1" and organization "FileNet". Its IP address is 123.45.6.78. For this system, the hosts file entry is:

123.45.6.78 ace-1 ace1-filenet-nch-server

NOTE The hyphen in the *nch_domain-organization* name has been removed and the F and N in the “FileNet” organization name have been converted to lower case.

To configure Process Engine clients for ORB (Windows)

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

Therefore, if you have Windows-based non-J2EE and custom applications, configure the Java installation on those servers with either the IBM or the Sun ORB.

Content Search Engine on Windows

To configure security for Content Search Engine servers

Set *K2 Operating System User* as an administrator on each Windows Content Search Engine Server. See [“Specify IBM FileNet P8 accounts” on page 65](#) for details on creating required users and groups.

Configure network

The requirements in this Task apply to the network on which FileNet P8 servers are running.

To configure 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 FileNet P8 so that they can all communicate with one another.
- Ensure NetBIOS over TCP/IP is enabled on Windows.
- Ensure availability of required port numbers. Several port numbers are required by the various IBM FileNet P8 components. [See "IBM FileNet P8 ports" on page 265.](#)
- For information about proxy/firewall configuration requirements, see the *IBM FileNet P8 Hardware and Software Requirements* for support information related to IBM FileNet P8 components and database engines. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.](#)

To configure time and date

Synchronize the time and date on all servers. 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 need to 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.

Prepare storage areas for object stores

This task has two purposes:

- To prepare locations of initial storage areas for the object stores you will be creating, via the Create Object Store wizard (see [“Create the initial object store” on page 60](#)).
- To prepare locations of additional file storage areas (see [“Create Additional File Storage Areas” on page 171](#)) for object stores you already created.

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

- 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 [FileNet P8 Administration > Content Engine Administration > Content storage > File storage areas](#).

- 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 [FileNet P8 Administration > Content Engine Administration > Content storage > Fixed storage areas](#).

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

NOTES

- In this document, *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, Vormetric, 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 [FileNet P8 Administration > Content Engine Administration > Content storage > Fixed 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.

Configure File Servers for File Storage Areas

In this section you will 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. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

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:

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, as 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 as follows:

- Common Internet File System (CIFS)
- Network File System (NFS)
- Distributed File System (DFS)

NOTE DFS is supported if you are using it to manage a file storage area; however, the replication feature of DFS is not supported. For details on setting up a link to DFS, see IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content storage > File storage areas > How to...Create DFS link](#).

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

NOTE 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 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.

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

Users and Groups	Role
Content Engine operating system user (<i>ce_os_user</i>)	The user under which Content Engine Server executes (typically, the user that starts Content Engine Server).
K2 operating system user (<i>k2_os_user</i>)	The user under which Content Search Engine executes (typically, the user that starts Content Search Engine).
Content Engine operating system group (<i>ce_os_group</i>)	The group in which the Content Engine operating system user and the K2 operating system user are members.

The following procedures, which use the abbreviations in the following table, show how to configure UNIX-based and Windows-based file servers. Do the procedures that apply to your environment.

Abbreviation	Meaning
<i>ce_os_user</i>	Content Engine operating system user
<i>ce_os_group</i>	Content Engine operating system group
<i>fsa1</i>	Directory where content will be stored

For details on file storage area security, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Storage Area Security](#).

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 *fsa1* where content will be stored. 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
```

NOTE 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 Server and Content Search Engine are running. This will normally be true if all machines use the same directory service, but they may 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.

NOTE IBM recommends that trusted hosts be restricted 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.

To configure a Windows-based file server for a Windows client using CIFS

1. Log on to the Windows file server as *ce_os_user*.
2. Create (or designate) a directory *fsa1* where content will be stored. For example:

```
C:\> md c:\filenet\file_stores\fsa1
```
3. Navigate in Windows Explorer to *fsa1*, right-click the file icon, and choose Properties.
4. In the Security tab, click **Advanced**.
5. In the Advanced Security Settings dialog box,
 - a. Grant Full Control to *ce_os_user* and *ce_os_group*, and select *This Folder, subfolders, and files* from the *Apply onto* drop-down list.
 - b. Remove all other users and groups in the *Permission entries* table.
 - c. Click **OK**.

6. In the Sharing tab, do the following:
 - a. Click **Share this folder** and click **Permissions**.
 - b. Grant Full control to *ce_os_user* and *ce_os_group*.
 - c. Remove all other users and groups in the *Permission entries* table.
 - d. Click **OK**.

To configure a Windows-based file server for a UNIX client using NFS

1. Do all the steps in [“To configure a Windows-based file server for a Windows client using CIFS” on page 56](#).
2. Use the procedures in Microsoft documentation to configure Windows Services for NFS to expose *fsa1*.

NOTES

- Windows Services for NFS is an optional Windows component bundled with Windows Server 2003 R2.
- As part of configuring Windows Services for NFS, you must set up a mapping of Windows users and groups to UNIX users and groups. When setting up the mapping for *ce_os_user* and *ce_os_group*, you must specify the same UID (UNIX user ID) and GID (UNIX group ID) that these accounts have on the machine where Content Engine Server is installed.

Configure 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 and/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.

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, log on as the user who launched the application server.
2. Mount the exported NFS file system (from [Step 5 of “To configure a UNIX-based file server” on page 56](#)) 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 Server 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*.

To configure Windows-based Content Engine Server to talk to a Windows file server via CIFS

If both Content Engine Server and the file server are in the same Windows domain, no action is required. If they are in different domains, establish access to the file server machine from the machine where you will install Content Engine Server before creating the file storage areas.

Security Administrator tasks

As the Security Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform:

- Review all rows assigned to the Security Administrator (SA) in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- If you are installing in a non-English environment, review the considerations and procedures in [“Installing P8 Platform in a non-English environment” on page 248](#) before you begin your preparation tasks.
- Review the following planning information:
 - [“Security planning considerations” on page 60](#)
- Configure your directory server using the following task:
 - [“Configure directory server” on page 62](#)

Security planning considerations

The security 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.

- **Understand that authentication and authorization are separate processes.** As of the IBM FileNet P8 Platform 4.0.0 release, authentication (logon security) is separate from authorization (object and process security). Thus, you must configure your JAAS login on the Content Engine application server so that any user or group that can successfully authenticate to access (log on to) IBM FileNet P8 resources can also be authorized to work within IBM FileNet P8 interfaces, using Content Engine's directory service provider's connection.

The Content Engine Configuration Manager tool will capture 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 Content Engine's authorization by creating a Directory Configuration.

- **Understand that logins are done through JAAS.** IBM FileNet P8 leverages Java Authentication and Authorization Service (JAAS) for authentication only, 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 FileNet code.

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

- **Determine your single sign-on (SSO) requirements.** Content Engine's ability to leverage JAAS-based authentication means that if a single sign-on (SSO) provider writes a JAAS LoginModule for a supported application server, then clients of FileNet P8 applications hosted in that application server can leverage that SSO solution.
- **Determine Kerberos applicability.** You can use Kerberos for SSO authentication between 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](#).

This guide does not provide specific instructions for installing or configuring your SSO provider. For detailed reference information, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication](#).

- **Determine how many authentication realms you require.** At least one authentication realm is required, which you create during an initial installation. For an explanation of how to create and configure multiple realms, for example, multiple Windows domains, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms](#).

- **Ensure that you have a directory service provider in place.** Authorization in IBM FileNet P8 Platform is provided by one of the following supported directory servers:

- Microsoft Windows Active Directory
- Microsoft Active Directory Lightweight Directory Services (AD LDS, formerly named ADAM)
- Novell eDirectory
- Sun Java System Directory Server
- IBM Tivoli® Directory Server

This guide provides instructions for configuring the connections between Content Engine and the directory server. You can find additional detailed reference information in the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers](#).

- **Note that any WebLogic authentication provider should be dedicated to IBM FileNet P8.** For performance reasons, no authentication provider used by WebLogic for deployed IBM FileNet P8 components should be shared with applications used for other purposes.
- **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 Planning and Preparation Guide provides instructions for creating the administrative accounts required for installation and initial configuration. You can find additional detailed reference information for all users, groups, and administrative roles in the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups](#).
- **Note that Process Engine delegates authentication to Content Engine.** As of 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. Content Engine, in turn, runs under a J2EE application server, and relies on that server's JAAS-based facilities to authenticate users and groups against the chosen directory server.

Configure directory server

This task describes how to configure the directory server that will provide the authentication repository for your FileNet P8 system.

To configure your directory server, do one of the following:

- “Configure Windows Active Directory” on page 62
- “Configure Active Directory Lightweight Directory Services (AD LDS)” on page 63
- “Configure Sun Java System Directory Server” on page 63
- “Configure Novell eDirectory” on page 64
- “Configure IBM Tivoli Directory Server” on page 64

Configure Windows Active Directory

Use this task to configure Microsoft Windows Active Directory on a Windows server.

NOTES

- For a complete list of IBM FileNet P8-supported Windows Active Directory features, refer to [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Windows Active Directory](#).
- In a multi-domain Active Directory environment, a logon will fail for any account whose user name and password in a parent/child domain match those in a child/parent domain.

To enable DNS forwarders (when they are required for your network configuration)

DNS forwarders provide external DNS lookup functionality. If you are working in an “isolated” network, a DNS forwarder is not required. However, if you want to access the Internet or other network resources, then a DNS forwarder pointing to a DNS server that serves the external resources (for example, the Internet) is required. Take the following steps to enable DNS forwarders:

1. On the machine that is configured as the Windows DNS Server, log on with an account that can configure the DNS components.
2. Start DNS. For example, on Windows 2003, choose **Start > All Programs > Administrative Tools > DNS**.
3. Right-click the *your_computer_name* container and select Properties.
4. Select the *Forwarders* tab and verify the check box for *Enable forwarders* is selected.

NOTE If this feature is grayed out (unavailable), you will need to reconfigure your DNS server.

5. If you selected the check box, add an appropriate IP address and click **OK**.

NOTE This IP address may be the IP address of a DNS server that allows traffic to the Internet.

Configure Active Directory Lightweight Directory Services (AD LDS)

For IBM FileNet P8, no special settings are required for AD LDS (formerly ADAM). For a complete list of directory server features that IBM FileNet P8 supports, refer to the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers > Windows Active Directory Lightweight Application Mode \(AD LDS\)](#).

NOTES

- You can use AD LDS as a stand-alone directory service, or you can synchronize AD LDS with Active Directory, using Microsoft's built-in tools. Synchronization is invisible to IBM FileNet P8 applications and authentication. It is a best practice to establish the connection between Active Directory and AD LDS before installing IBM FileNet P8. Consult your AD LDS documentation for full information.

Configure Sun Java System Directory Server

You will use this task to configure Sun Java System Directory Server on a Windows or UNIX server. For a complete list of IBM FileNet P8-supported Sun Java System Directory Server features, refer to the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers > Sun Java System Directory Server](#).

NOTES

- On Windows servers, Sun Java System Directory Server should be installed on an NTFS hard drive partition.
- If there are more than 2,000 users in the Directory Server, you must increase the resource limits to correctly display users in IBM FileNet P8. IBM recommends setting this limit to -1 (unlimited). You can either set this limit for the entire LDAP server or for the individual IBM FileNet P8 users. Instructions for either procedure appear below.

To set the resource limits for the entire Directory Server (v 5.2)

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

- From the server where Sun Java System Directory Server is installed, log on with an account that has rights to modify the Sun Java System Directory Server environment.
- Run the Sun Java System Directory Server console and login.
- Expand the Domain > Server Group containers and select your **Directory Server**.
- Right-click and select **Open**.
- Select the Configuration tab.
- Select the Performance container.
- Select the Client Control tab.
- For the LDAP group box, ensure that *Size limit* and *Look-through limit* are both set to **Unlimited**.
- If changes were made, click **Save**.
- Select the Tasks tab and Restart the Directory Server if changes were made.

To set the resource limits for individual IBM FileNet P8 users

You will need to perform the steps below any time you add additional IBM FileNet P8 users.

1. From the Sun Java System Directory Server console, expand the Domain > Server Group containers and select your Directory Server. Then click **Open**.
2. Select the Directory tab.
3. From the left pane, select the *Object* (OU, etc.) that contains the user(s) you want to change.
4. For each IBM FileNet P8 user whose limit you want to change, complete the following steps:
 - a. From the right pane, double-click on the user name.
 - b. Select **Properties**.
 - c. On the left pane of the Properties dialog box, select **Account**.
 - d. Enter -1 in the **Look through limit** and **size limit** fields.
 - e. Click **OK**.
5. Restart the Directory Server.

Configure Novell eDirectory

For a complete list of IBM FileNet P8-supported Novell eDirectory features, refer to the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers > Novell eDirectory](#).

NOTES

- The Windows server where Novell eDirectory Server is installed, must have an NTFS hard drive partition.
- The Novell eDirectory administrator may have to create an index if the sorting attribute is not in the list of default attributes shipped by eDirectory.
- Access control settings in IBM FileNet P8 require that all users have Browse access on the directory server. If you do not want to set Browse access at the individual user level, it is a best practice to establish a Public trustee for the realm.
- IBM FileNet P8 supports cross-realm group memberships. This means that IBM FileNet P8 supports a configuration in which a group is in one realm while some or all of its users are in another.

Configure IBM Tivoli Directory Server

No special IBM FileNet P8-specific settings are required for IBM Tivoli Directory Server. For a complete list of IBM FileNet P8-supported IBM Tivoli Directory Server features, refer to the IBM FileNet P8 Help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers > Tivoli Directory Server](#).

Specify 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 [“Configure directory server” on page 62](#).

NOTE 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 [FileNet P8 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 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. An example of this is *pe_install_user*, the account you log in as to run the Process Engine Installer.

- By a name that looks like a variable but is not formatted in italics. Examples are the Process Engine's required accounts *f_maint* and *f_sw*, which, because they are not italicized, are not to be replaced by accounts of your choosing.

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.

Accounts for Content Engine

Task to be performed by: Database Administrator

To create Content Engine database accounts

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

User/Group	Description
<p>Database user name: Microsoft SQL Server <i>ce_db_user</i></p>	<p>The database owner accounts that Content Engine uses to access SQL Server, depending on whether you use one account for all Content Engine databases, or use one (for example, <i>ce_db_user</i>) for the GCD database and different accounts for each object store (for example, <i>ce_db_user1</i>, <i>ce_db_user2</i>, and so on).</p> <p><i>ce_db_user</i> can be a local account or a Windows domain account. It does not have to be an account in the configured directory service.</p> <p>Grant each <i>ce_db_user</i> at least the following database access permissions:</p> <ul style="list-style-type: none"> • public • db_owner <p>Add these accounts to SQL Server's master database and grant the public role to each. When you perform the procedure "To configure the JDBC Distributed Transaction Components" on page 93 these accounts will also be granted the SqIDBCXAUser role.</p>
<p>Database user name: Oracle <i>ce_db_user</i></p>	<p>The tablespace owner accounts that Content Engine uses to access Oracle. Use one account for each object store tablespace and one for the GCD tablespace.</p> <p><i>ce_db_user</i> does not have to be an LDAP account.</p> <p>Grant each <i>ce_db_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> • CREATE SESSION • CREATE TABLE • CREATE SEQUENCE (object store tablespaces only) • SELECT on pending_trans\$ • SELECT on dba_2pc_pending • SELECT on dba_pending_transactions • EXECUTE on dbms_system

Task to be performed by: Application Server Administrator

To create Content Engine application server accounts

1. 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 Administrative Console User <i>ce_appserver_admin</i> <i>ce_appserver_console_admin</i>	<p>An application server administrative account or accounts that can log on to the application server administration console.</p> <p>WebSphere: If your site uses a FederatedLDAP registry, <i>ce_appserver_console_admin</i> must be a unique user across all federated realms. If not, <i>ce_appserver_admin</i> and <i>ce_appserver_console_admin</i> can be the same account.</p> <p>WebLogic: <i>ce_appserver_admin</i> and <i>ce_appserver_console_admin</i> must be different accounts.</p> <p>(JBoss does not require an administrative account.)</p> <p>Login to the application server as <i>ce_appserver_admin</i> to perform application tasks such as the following:</p> <ul style="list-style-type: none"> • Create directory service providers within the application server. • Create JDBC Providers and/or data sources for db connectivity. • Deploy the Content Engine application. • Stop and restart servers and cluster members (WebSphere). • Stop and restart managed servers (WebLogic).

Task to be performed by: IT Administrator

To create Content Engine operating system accounts

1. Create new or designate existing installation accounts for Content Engine, as shown in the following table:

User/Group <i>unique_name</i>	Description
Application Server Installation Administrator: WebSphere WebLogic JBoss <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 Windows, <i>ce_appserver_install_user</i> must be a member of the local administrators group. • For UNIX, <i>ce_appserver_install_user</i> must have read, write, and execute permissions to the Content Engine installation directory. <p>Log in as <i>ce_appserver_install_user</i> to do the following:</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 "ce_appserver_install_group" on page 69.</p>

User/Group <i>unique_name</i>	Description
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"> • “ce_appserver_install_user” on page 68 • “ce_install_user” on page 70 user (Windows) or “ce_install_user” on page 70 (UNIX). • “config_mgr_user” on page 72 <p>Use the user accounts in <i>ce_appserver_install_group</i> to do the following:</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.
Content Engine Installer: Windows <i>ce_install_user</i>	<p>An operating system account you will use to log on to a machine to launch Content Engine install wizard.</p> <p><i>ce_install_user</i> can be the same user as the “ce_appserver_admin” on page 67.</p> <p>Use Windows administrative tools to add <i>ce_install_user</i> to the Local Administrators group and to the “ce_appserver_install_group” on page 69.</p>

User/Group <i>unique_name</i>	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 /tmp directory. • Membership in the “ce_appserver_install_group” on page 69. <p><i>ce_install_user</i> can be the same user as the “ce_appserver_admin” on page 67.</p>
Database user name: DB2 for Linux, UNIX, and Windows <i>ce_db_user</i>	<p>Operating system user accounts on the database server. These users are granted database permissions for Content Engine access to the DB2 database. Use one account for the GCD tablespace (for example, <i>ce_db_user1</i>) and one for each object store tablespace (for example, <i>ce_db_user2</i>, <i>ce_db_user3</i>, and so on).</p> <p>Use <i>ce_db_user</i> to do the following:</p> <ul style="list-style-type: none"> • Connect to the database • Create tables in the tablespace (CREATETAB) • Use the tablespace (USE OF) for User and User Temp tablespaces

User/Group <i>unique_name</i>	Description
<p>Database user name: DB2 for z/OS</p> <p><i>ce_db_user</i></p> <p>NOTE DB2 does not allow underscores in accounts names.</p>	<p>Operating system user accounts on the database server that are granted database permissions as follows:</p> <p>GRANT DBADM ON DATABASE <i>databasename</i> TO <i>operatingsystemuser</i>;</p> <p>GRANT USE OF STOGROUP <i>storagegroupname</i> TO <i>operatingsystemuser</i>;</p> <p>GRANT USE OF BUFFERPOOL BP32K TO <i>operatingsystemuser</i>;</p> <p>Use one account for the GCD (for example, <i>ce_db_user1</i>) and one for each object store (for example, <i>ce_db_user2</i>, <i>cce_db_user3</i>, and so on).</p>
<p>Accounts required for DB2 (Linux, UNIX, and Windows) database installation and administration:</p> <ul style="list-style-type: none"> Groups: instance owner primary group Users: instance owner <p><i>ce_db_db2_group</i></p> <p><i>ce_db_db2_instanceowner</i></p>	<p>Operating system users and groups that must exist on the database server.</p>

User/Group <i>unique_name</i>	Description
Configuration Manager user: WebSphere WebLogic JBoss <i>config_mgr_user</i>	<p>The operating system account you will use to run Configuration Manager. You will be instructed to grant additional permissions to <i>config_mgr_user</i> at several times in the <i>Planning and Preparation Guide</i>:</p> <ul style="list-style-type: none"> • “To set permissions for the Configuration Manager user” on page 133 (WebSphere) • “Assign directory permissions” on page 139 (JBoss) • “Give the Configuration Manager user the following permissions:” on page 137 (WebLogic) <p>Add <i>config_mgr_user</i> to the “<i>ce_appserver_install_group</i>” on page 69.</p> <p>(Windows only) Using Active Directory tools, add <i>config_mgr_user</i> to either the Power Users group or the Local Administrators group.</p> <p>After Content Engine is installed you must grant additional permissions to <i>config_mgr_user</i>:</p> <ul style="list-style-type: none"> • Execute permission to the Configuration Manager tool executable file, <i>configmgr.bat</i> (Windows) or <i>configmgr.sh</i> (UNIX). • Write permission to the directory where the Configuration Manager tool will create the configuration XML files. This directory is one of the following: <ul style="list-style-type: none"> – the directory you specify using the optional <i>-path</i> 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 <i>-path</i> parameter, or the default directory.

User/Group <i>unique_name</i>	Description
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 “Prepare storage areas for object stores” on page 53.</p> <p>Windows</p> <p>For Windows-based Content Engine and file storage areas, <i>ce_os_user</i> must reside in the same Windows domain or in trusted Windows domains as the servers that host Content Engine and the file storage area.</p> <p>UNIX</p> <p>For Unix-based Content Engine and file storage areas, configuring security requires the use of NFS.</p>

Task to be performed by: Security Administrator

To create Content Engine directory server accounts

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

User/Group	Description
<p>Bootstrap user name</p> <p><i>ce_bootstrap_admin</i></p>	<p>A directory service and application server account that is stored in the CEMPBoot.properties 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 login 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>CAUTION 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 FileNet P8 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/Group	Description
GCD Administrator <i>gcd_admin</i>	<p>A directory service account that has Full Control access to the Content Engine's domain object.</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>Logon 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 “Establish the FileNet P8 domain and Global Configuration Data (GCD)” on page 56 in the <i>IBM FileNet P8 Installation and Upgrade Guide</i>). • Carry out administrative tasks for the P8 domain. <p>For more information, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups.</p>
Object Store Administrator <i>object_store_admin</i> <i>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_admins</i> (see “To create an object store” on page 60 in the <i>IBM FileNet P8 Installation and Upgrade Guide</i>). Each object store could have a different set of object store administrators, depending on your security design.</p>

User/Group	Description
<p>Directory service user: Active Directory</p> <p>(Referred to as "Directory service bind user name" in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<p>An Active Directory user account that Content Engine uses to connect to Active Directory.</p> <p><i>ce_service_user</i> performs the following roles:</p> <ul style="list-style-type: none"> Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace. Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object once a user has been authenticated. <p>Using Active Directory tools, grant <i>ce_service_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> Member of the Pre-Windows 2000 Compatible Access Group in each desired domain in the Active Directory forest. <p>Provide the fully qualified distinguished name of <i>ce_service_user</i> as the Directory service bind user name while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>
<p>Directory service user: Windows Active Directory Lightweight Directory Services (AD LDS, formerly known as ADAM)</p> <p>(Referred to as Directory service bind user name in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<p>An AD LDS user account that Content Engine uses to connect to a single Microsoft AD LDS partition.</p> <p><i>ce_service_user</i> performs the following roles:</p> <ul style="list-style-type: none"> Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace. Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object once a user has been authenticated. <p>Using AD LDS administrative tools, grant <i>ce_service_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> Ability to see the other users in the partition. (For a procedure, see the entry for the AD LDS directory service user in FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups.) <p>Provide the fully qualified distinguished name of <i>ce_service_user</i> as the <i>LDAPBindDN</i> while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>

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

User/Group	Description
<p>Directory service user: IBM Tivoli Directory Server</p> <p>(Referred to as Directory service bind user name in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<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 FileNet P8 object once 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 <i>LDAPBindDN</i> while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>

Accounts for Process Engine (Windows)

To create Process Engine accounts (Windows)

Depending on your response to its prompt, the Process Engine installation program creates default user and group accounts or creates the accounts with your specified alias names.

There are two options:

- Allow the Process Engine installation program to create the Process Engine accounts or aliases.
- Pre-create the default accounts or aliases before running the Process Engine installation program.

The Process Engine installation program will ask if you want to configure aliases:

- A no answer indicates that the Process Engine installation program should create the default accounts and groups.
- A yes answer brings up a second screen allowing you to define aliases for each of the default accounts and groups. If you choose to create an alias for one Process Engine account, you must create aliases for all accounts, even if you set the alias to the default value of a Process Engine account.

1. Create the following user and group accounts.

User/Group	Description
fnadmin (group account)	Operating system group whose members have all privileges on Process Engine files and databases
fnusr (group account)	Operating system group whose members have non-administrator privileges on Process Engine files and databases
fnop (group account)	Operating system group whose members have operator non-administrator privileges on Image Services utilities used by Process Engine
fnsu (user account)	Operating system user account who executes Process Engine software
f_sw (user account)	Database runtime user (operating system user for DB2)
f_maint (user account)	Database maintenance user (operating system user for DB2)

NOTE The f_sw and f_maint database users will be created if the SQL scripts for SQL Server and Oracle databases are manually executed before running the Process Engine installation program. These scripts create users, passwords and stored procedures. There are no comparable scripts for DB2 databases. See [“Process Engine SQL scripts” on page 218](#) for details on the SQL scripts.

2. If you are going to run the Process Engine installation program while logged on as a Windows domain user, do the following:
 - a. Create the following security group accounts, with Windows domain local scope, on the Windows domain controller:

Default Name	Description
fnadmin or alias	Members have all privileges on IBM FileNet files and databases
fnusr or alias	Members have normal privileges on IBM FileNet files and databases
fnop or alias	Members have operator non-administrator privileges on Image Services utilities used by Process Engine

- b. Create the following user accounts on the Windows domain controller:

Default Name	Description
Process Engine installation account <i>pe_install_user</i>	User who will run the Process Engine installation program
fnsf or alias	Primary IBM FileNet software user NOTE Set the password to BPMtemp1pzd. This password is case-sensitive. You can change the password after installing Process Engine.

- c. Add the *pe_install_user* and fnsf users to the fnadmin, fnusr, and fnop groups.
 - d. Log on to the machine where you will install Process Engine as a member of the Domain Admins group.
 - e. Add the *pe_install_user* and fnsf users to the local Administrators group.

Accounts for Process Engine (UNIX)

On UNIX platforms, Process Engine accounts, or their aliases, must always be created on the Process Engine server before installing the Process Engine software. If the database is DB2, the f_sw and f_maint accounts, or their aliases, must be created as operating system users. If the database is DB2 for z/OS, the f_sw and f_maint users must have aliases, and underscore characters are not allowed on z/OS. Assigning the default f_sw and f_maint user names will result in a failure to connect to the database. See [“Prepare DB2 Server for z/OS” on page 109](#) for additional information.

NOTE The f_sw and f_maint database users will be created if the SQL scripts for SQL Server and Oracle databases are manually executed before running the Process Engine installation program.

These scripts create users, passwords and stored procedures. There are no comparable scripts for DB2 databases. See [“Process Engine SQL scripts” on page 218](#) for details on the SQL scripts.

See [“To create or modify Process Engine users and groups \(UNIX\)” on page 81](#) for specific information on how to create these accounts on your operating system.

User/Group	Description
fnadmin (group account)	Operating system group account whose members have all privileges on Process Engine files and databases
fnusr (group account)	Operating system group account whose members have non-administrative privileges on Process Engine files and databases
fnop (group account)	Operating system group account whose members have operator non-administrator privileges on Image Services utilities used by Process Engine
fnsu (user account)	Operating system user account who executes Process Engine software
f_sw (user account)	Database runtime user. For DB2 databases, this is an operating system user of DB2. For Oracle and SQL Server databases, this is a database user.
f_maint (user account)	Database maintenance user. For DB2 databases, this is an operating system user of DB2. For Oracle and SQL Server databases, this is a database user.

To create or modify Process Engine users and groups (UNIX)

1. Perform the task in the Action column in the following table for the required Process Engine users and groups listed. Then log out and log back in to ensure the changes are picked up.

(AIX only) Because AIX does not allow users to create a group with an empty member list, you must create these groups and users in the following order:

- Create the fnusr or alias group with root and Oracle User as its members.
- Create the fnsu or alias user with fnusr as its primary group.

- Create the fnadmin group with fnsw and root as its members.

User Name	Action	User Type	Default Shell	Primary Group	Secondary Group
fnsw or an alias	Create	Operating System	Ksh	fnusr	Oracle Database Administrators Group; default = dba, fnadmin, fnop
root	Modify	Operating System	Ksh		fnusr, fnadmin
Oracle User; default = oracle	Modify, if database is local Create, if database is remote	Operating System		Oracle Database Administrators Group; default = dba	fnusr

Group Name	Action	Description	Members
fnadmin or an alias	Create	Members have all privileges on Process Engine files and databases.	fnsw, root
fnusr or an alias	Create	Members have non-administrator privileges on Process Engine files and databases.	fnsw, root, Oracle User; default = oracle
fnop or an alias	Create	Members have operator non-administrator privileges on Image Services utilities.	fnsw
Oracle Database Administrators Group; default = dba	Modify, if database is local Create, if database is remote	Members act as database administrators.	fnsw, Oracle User; default = oracle

To create Oracle accounts for Process Engine (UNIX)

1. If the Oracle database is local to Process Engine, ensure that the following operating system user and group already exist as a result of installing the Oracle software. The Process Engine installation program will prompt for the values.

If the database is remote, define the user and group on the server where Process Engine will be installed if it hasn't been done already.

In either case, once the Oracle user and group have been created, modify the account information as described.

The Process Engine installation program prompts for the user and group names but does not allow assignment of aliases for them.

User/Group	Description
Oracle Database Administrators Group; default = dba	Members act as database administrators Action: Modify, if database is local. Create, if database is remote. Members: fnsr, Oracle user; default = oracle

User/Group	Description
Oracle User; default = oracle	Action: Modify, if database is local. Create, if database is remote. User Type: Database Primary Group: Oracle Database Administrators Group; default = dba Secondary Group: fnusr

To create other Process Engine accounts

1. Create new (or designate existing) directory server accounts for Process Engine, as shown in the following table:

User/Group	Description
Process Engine service user <i>pe_service_user</i>	Process Engine uses the service username when connecting the Content Engine server. This user must belong to the Process Engine Administrator group

User/Group	Description
Process Engine administrators group <i>pe_admin_group</i>	Members of this group automatically have administrative privileges for Process Engine.
Process Engine configuration group <i>pe_config_group</i>	<p>(Optional) A valid group name. Members of this group automatically have configuration privileges for the PE workflow database.</p> <p>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.</p>

Accounts for Application Engine or Workplace XT

Task to be performed by: Application Server Administrator

To create Application Engine or Workplace XT accounts

1. Create new or designate existing directory server accounts for Application Engine or Workplace XT, as shown in the following table:

User/Group	Description
Application Engine or Workplace XT installer account (Windows) <i>ae_install_user</i> <i>wpxt_install_user</i>	The account you will use to log on to a Windows machine and launch the Application Engine or Workplace XT installation program. This account must be a Windows Local administrator or a user with equivalent permissions.
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 access to the /bin directory and 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 may be the same as the Application Engine or Workplace XT installer account.
Application Engine Administrators <i>ae_admin_user</i> <i>wpxt_admin_user</i>	These accounts will serve in the role of Application Engine administrator. The role applies to both Application Engine and Workplace XT. You will specify these accounts as members of the AE administrator role when you set bootstrap preferences. These accounts must have passwords.

In addition to the requirements above, the Application Engine or Workplace XT installer account and Application server account need the read/write permissions to these directories:

WebSphere

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

WebLogic 9.x

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

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

WebLogic 10.x

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

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

JBoss

JBoss_home/bin/run.sh or run.bat

JBOSS_home/server/default/conf/login-config.xml (on both CE and AE servers)

NOTE 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.

Accounts for Content Search Engine

Task to be performed by: IT Administrator

To create Content Search Engine accounts

1. If you are installing Content Search Engine, create new (or designate existing) Autonomy K2 security accounts as shown in the following table:

User/Group	Description
K2 Security Group <i>k2_sec_group</i>	Autonomy K2 security group used to secure K2 collections. You will specify this group in the User Group field in the Verity Domain Configuration when you configure CBR in Enterprise Manager.
K2 Security User <i>k2_sec_user</i>	Autonomy K2 security user account, used when logging on to perform Content-Based Retrieval (CBR). 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 K2 Security Group and must be defined as an authorized K2 administrator in the K2 dashboard.
K2 Operating System User <i>k2_os_user</i>	Autonomy K2 services will run as this user. This user must be an operating system administrator on the machine where the Autonomy K2 Master Administration server is installed. Additionally, this user must have access to the file system that contains the file storage areas and the full text index collections, because 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.

NOTE Both K2 Security User and K2 Operating System User can be the same user. All permissions listed above must be assigned.

Database Administrator tasks

As the Database Administrator, perform the following tasks to prepare the required database environment for IBM FileNet P8 Platform:

- Review all rows assigned to the Database Administrator (DBA) in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. (Your organization may have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.)

HINT 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 the Role you are interested in.
- 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).
- For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- If you are installing in a non-English environment, review the considerations and procedures in [“Installing P8 Platform in a non-English environment” on page 248](#) before you begin your preparation tasks.
- Create the database accounts required by FileNet P8, which are listed for the Database Administrator in the following:
 - [“Task to be performed by: Database Administrator” on page 66](#)
- Set up your database server for Content Engine and Process Engine, depending on the database type required by your installation:
 - [“Prepare Microsoft SQL Server” on page 89](#)
 - [“Prepare Oracle Server” on page 97](#)
 - [“Prepare DB2 Server” on page 78](#)
- Set up your database client software for Process Engine, depending on the database type for your installation:
 - [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#)
 - [“Configure Oracle Client for Process Engine” on page 106](#)
 - [“Configure DB2 Client” on page 121](#)

Prepare Microsoft SQL Server

Microsoft SQL Server planning considerations

- **Determine whether you want to use a dedicated or shared database.** In this regard:
 - Content Engine, Process Engine, and Rendition Engine can share a database engine, or they can each have a dedicated (unique) database engine. The Process Analyzer expansion product must have its own database engine.
 - Content Engine, Process Engine, and Rendition Engine components can each have a dedicated SQL Server database instance, or they can share a database instance with one another or with non-IBM FileNet applications. You can use the default instance or a named instance of your choosing.
 - Content Engine, Process Engine, and Rendition Engine must each have their own databases, even if they share a database engine.
 - Each Process Engine isolated region configured for recovery must have dedicated filegroups.
- **Determine when to execute SQL Server scripts for Process Engine.** A number of SQL scripts must be executed that create a number of stored procedures. These stored procedures will then be used to create database users and other stored procedures required for Process Engine production.

These scripts can be executed manually, before starting Process Engine installation, or executed from the Process Engine installation program. See [“Process Engine SQL scripts” on page 218](#) for information on execution modes and associated security requirements as well as details about the scripts.

The Process Engine installation program will complete only if all these scripts run successfully.

- **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.

NOTE 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.

Verify that Microsoft SQL Server is installed for IBM FileNet P8

The procedures in this task describe how to install and configure a SQL Server database that is dedicated or shared by one or more of the following IBM FileNet P8 components:

- dedicated to Content Engine
- dedicated to Process Engine

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

For a SQL Server database dedicated to Rendition Engine, see the IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

In a shared configuration, the IBM FileNet P8 components use the same database instance. You can also share the database with other non-IBM FileNet applications. In a dedicated configuration, Content Engine, Process Engine, and Rendition Engine use separate database instances.

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

NOTES

- Process Engine requires an ODBC data source for connection to the SQL Server database. See [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#) for details on configuring the data source.
- If your Microsoft SQL Server database will be remote, then after completing the procedures in this topic you might also need to install SQL Server Client software on the machine where you will install Process Engine. See [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#) for details.
- The following components can optionally share a SQL Server instance:
 - Content Engine (object store databases and GCD database)
 - Process Engine
 - Rendition Engine
- A database should not be shared by multiple Content Engine object stores.
- Process Engine and the optional Process Analyzer cannot share an instance. For information on SQL Server requirements for Process Analyzer Engine, see the *Process Analyzer Installation and Upgrade Guide*.
- 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” on page 216](#), as appropriate. Be aware that the Process Engine installation program allows only alphanumeric characters and underscores.
 - Server name
 - Instance name (for example, P8_inst)
 - Dedicated database name (for example, VWdb)
 - Default file group (for example, vwdata_fg)
 - TCP/IP port number assigned

Get Microsoft SQL Server ready for IBM FileNet P8

To install and configure Microsoft SQL Server

1. Create a database instance for use by IBM FileNet P8 software, or verify that such an instance already exists.
2. If creating a new instance, indicate an appropriate name based on whether Content Engine (object store), GCD, Process Engine, or Rendition Engine will use the instance.
3. Verify the authentication mode you specify is for Mixed Mode.
4. Select the database collation settings. Specify one of the following:
 - **Dictionary order, case-insensitive, for use with 1252 Character Set** (or any case-insensitive MS SQL Server collation). Case-insensitive collation is the Microsoft default and the setting most used in IBM FileNet P8 environments (because it offers search results without regard to character case).
 - **Dictionary order, case-sensitive, for use with 1252 Character Set** (or any case-sensitive MS SQL Server collation). Select case-sensitive MS SQL Server collation only if you are sure your site actually requires (and will continue to require) searches that must differentiate upper-case from lower-case characters (in property choice lists, folder names, etc.). If you plan to use the Content Engine with CFS/IS, you must configure case-sensitive. The Image Services database is configured as case-sensitive and the Content Engine database must match.

CAUTION Select your MS SQL Server collation setting carefully. Switching collation settings after installation can be difficult and time-consuming, especially if you want to switch from case-sensitive to case-insensitive collation after significant user activity. Also, be aware that if you have a case-sensitive database, and you want to perform a case-insensitive search (programmatically or otherwise), you will likely encounter serious performance degradation on MS SQL Server because the database cannot use column (that is, property) indexes in these cases.

5. (MS SQL Server 2005 only) The TCP/IP port number cannot be assigned during installation of SQL Server 2005. You must assign the port number after installation is complete by using the SQL Server Configuration Manager application to modify the network configuration.
6. Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. Verify that the required service pack has been installed before proceeding. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
7. If you want to disable the Named Pipes protocol for the database instance to be used by Process Engine, wait until after Process Engine installation and configuration is complete. Disabling this protocol too early might cause Process Engine initialization to fail.

WARNING The default on a SQL Server 2005 installation is to disable Named Pipes. Use the SQL Server Configuration Manager application to modify this network configuration parameter after SQL Server 2005 is installed and the instance has been created.

To create a SQL Server database for the GCD

Create a SQL Server database for GCD database, which is required for Content Engine installation. Create the database with an initial size of 100MB, minimum. Make note of the database name as it will be required later when installing Content Engine software.

NOTE Record the database name in the [“Installation and upgrade worksheet” on page 216](#).

To create a MS SQL Server database for Content Engine object stores

Create a MS SQL Server database for a Content Engine object store. Each object store you create will require its own, empty database. Create the database with an initial size of 200MB, minimum.

NOTE Record this value in the [“Installation and upgrade worksheet” on page 216](#).

To create the Process Engine database

Create a MS SQL Server database for Process Engine. The default name assigned by the Process Engine installation program is VWdb. Create the database with an initial size of 200MB, minimum. Assign a new file name for the database (for example, vw_data). Specify a filegroup (for example, vwdata_fg.) IBM recommends that the Primary filegroup is not used. If you will be using the region recovery feature, create an additional filegroup for every region configured for recovery.

NOTE You will need the default database and filegroup names for the Process Engine installation. Record the values in the [“Installation and upgrade worksheet” on page 216](#).

To modify the tempdb database for Process Engine

Verify that the space allocated for the tempdb is at least 80MB.

To execute Process Engine pre-installation scripts manually

The Process Engine installation program runs several pre-installation SQL scripts to create Process Engine users and passwords and stored procedures. These scripts for MS SQL Server databases can be executed in one of three ways.

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

See [“Process Engine SQL scripts” on page 218](#) for detailed information on the scripts and modes of execution.

To enable XA Transactions

Perform these steps on every MS SQL Server that will contain a Content Engine database.

1. From Control Panel, open Administrative Tools, and then open Component Services.
2. Expand Component Services, right-click **My Computer**, and then select Properties.
3. Click the **MSDTC** tab, and then click **Security Configuration**.
4. Select the *Enable XA Transactions* check box, and then click **OK**. This will restart the MS DTC service.
5. Click **OK** again to close the Properties dialog box, and then close Component Services.
6. Stop and then restart the MS SQL Server.

To configure the JDBC Distributed Transaction Components

Execute these steps on every SQL Server that will have a Content Engine database.

1. Download the Microsoft SQL Server 2005 JDBC Driver that is referenced in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19](#).
2. Copy the sqljdbc_xa.dll from the JDBC installation directory to the Program Files\Microsoft SQL Server\80\Tools\Binn directory if there is only a default instance or the Program Files\Microsoft SQL Server\MSSQL\$*instance name*\Binn if you are using a named instance. If you are on a 32-bit processor, use the sqljdbc_xa.dll file in the x86 folder. If you are on a 64-bit processor, use the sqljdbc_xa.dll file in the x64 folder.
3. Log on as a database administrator and execute the database script xa_install.sql on every SQL Server instance that will participate in distributed transactions. This script installs sqljdbc_xa.dll as an extended stored procedure and creates the SqlJDBCXAUser role in the Master database.
CAUTION Use SQL Server database credentials, not Windows credentials, to log on. Windows Integrated Logon to SQL Server is not supported with IBM FileNet P8.
4. Navigate to **Component Services > My Computer > Properties > , MSDTC > Security Configuration**, and make sure that XA transactions are enabled.

Configure Microsoft SQL Server Client for Process Engine

Use the procedures in this topic to:

- create the ODBC data sources required by Process Engine.
- install SQL Server Client software for Process Engine, if required.

NOTES

Only Process Engine has a possible requirement for SQL Server Client software installation. You must execute a number of SQL scripts for the SQL Server database to create stored procedures. Those scripts can either be executed manually on the SQL Server database before running the Process Engine installation program, or they can be executed from the Process Engine installation program.

- If the scripts are executed from the Process Engine installation program, SQL Server Client software must be installed on the Process Engine server.
- If the scripts are executed from the Process Engine installation program on a 64-bit Windows 2008 server, a 64-bit ODBC data source must be created, in addition to the 32-bit ODBC data source required for Process Engine run time software.
- If the scripts are executed manually on the remote database server before you install Process Engine, you need not install SQL Server Client. If SQL Server Client software is installed on the Process Engine Server, that software can be removed after successful installation and configuration of Process Engine software.

See [“Process Engine SQL scripts” on page 218](#) for details on execution of the SQL Scripts.

- You can install Microsoft SQL Server and Process Engine servers in different Active Directory forests. In this case, however, you must use SQL Authentication rather than Windows Authentication.

To create the Process Engine ODBC data source and test the connection

SQL Server connections for Process Engine at run time are handled through an ODBC data source. The 32-bit ODBC data source is required for both local and remote databases and must be created on the Process Engine server. The following steps apply on a 32-bit version of the operating system.

NOTE On a 64-bit operating system locate and manually execute the 32-bit version, typically located in: C:\WINDOWS\SysWOW64\odbcad32.exe.

1. Start **Program > Administrator Tools > Data Source (ODBC)**.
2. Click **Add** on the **System DSN** tab.
3. Select **SQL Server** as the driver to use for the new data source and click **Finish**.
4. Enter a name and description for the data source. The name will be required input for the Process Engine installation program when configuring for a SQL Server database.
5. Choose the SQL Server to connect to from the drop-down list of servers and click **Next**.

NOTE If only a server name appears in the list, the connection will be with the default instance. If there are named instances in the database, the name will appear as *server/instance name*.

6. Do the following, and then click **Next**:
 - a. Choose SQL Server authentication.
 - b. Select the option to get default settings for additional configuration options by connecting to the SQL Server.
 - c. Indicate the Login ID and Password to connect to the database.

NOTE This database login ID information need not be for an administrator and it is only used to connect to the database to get the default values for the remaining settings required to configure the data source.

7. Change the default database to be the Process Engine database created earlier in [“Verify that Microsoft SQL Server is installed for IBM FileNet P8” on page 89](#).
8. Turn on **Use ANSI null, paddings, and warning** and turn on **Use ANSI quoted identifiers**.
9. Turn on **Perform translation for character data** and click **Finish**.
10. Verify the settings for the data source configuration and click **Test Data Source**. If the test is successful click **OK**. Otherwise resolve any problem before continuing.
11. Double-click **SQL Server** on the **Connection Pooling** tab.
12. Select **Don't pool connection to this driver** and click **OK**.
13. Click **OK** on the ODBC Data Source Administrator window to finish configuration of the data source.

On the summary screen click **Test Data Source**. If error messages display, resolve them before proceeding.

To install SQL Server Client software for remote database access

Ensure that all users and groups defined on the local server are also defined and granted security permissions on the database server.

Install SQL Server Client software

1. Log on with an account that has local administrator privileges on the computer where the SQL Server client software will be installed.
2. (SQL Server 2000 only) Install the SQL Server Client Tools Only.
3. (SQL Server 2005 only) Install “Workstation components, Books Online and development tools” and from the Advanced options, select **Client Components**.
4. When the installation is complete, start the “Client Network Utility” and clear the “Automatic ANSI to OEM conversion” checkbox on the DB-Library Options tab.
5. Test the database connection.

Install SQL Server patches and service packs

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#). Verify that the required service pack has been installed before proceeding.

Prepare Oracle Server

Use the procedures in this topic to plan and prepare your Oracle servers.

Oracle Planning Considerations

General

- **Determine whether you want to use a dedicated or shared database.** In this regard:
 - Content Engine, Process Engine, and Rendition Engine can share a database engine, or they can each have a dedicated (unique) database engine.
 - Content Engine, Process Engine, and Rendition Engine can each have a dedicated Oracle database instance, or they can share a database instance with one another or with non-IBM FileNet applications.
 - Content Engine, Process Engine, and Rendition Engine must each have their own tablespaces, even if they share a database engine.
 - Each Process Engine isolated region configured for region recovery must have dedicated tablespaces.

NOTE For detailed information regarding installation of Rendition Engine, see IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

- **Plan to use locally managed tablespaces.** For performance reasons, IBM recommends that you create locally managed, rather than dictionary managed, tablespaces for Process Engine and Content Engine. (The tablespaces you create via Oracle Database assistant (dbca) are locally managed by default.)
- **Be aware of database client software requirements.** For Process Engine, if the database is remote, you must install database client software on the Process Engine server.
- An Oracle database must be remote if it is installed on a Linux server.

Process Engine

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

These scripts can be executed manually, before starting Process Engine installation, or executed from the Process Engine installation program. See [“Process Engine SQL scripts” on page 218](#) for information on execution modes and associated security requirements as well as details about the scripts.

The Process Engine installation program will complete only if all these scripts run successfully.

Verify that Oracle Server is installed for IBM FileNet P8

The procedures in this task describe how to install and configure an Oracle database that is dedicated or shared by one or more of the following IBM FileNet P8 components:

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

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

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

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

For information regarding installation of Oracle Server and Rendition Engine, see the IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

NOTES

- If your Oracle database will be remote, then after completing the procedures in this topic you must also complete the procedures in [“Configure Oracle Client for Process Engine” on page 106](#) on each machine where you will install Process Engine. For Rendition Engine instructions, see the IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).
- Make sure the machine that will host the database satisfies all pre-installation requirements specified in the Oracle9i or Oracle 10g installation documentation.
- For Content Engine and Process Engine, IBM FileNet P8 supports the Oracle Advanced Security functionality of secure data transfer across network protocol boundaries.
- If you will be installing Process Engine on a UNIX machine hosting the Oracle database, be sure that the value of the Oracle environment variable ORACLE_HOME (the path name for the Oracle Server software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installer will not find the Oracle software, causing the installation to fail.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#). The Oracle patches are available at the Oracle website. The Oracle patch-installation procedure may be less complicated if done before you create any databases.

- Transaction Processing is the required configuration type for the database that supports Content Engine. Choose this configuration type if your database will be dedicated to Content Engine or shared with Process Engine.
- Record the values for the following settings as you work through the database installation. Enter this information in the PE installer or CE installer sections of the [“Installation and upgrade worksheet” on page 216](#), as appropriate. This information must be entered during subsequent installations. Be aware that the Process Engine installation program allows only alphanumeric characters and underscores.
 - Oracle Home
 - Global Database Name
 - Oracle temporary tablespace name
 - Oracle data tablespace name
 - Oracle index tablespace name (optional)
 - Oracle SID

To install an Oracle database engine

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

- Choose the following from the list of available product components.
 - Oracle Server
 - Oracle Net Services
 - Oracle Net Listener
 - Oracle Development Kit
 - Oracle Call Interface (OCI)
 - (Windows) Oracle Windows Interfaces
 - Oracle Services for Microsoft Transaction Server
 - Oracle Documentation (recommended)
- If you are going to install Process Engine on this machine, verify/add/edit/uncomment the following lines in the file `sqlnet.ora` (create the file if it doesn't exist) while the Oracle services/processes are stopped:

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

NOTE If Oracle is configured to use LDAP, TNSNAMES must appear in the `names_directory_path` ahead of LDAPNAMES.

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

1. Install the latest Oracle patch sets, as specified in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
2. Start the listener and the Oracle database service/processes (Windows/UNIX) if they haven't started automatically.

To create an Oracle database

Oracle documentation describes several ways to create a database. IBM recommends you use the Database Configuration Assistant (DBCA).

IBM FileNet requires the following settings:

- Database configuration type
If this database is dedicated to Content Engine, or if it will be shared by Content Engine and Process Engine, then Transaction Processing (also known as OLTP) is the required configuration type.
- Server process type
Dedicated Server Mode
- Database character set
Choose a database character set as specified in [“IBM FileNet P8 database character sets” on page 262](#).

To create tablespaces for the GCD

Using Oracle Enterprise Manager or SQL*Plus, create a user, password, and default tablespace in the Oracle database for the GCD that Content Engine will access. Grant CONNECT and RESOURCE roles to the user. These two roles combine to include the minimal privileges required by Content Engine: CREATE SESSION, CREATE TABLE, and CREATE SEQUENCE.

NOTE Because these two roles include other privileges as well, IBM recommends that you design your own roles if you prefer to grant only the minimal privileges required by the GCD.

Grant the following additional permission to the user:

- select on pending_trans\$
- select on dba_2pc_pending
- select on dba_pending_transactions
- execute on dbms_system

WARNING The Oracle user you create for the permanent and temporary tablespaces of the GCD must be unique. That is, the Oracle user for the GCD must not be the same as that of the user for any object store. Otherwise, the objects you intend to add only to the GCD will show up in all object stores that share the same Oracle user.

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

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

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

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

NOTE Record the tablespace name in the [“Installation and upgrade worksheet” on page 216](#).

To create Oracle tablespaces for Content Engine object stores

Using Oracle Enterprise Manager or SQL*Plus, create a user, password, and default tablespace in the Oracle database for each object store that Content Engine will access. Grant CONNECT and RESOURCE roles to the user. These two roles combine to include the minimal privileges required by Content Engine: CREATE SESSION, CREATE TABLE, and CREATE SEQUENCE.

NOTE Because these two roles include other privileges as well, IBM recommends that you design your own roles if you prefer to grant only the minimal privileges required by Content Engine.

Grant the following additional permission to the user:

- select on pending_trans\$
- select on dba_2pc_pending
- select on dba_pending_transactions
- execute on dbms_system

WARNING The Oracle user you create for the permanent and temporary tablespaces of an object store must be unique. That is, multiple object stores and the GCD must *not* share the same Oracle user. Otherwise, the objects you intend to add only to one object store will show up in all object stores (and the GCD) that share the same Oracle user.

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

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

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

Tablespace Name	Tablespace Type	Minimum Size	Description
<i>objectstore1</i>	Permanent	200 MB	Permanent tablespace for object store
<i>tempobjectstore1</i>	Temporary	2 GB	Temporary tablespace for object store

To create tablespaces for Process Engine

Using Oracle Enterprise Manager or SQL*Plus, create the tablespaces shown in the following table for the Process Engine. Note that the indexing tablespace (vwindex_ts) is optional. Tablespace names used by Process Engine can contain only alphanumeric and underscore characters. Names must start with an alphabetic character and must be at most 18 characters long.

If you will use the region recovery feature in Process Engine, you must create the default tablespaces and an additional data and index tablespace for each region to be configured for recovery. The same run-time and maintenance users will be used for all Process Engine tablespaces.

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

Tablespace Name	Tablespace Type	Minimum Size (MB)	Description
vwdata_ts	Permanent	200	Default name of the dedicated IBM FileNet data tablespace. This is the only data tablespace configured during Process Engine installation.
vwtemp_ts	Temporary	400	Default name of the dedicated IBM FileNet temporary tablespace. This is the only temp tablespace configured during Process Engine installation.
vwindex_ts (optional)	Permanent	200	Default name of the optional default index tablespace. This is the only index tablespace configured during Process Engine installation.
<i>region X data</i>	Permanant	200	Data tablespace to be used by an individual region configured for recovery. Cannot be shared by any other region.
<i>region X index</i>	Permanent	200	Index tablespace to be used by an individual region configured for recovery. Cannot be shared by any other region.

NOTES

- If you don't create vwindex_ts, vwdata_ts will be used for indexes.
- Only the default tablespace names need to be entered during Process Engine installation.

To set environment variables for the Oracle and root users on UNIX database server

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

- For the oracle user, set the following:
 - ORACLE_SID
 - ORACLE_HOME
 - Set PATH to:
\$ORACLE_HOME/bin (HP-UX and 32-bit Solaris only)
 - Set LD_LIBRARY_PATH to:
\$ORACLE_HOME/lib (HP-UX and 32-bit Solaris only)
 - Set LD_LIBRARY_PATH to:
\$ORACLE_HOME/lib32 (64-bit Solaris only)
 - Set LIBPATH to:
\$ORACLE_HOME/lib32:\$ORACLE_HOME/lib (AIX only)
 - Set SHLIB_PATH to:
\$ORACLE_HOME/lib32 (HP-UX only)
- For the root user, set ORACLE_HOME.

To configure automatic transaction recovery

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

To enable automatic transaction recovery, perform the tasks shown in the section "Scheduling Automatic Microsoft Transaction Server Recovery" in *Oracle Services for Microsoft Transaction Server Developer's Guide* (Oracle Part Number A95496-01).

In addition, if you are using an Oracle Fail Safe configuration, perform the procedure shown in "Modifying Registry Values for Oracle Fail Safe Configurations" in *Oracle Services for Microsoft Transaction Server Developer's Guide* (Oracle Part Number A95496-01).

To verify the listener

Verify that the listener is configured and running.

To execute Process Engine pre-installation scripts manually

The Process Engine installation program runs several pre-installation SQL scripts to create Process Engine users and passwords, tablespace defaults and privileges for the users, stored procedures, and synonyms for the stored procedures. These scripts for Oracle databases can be executed in one of three ways.

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

See [“Process Engine SQL scripts” on page 218](#) for detailed information on the scripts and modes of execution.

To turn off Oracle Password Complexity Verification

Process Engine does not support Oracle Password Complexity Verification during the installation process. If you are using the default passwords for the run-time (f_sw or alias) or maintenance (f_maint or alias) users, turn off this Oracle feature and do not re-enable it until you have installed and configured Process Engine.

Configure Oracle Client for Process Engine

Use the procedures in this topic to install the Oracle Client software to prepare for the installation of Process Engine. Content Engine does not require installation of Oracle client software. For information regarding installation of Oracle Client and Rendition Engine, see IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

NOTES

- Install Oracle Client on any machine that will host Process Engine, Rendition Engine, or any other IBM FileNet P8 component except Content Engine (such as Enterprise Manager) that needs to access an Oracle database.
- If you will be installing Process Engine on a UNIX machine hosting Oracle Client software, be sure that the value of the Oracle environment variable ORACLE_HOME (the path name for the Oracle Client software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installer will not find the Oracle software, causing the installation to fail.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating system and database patch sets, and Service Packs. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#). The Oracle patches are available at the Oracle website. The Oracle patch installation procedure may be less complicated if done before you create any databases.

To install the Oracle client

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

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

Windows

Set/modify the value of the NLS_LANG key via System Properties in the Control Panel.

UNIX

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

NOTE To affect the environment for only Process Engine, set NLS_LANG for just the fnsw user. (On Windows platforms, fnsw is created by the Process Engine installer; on UNIX platforms, you manually create fnsw as part of the Process Engine installation task.)

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

Windows

LOCAL

UNIX

TWO_TASK

5. If you are going to install Process Engine software on a UNIX machine to connect to a remote Oracle database, set the following environment variables in the startup file of the default *oracle* user.
 - ORACLE_SID
 - ORACLE_HOME
 - Set PATH to:
\$ORACLE_HOME/bin (HP-UX and 32-bit Solaris only)
 - Set LD_LIBRARY_PATH to:
\$ORACLE_HOME/lib (HP-UX and 32-bit Solaris only)
 - Set LD_LIBRARY_PATH to:
\$ORACLE_HOME/lib32 (64-bit Solaris only)
 - Set LD_LIBRARY_PATH_64 to:
\$ORACLE_HOME/lib (64-bit Solaris only)
 - Set LIBPATH to:
\$ORACLE_HOME/lib32:\$ORACLE_HOME/lib (AIX only)
 - Set SHLIB_PATH to:
\$ORACLE_HOME/lib32 (HP-UX only)

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

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

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

7. Install all required Oracle patches, as specified in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#). These patches are available at the Oracle website.

Prepare DB2 Server for z/OS

Use the procedures in this topic to plan and prepare your DB2 Server for z/OS servers.

DB2 for z/OS planning considerations

- **The database must be remote.** The DB2 for z/OS database must be remote from Process Engine and Content Engine. Process Engine and Content Engine do not run on z/OS. See the *IBM FileNet P8 Hardware and Software Requirements* for specifics on supported software versions and platforms. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.](#)
- **Determine whether you want to use a dedicated or shared database resource.** In this regard:
 - Content Engine and Process Engine can share a database engine, or they can each have a dedicated (i.e., 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 applications.
 - Each Content Engine and Process Engine must have their own databases and tablespaces.
- **Assign unique databases to the Content Engine global configuration data (GCD) and object stores.** It is a best practice to create a separate database for the GCD and each object store for ease of maintenance, security control, and the like.
- **Use System Managed Storage (SMS).** Configure system managed storage for Process Engine and Content Engine databases.
- **Use UTF-8 collation.** For both Content Engine and Process Engine, use UTF-8 collation settings by configuring CCSID UNICODE.
- **Add SDSNLOD2 into the LNKLIST.** The SDSNLOD2 library must be added into the LNKLIST when z/OS is configured.
- **Use TCP/IP as the default protocol.** For both Content Engine and Process Engine, set TCP/IP as the default network protocol.
- **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.

NOTE 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.

Verify that DB2 for z/OS Server is installed for IBM FileNet P8

The procedures in this task describe the requirements for a DB2 instance that is dedicated or shared by one or more of the following IBM FileNet P8 components. In a shared configuration, the

IBM FileNet P8 components use the same instance, but different databases. You can also share the instance with other (non-IBM FileNet) applications.

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

Record the values for the following settings as you work through the database installation. This information must be entered during subsequent installations of Process Engine and Content Engine. Be aware that z/OS allows only alphanumeric characters.

- DB2 Server name
Record both the TCP/IP address and the fully qualified domain name
- DB2 server database instance name(s)
- Content Engine and Process Engine dedicated database names
- DB2 instance port numbers
- Process Engine run-time database user (fsw or alias) password (Process Engine only)
- Process Engine maintenance database user (fmaint or alias) password (Process Engine only)
- User IDs and passwords for Content Engine DB2 users

These are operating system users who have been granted permissions on the database.

To create DB2 users for Content Engine

Content Engine requires a separate user for the GCD and every object store. Create new operating system users or identify existing operating system users to grant permissions for. These user IDs and passwords will be required later when creating connection pools, JDBC data sources, and object stores.

The permissions required for each user are:

```
GRANT SYSADM TO operatingsystemuser;
```

```
GRANT USE OF STOGROUP storagegroupname TO operatingsystemuser;
```

```
GRANT USE OF BUFFERPOOL BP32K TO operatingsystemuser;
```

```
GRANT SELECT ON SYSCAT.DATATYPES TO operatingsystemusername;
```

```
GRANT SELECT ON SYSIBM.SYSVERSIONS TO operatingsystemusername;
```

where:

operatingsystemuser is the operating system user name

storagegroupname is the name of the storage group

To determine 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, and it also affects 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.
- 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 8KB minimum pagesizes for Process Engine databases and 32KB minimum pagesizes for Content Engine databases.

To create and update the DB2 databases for Content Engine and Process Engine

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.

To create users for Process Engine

Create the following two operating system users on the database server. Note that underscores are not legal characters in these names so the default f_sw and f_maint users cannot be used.

- Process Engine runtime database user (fsw or other alias)
 - Primary runtime Process Engine database user.
 - Used only by the Process Engine software to access the DB2 database.
 - Must have DBADM authority of the DB2 instance that will be used by the Process Engine software.
- Process Engine maintenance database user (fmaint or other alias)
 - Process Engine Database maintenance user.
 - Mainly used by the customer for database maintenance.
 - Recommended to have DBADM authority of the DB2 instance that will be used by the Process Engine software.

The user names and passwords will be required input to the Process Engine installation program.

The permissions required for each user are:

```
GRANT DBADM ON DATABASE databasename TO operatingsystemuser;  
GRANT USE OF STOGROUP storagegroupname TO operatingsystemuser;  
GRANT USE OF BUFFERPOOL BP8K TO operatingsystemuser;
```


Prepare DB2 Server for Linux, UNIX and Windows

Use the procedures in this topic to plan and prepare your DB2 Server for Linux, UNIX and Windows servers.

DB2 for Linux, UNIX and Windows planning considerations

- **DB2 server instances must be 64-bit on Unix, 32-bit on Windows.** Instances created on a remote UNIX database server must be 64-bit. On Windows database servers, instances must be 32-bit.
- **DB2 client instances must be 32-bit.** Instances created on the Process Engine server must be 32-bit.
- **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 (i.e., 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 applications.
 - Each Content Engine and Process Engine must have their own databases and tablespaces.
 - The Content Engine global configuration data (GCD) and each object store must have its own database.
 - Each Process Engine isolated region configured for region recovery must reside in a dedicated data, blob, and index tablespace, separate from the default tablespaces.
- **Plan to use Database Managed Storage(DMS) for tablespaces.** For performance reasons, IBM recommends that you create database managed user and user temp tablespaces rather than 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

- **IBM recommends 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.

NOTE 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.

Verify that DB2 Server is installed for IBM FileNet P8

The procedures in this task describe how to install and configure a DB2 instance that is dedicated or shared by one or more of the following IBM FileNet P8 components. You can also share the instance with other (non-IBM FileNet) applications.

- dedicated to Content Engine
- dedicated to Process Engine
- shared by two or more of Content Engine and 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 [“Configure DB2 Client” on page 121](#).

NOTES

- Content Engine and Process Engine have different requirements for DB2 users and groups.

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*. See the appendix [“Installation and upgrade worksheet” on page 216](#). 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 name(s) (e.g. P8inst)
- Content Engine and Process Engine dedicated database names (e.g. VWdb)
- Dedicated tablespace names (e.g. vwdata)
- DB2 instance port numbers
- Process Engine run-time database user (f_sw or alias) password (Process Engine only)
- Process Engine maintenance database user (f_maint or alias) password (Process Engine only)
- User ID and password for Content Engine DB2 user

To create DB2 users and groups

You must create operating system users and groups to function as instance owners and a primary group for the instance owner for both Content Engine and Process Engine. A fenced user and a primary group for the fenced user must be created for Process Engine. You may specify your own user and group names as long as they adhere to system naming rules and DB2 naming rules.

- Instance Owner - one is required if the database instance is being shared by Content Engine and Process Engine, two instance owners are required if Content Engine and Process Engine will have separate database instances.
- Instance Owner primary group - one is required if the database instance is being shared by Content Engine and Process Engine, two instance owner primary groups are required if Content Engine and Process Engine will have separate database instances.

The *instance owner* home directory is where the DB2 instance will be created.

- Fenced User
- Fenced User primary group

The *fenced user* runs user-defined functions (UDFs) and stored procedures outside the address space used by the DB2 database.

Each instance should have its own home file system.

Each *instance owner* must have a unique home directory. All of the files necessary to run the instance are created in the home directory of the instance owner's user ID/username.

The *instance owner* and its primary group are associated with every instance. The *instance owner* is assigned during the process of creating the instance.

The primary group of the *instance owner* automatically becomes the system administration group for the instance and gains SYSADM authority as a database administrator (DBA) over the instance. Other user IDs or user names that are members of the primary group of the *instance owner* also gain this level of authority.

NOTE The root user cannot act as a DBA. You must log on as the *instance owner* to act as the DBA.

When the database authentication type is set to SERVER or SERVER_ENCRYPT, two additional operating system users must be created for Process Engine on the database server where the DB2 database resides. See [“To create additional OS users and groups \(PE, SERVER or SERVER_ENCRYPT authentication only\)”](#) on page 119.

Content Engine requires a separate user for every tablespace if object stores are going to be created within the same database. Create a new operating system user or identify an existing operating system user to grant DB2 permissions for. This user ID and password will be required later when creating connection pools and Object Stores. The DB2 permissions required for this user are:

- Permission to connect
- Permission to create tables in the tablespace (CREATETAB)
- Permission to use the tablespace (USE OF) (for User and User Temp tablespaces)

To install DB2 ESE

Install the DB2 UDB Enterprise Edition. Content Engine and Process Engine both need 64-bit instances on a Unix servers, 32-bit instances on Windows servers, and a single instance can be shared.

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. Continue with the next section.

To create DB2 instances

Content Engine and Process Engine can share an instance, or each engine can have its own instance. Both Content Engine and Process Engine require 64-bit instances on Unix servers, 32-bit instances on Windows servers. Create the appropriate instances if they don't exist.

To set the optprofile value (Content Engine)

Set the DB2 option DB2_OPTPROFILE=YES. The default value is NO.

To set TCP/IP as the default protocol

Log on as the instance owner and set DB2COMM by executing:

```
db2set DB2COMM=tcPIP
```

To determine 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, and it also affects 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.
- 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 8KB minimum pagesizes for Process Engine databases and 32KB minimum pagesizes for Content Engine databases.

To create and update the DB2 databases for Content Engine and Process Engine

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. For example, VWdb is the default database name in the Process Engine installer.

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

To create the DB2 tablespaces

If you will use the region recovery feature in Process Engine, you must create the default data tablespace and an additional data, blob, and index tablespace for each region to be configured for recovery. If region recovery is not enabled, only a single data tablespace is required for Process Engine. Separate Index and blob tablespaces are not supported for the default configuration. Only the default data tablespace name will be required input to the Process Engine installation program.

IBM FileNet Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
user temporary ts(for PE)		40		must match pagesize of PE data tablespace
system temporary (for PE)		40		must match pagesize of PE data tablespace
vwdata_ts (for PE) NOTE This is the only tablespace name entered during Process Engine installation.		200		8 (recommended)
<i>region X data</i> (Data tablespace to be used by an individual PE region configured for recovery. Cannot be shared by any other region.)		200		8 (recommended)
<i>region X index</i> (Index tablespace to be used by an individual PE region configured for recovery. Cannot be shared by any other region)		200		8 (recommended)

IBM FileNet Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
<i>region X blob</i> (Blob tablespace to be used by an individual PE region configured for recovery. Cannot be shared by any other region)		200		8 (recommended)
GCD_ts (for the GCD database)		256		32 (required)
cedata_ts (for a single CE object store)		512		32 (required)
user temporary ts (for CE)		40		32 (required)
system temporary ts (for CE)		40		32 (required)

To create additional OS users and groups (PE, SERVER or SERVER_ENCRYPT authentication only)

Create the following two additional operating system users with SYSADM authority to access the DB2 database.

- Process Engine runtime database user (f_sw or alias)
 - Primary runtime Process Engine database user.
 - Used only by the Process Engine software to access the DB2 database.
 - Must be a member of operating system group having SYSADM authority of the DB2 instance that will be used by the Process Engine software.
- Process Engine maintenance database user (f_maint or alias)
 - Process Engine Database maintenance user.
 - Mainly used by the customer for database maintenance.

- Recommended to be member of group having SYSADM authority of the DB2 instance that will be used by the Process Engine software.

Unlike the instance owner user, these users don't need to have separate file systems for their home directories. They must belong to the primary group of the instance owner.

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.

The user names and passwords will be required input to the Process Engine installation program.

To grant database permissions for f_sw and f_maint (PE, CLIENT authentication only)

If you will be using DB2 client authentication, the Process Engine runtime and maintenance database users do not exist on the database server but will need to be created on the client. These two users will need to be granted the following permission in the DB2 database for Process Engine.

- Process Engine runtime database user
 - Primary runtime PE database user.
 - Used only by the PE software to access the DB2 database.
 - Must be granted the following permissions in the DB2 database:
 - Connect
 - Createtab
 - Bindadd
 - db2set DB2_SNAPSHOT_NOAUTH=on
- Process Engine maintenance database user
 - Process Engine Database maintenance user.
 - Mainly used by the customer for database maintenance.
 - Must be granted the following permissions in the DB2 database:
 - dbadm
 - db2set DB2_SNAPSHOT_NOAUTH=on

Configure DB2 Client

Use the procedures in this topic that are appropriate to your configuration.

Content Engine

- DB2 for Linux, UNIX, and Windows database
There is no configuration required.
- DB2 for z/OS database
[“To install the z/OS license and modify the classpath \(z/OS only, Content Engine and Process Engine\)” on page 124](#)

Process Engine

- DB2 for Linux, UNIX, and Windows database.
[“To install DB2 client software \(Process Engine\)” on page 122](#)
[“To create DB2 users and groups \(Process Engine, UNIX only\)” on page 122](#)
[“To create DB2 Client instances for Process Engine \(UNIX\)” on page 123](#)
[“To create additional DB2 users and groups \(Process Engine CLIENT authentication only\)” on page 123](#)
[“To catalog the DB2 server node” on page 123](#)
[“To create the Process Engine DB2 database alias” on page 123](#)
[“To verify the connection to the DB2 database” on page 124](#)
[“To verify the ability to log on to the DB2 database” on page 124](#)
- DB2 for z/OS database.
[“To install DB2 client software \(Process Engine\)” on page 122](#)
[“To create DB2 users and groups \(Process Engine, UNIX only\)” on page 122](#)
[“To create DB2 Client instances for Process Engine \(UNIX\)” on page 123](#)
[“To create additional DB2 users and groups \(Process Engine CLIENT authentication only\)” on page 123](#)
[“To catalog the DB2 server node” on page 123](#)
[“To create the Process Engine DB2 database alias” on page 123](#)
[“To install the z/OS license and modify the classpath \(z/OS only, Content Engine and Process Engine\)” on page 124](#)
[“To verify the connection to the DB2 database” on page 124](#)
[“To verify the ability to log on to the DB2 database” on page 124](#)

Before starting these procedures for Process Engine you will need the following information:

- The database name for the Process Engine database.

- DB2 instance port number from the database server.
- Process Engine runtime database user (f_sw or alias) password .
- Process Engine maintenance database user (f_maint or alias) password.

To install DB2 client software (Process Engine)

Install the DB2 Administration Client or the DB2 Run-Time Client. Process Engine needs a 32-bit DB2 Client instance.

NOTE On Windows, an DB2 Client instance is automatically created during installation of the client software.

To create DB2 users and groups (Process Engine, UNIX only)

Create the following operating system user and group accounts for the DB2 Client instance, but only if Process Engine will be installed on UNIX.

NOTE You can specify your own user and group names, but they must adhere to system and DB2 naming rules.

- Instance Owner
- Instance Owner primary group

The *instance owner* home directory is where the DB2 Client instance will be created.

- Fenced User
- Fenced User primary group

The *fenced user* runs user-defined functions (UDFs) and stored procedures outside the address space used by the DB2 database.

Each instance must have its own home file system.

Each instance owner must have a unique home directory. All of the files necessary to run the instance are created in the home directory of the instance owner's user ID/name.

The instance owner and its primary group are associated with every instance. The instance owner is assigned during the process of creating the instance.

The primary group of the instance owner automatically becomes the system administration group for the instance and gains SYSADM authority as a database administrator (DBA) over the instance. Other user IDs or user names that are members of the primary group of the instance owner also gain this level of authority.

NOTE The *root* user cannot act as a DBA. You must log on as the instance owner to act as the DBA.

When the database authentication type is set to SERVER or SERVER_ENCRYPT, two additional operating system users must be created on the database server where the DB2 database resides. You will be provided details in [“To create additional DB2 users and groups \(Process Engine CLIENT authentication only\)” on page 123.](#)

To create DB2 Client instances for Process Engine (UNIX)

Depending upon your configuration, create at least one 32-bit DB2 Client instance for Process Engine. The client instance is referred to as the local instance in the Process Engine installation program.

To create additional DB2 users and groups (Process Engine CLIENT authentication only)

Create the following two additional operating system users with SYSADM authority to access the DB2 database.

- Process Engine runtime database user (f_sw or alias)
 - Primary runtime Process Engine database user
 - Used only by the Process Engine software to access the DB2 database.
- Process Engine maintenance database user (f_maint or alias)
 - Process Engine Database maintenance user
 - Mainly used by the customer for database maintenance

Unlike the instance owner user, these users don't need to have separate file systems for their home directories.

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. For DB2 Client authentication these are the users and passwords that will be entered into the Process Engine installation program to be used by the Process Engine software. Currently, you must ensure that each Process Engine in a farm environment using DB2 Client authentication has the same runtime user name.

To catalog the DB2 server node

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 server name server server side instance tcpip port #
```

For example:

```
db2 catalog tcpip node sampnod hqvmais20 server 60004
```

To create the Process Engine DB2 database alias

Use the db2ca tool, or create aliases for the DB2 database.

For example:

```
db2 catalog database PEDBAIX at node aix20nod [as alias name]
```

Make note of the DB2 client database alias created during this task as it will be needed in the Process Engine installation steps. This information will be entered in response to the prompt for the database alias name in the Process Engine installation program.

To verify the connection to the DB2 database

Verify that TCP/IP communications have been configured successfully on both server and client computers.

To do so on a Windows client, you can use the DB2 Configuration Assistant by entering one of the following commands at a command prompt to connect the DB2 Client to the DB2 database on the remote database server:

db2ca

db2

To install the z/OS license and modify the classpath (z/OS only, Content Engine and Process Engine)

Install the following file on the Content Engine and Process Engine server and add it to the classpath. This needs to be the Data Source configuration within WebSphere and the same place that the JDBC driver jar is referenced.

db2jcc_license_cisuz.jar

To verify the ability to log on to the DB2 database

After successfully connecting to the remote DB2 database, check the connection using either the Command Line Processor (CLP) or db2. For example, launch the Command Line Processor and enter the following command:

```
db2 connect to database_name or instance_name user PE runtime user using password
```

where:

database_name is the Process Engine DB2 for Linux, UNIX and Windows database name

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

PE runtime user is the Process Engine runtime user (f_sw or alias)

password is the Process Engine runtime user password in the Process Engine database

Verify the ability to connect to the database

Use the procedures in this topic to verify the ability to connect to the database. These procedures can be executed after initial configuration of the database and immediately prior to installation of Process Engine software. Execute these steps on the database server or the client according to whether the database is local to or remote from Process Engine.

To verify the Process Engine database connection (Oracle)

Take the following steps to verify that the Oracle database instance used by Process Engine is accessible. How you log on to sqlplus will vary, depending upon how you will choose to execute the SQL scripts. This procedure will verify that you can connect to the Oracle database in the same way the Process Engine installation program will. Make whatever corrections are necessary before proceeding.

To verify the Process Engine database connection (Oracle)

1. Execute the following at a command prompt:

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

2. Enter one of the following commands at the SQLPlus prompt, as follows:

- If the Process Engine pre-installation SQL scripts will be run from the Process Engine installation program by prompting for the sys password, type the following command:

```
sys/password as sysdba
```

- If the SQL scripts will be run from the Process Engine installation program by using operating system authentication, type the following command:

```
/ as sysdba
```

3. At the prompt, enter the following SQL command:

```
SQL> select instance_name, host_name, version from v$instance;
```

The following represents an example of the information returned:

```
INSTANCE_NAME
```

```
-----
```

```
HOST_NAME
```

```
-----
```

```
VERSION
```

```
-----
```

```
p8dbshr  
HQVWBUCS  
10.2.0.2.0
```

where:

p8dbshr is the instance ORACLE_SID.

*hqvwbu*cs is the database server name.

10.2.0.2 is the Oracle server version.

To verify the Process Engine database connection (DB2)

Verify the connection to the DB2 Process Engine database using the DB2 Control Center tool, or by executing the following commands.

To verify the Process Engine database connection (DB2)

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

Windows

At a command prompt, start the DB2 Command Line Processor by typing the following command:

```
db2cmd
```

and, at the subsequently displayed prompt, enter the following command:

```
db2
```

UNIX

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

```
db2
```

2. At the DB2 prompt, enter the following command:

```
connect to database alias or instance_name user f_sw using f_sw password
```

where:

database_name is the Process Engine DB2 for Linux, UNIX and Windows 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.

Following is an example of the database connection command and the information returned:

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

Database Connection Information

Database server = DB2/AIX64 9.1.0

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.

To verify the Process Engine database connection (SQL Server)

Take the following steps to verify that the SQL Server database instance used by Process Engine is accessible. You will need to know both the Process Engine database and filegroup names. Make whatever corrections are necessary before proceeding.

In this example, the database is *VWdb* and the filegroup name is *vwdata_fg*. Both the database name and filegroup name must match what was defined when the database MS SQL server was installed and configured.

To verify the Process Engine database connection (SQL Server)

1. Log on as a member of the local Administrators group or a user with equivalent permissions. The user you log on as must also be a database administrator. If the database is remote, the SQL connection must also be a trusted connection.

2. At a command prompt, enter:

```
osql -E -D DSN
```

where *DSN* is the ODBC data source name created in [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#) for Process Engine use.

This command puts Process Engine into *osql* interactive mode.

3. At the *osql* prompt, enter:

```
1> use VWdb
```

```
2> go
```

where *VWdb* is the Process Engine database name created in [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#).

This command verifies that the Process Engine database has been created. If you get another prompt with no error, you are attached to that database.

4. Verify that the correct Process Engine filegroup was created. At the *osql* prompt, enter:

```
1> select substring(groupname,1,20) from sysfilegroups where groupname = 'defined filegroup'
```

```
2> go
```

where *defined filegroup* is the default filegroup created in [“Configure Microsoft SQL Server Client for Process Engine” on page 94](#).

A listing of the Process Engine filegroups will display, for example:

```
vwdata_fg
```

Application Server Administrator tasks

As the Application Server Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform:

- Review all rows assigned to the Application Server Administrator (ASA) in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

HINT 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 the Role you are interested in.
- 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).
- For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- If you are installing in a non-English environment, review the considerations and procedures in [“Installing P8 Platform in a non-English environment” on page 248](#) before you begin your preparation tasks.
- Review the following planning considerations:
 - [“Application Server Planning Considerations” on page 129](#)
- Create the accounts required by FileNet P8, which are listed for the Application Server Administrator in the following:
 - [“Task to be performed by: Application Server Administrator” on page 67](#)
- Configure your application server for Content Engine according to your application server type:
 - [“Configure WebSphere for Content Engine” on page 131](#)
 - [“Configure WebLogic for Content Engine” on page 134](#)
 - [“Configure JBoss for Content Engine” on page 138](#)
- Configure your application server for Application Engine or Workplace XT according to your application server type:
 - [“Configure WebSphere for Application Engine or Workplace XT” on page 141](#)
 - [“Configure WebLogic for Application Engine or Workplace XT” on page 142](#)
 - [“Configure JBoss for Application Engine or Workplace XT” on page 143](#)
- Configure application server-based Process Engine clients:
 - [“Configure Process Engine clients for ORB” on page 144](#)

Application Server Planning Considerations

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication](#) for reference information about support for EJB and Web Services transports.

All platforms

Note the requirement for J2EE application servers. 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 (IBM WebSphere, BEA WebLogic, or JBoss) and their version numbers are identical for both components. Also, the applications must use Enterprise Java™ Bean (EJB) transport.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication](#) for reference information about support for EJB and Web Services transports.

Note that the Java Virtual Machine determines the maximum number of object stores. 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.

Note the impact of deploying Content Engine and other applications on the same machine. 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, as configurations will not match what has been qualified by IBM FileNet 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.

Note the impact of deploying Content Engine on multiple application server nodes. If you intend to deploy Content Engine to multiple server nodes in your production environment, the type of environment will determine how you install Content Engine, as described in “WebSphere and WebLogic” on page 33 and “JBoss” on page 33.

WebSphere and WebLogic - Content Engine Multi-Server Deployment

Note the impact of deploying centrally managed servers in farms or clusters. 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 (WebSphere) or the Administrator node (WebLogic). To install Content Engine in such an environment, see the IBM FileNet P8 Platform High Availability Technical Notice.

Note the impact of deploying centrally managed servers not in farms or clusters. In an environment in which multiple Content Engine instances may be geographically dispersed, and

where each instance may have its own local Application Engine or Workplace XT server, you will initially install Content Engine on the Deployment Manager node (WebSphere) or the Administrator node (WebLogic). You will then perform post-deployment steps using the administrative console of the application server to deploy Content Engine to other managed servers (see “Deploy Content Engine into a Managed Environment” on page 207).

JBoss Content Engine Multi-Server Deployment

Note the impact of deploying non-managed servers that are farmed or clustered. JBoss does not have a central management server. In an environment of multiple physical servers that are farmed or clustered, you will install Content Engine initially on one physical server (using the "all" instance), and then copy pertinent files and directories to the other servers. To install Content Engine in such an environment, see the IBM FileNet P8 Platform High Availability Technical Notice.

Note the impact of deploying non-managed servers that are not farmed or clustered. In an environment in which standalone application servers are individually managed and not farmed or clustered, and where the servers may be geographically dispersed, you will install Content Engine on each server. Each Content Engine instance will point to the same directory service and GCD database. The data sources you create for each instance will also point to the same object store databases.

Configure WebSphere for Content Engine

NOTE This task assumes you have already installed WebSphere Application Server on the machine where you are going to install and deploy Content Engine.

Content Engine is deployed into a profile in WebSphere Application Server. The profile serves as an application server environment for Content Engine. A default profile (named *AppSrv01* on WebSphere 6.1) is part of the initial WebSphere Application Server installation. You can deploy Content Engine into this profile, or create another profile for this purpose.

Perform the following procedure to specify the WebSphere environment variables for Content Engine and to set permissions on the profile directory for the user who will run Configuration Manager to deploy Content Engine. Among the variables you will specify is the JDBC driver corresponding to the type of database (DB2, Oracle, or SQL Server) where the Global Configuration Data (GCD) database will reside.

To specify the WebSphere environment variables

1. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
2. Install JDBC drivers on the WebSphere machine, as follows:
 - a. Obtain the JDBC drivers, depending on your database type.

DB2

Find the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](#) (<http://www.ibm.com>) by searching for “JDBC Type 4”.

Microsoft SQL Server

Find the Microsoft SQL Server Driver 2005 JDBC Driver, [sqljdbc.jar](#), at Microsoft Support.

Oracle

Access the [Oracle JDBC Driver Downloads](#) web site (http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) and find the JDBC driver file that matches the version of the JDK on the WebSphere machine.

- b. Copy the JDBC driver file from [Step a](#) to one of the following locations:

UNIX

[/opt/jars](#)

CAUTION Do not copy the file to [...WebSphere/AppServer/lib/ext](#).

Windows

[C:\jars](#)

CAUTION Do not copy the file to [...WebSphere\AppServer\lib\ext](#).

3. If you already have a profile for Content Engine, continue at [Step 4](#); or run the command script at one of the following (default) locations to create a new profile.

NOTE Make a note of your profile name, as you will need to specify it when you run Configuration Manager.

AIX

/usr/IBM/WebSphere/AppServer/wasprofile.sh

Other UNIX

/opt/IBM/WebSphere/AppServer/wasprofile.sh

Windows

C:\Program Files\IBM\WebSphere\AppServer\bin\wasprofile.bat

4. Start the WebSphere administrative console and log on to your profile as *ce_appserver_console_admin*, the Administrator Console User. For details on required accounts and related permissions, see [“Accounts for Content Engine” on page 66](#)
5. (WebSphere 6.1) Navigate to **Environment > WebSphere Variables** and perform these substeps to specify the JDBC driver path:
 - a. Choose the Cell scope entry from the *All scopes* drop-down list.
 - b. Click **New** to create a new WebSphere variable whose name is one of the JDBC environment variables shown in the following table, depending on your database type.
 - c. Set the value of the variable to the JDBC driver path you specified in [Step 2](#) and save your change to the master configuration.
 - d. Choose the Node scope entry from the *All scopes* drop-down list.
 - e. In the table of substitution variables, click the item in the Name column that corresponds to the JDBC environment variable in the following table for your database type.

Database	JDBC Environment Variable
MS SQL Server	WebSphere 6.x MSSQLSERVER_JDBC_DRIVER_PATH
	WebSphere 7.x MICROSOFT_JDBC_DRIVER_PATH
Oracle	ORACLE_JDBC_DRIVER_PATH
DB2	DB2UNIVERSAL_JDBC_DRIVER_PATH

- f. Set the value of the item to the JDBC driver path you specified in [Step 2](#), and save your change to the master configuration.
6. Navigate to **Servers > Application servers > server1 > Java and Process Management > Process Definition > Java Virtual Machine** and perform the following substeps:

- a. Set the values for initial and maximum heap sizes (in megabytes), as follows, where *server1* is the name of the server in which you will deploy Content Engine:

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

- b. Save your changes to the master configuration.
7. If any of your object stores will be of a database type that differs from those whose JDBC environment variables you have already specified in this procedure, return to [Step 1](#); otherwise, continue at [Step 8](#).
8. (Optional) Navigate to **Servers > Application Servers > *server1*** and perform the following substeps to set the transaction timeout value:
 - a. Click the **Runtime** tab and then click **Transaction Service**.
 - b. Change the Total transaction lifetime timeout parameter value to at least 600 (seconds).

CAUTION If the timeout value is not large enough, some administrative processes (such as upgrading Content Engine from version 3.5.x or adding an expansion product) may fail.

To set permissions for the Configuration Manager user

1. Set the Configuration Manager user (*config_mgr_user*) permissions on the profile directory (and all its subdirectories) in which Content Engine will be deployed, as follows:

UNIX

Read, write, and execute permissions

Windows

Read & Execute, and Write permission

Configure WebLogic for Content Engine

NOTE This task assumes you have already installed WebLogic Server on the machine where you are going to install and deploy Content Engine.

To configure WebLogic

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

1. Use the WebLogic Configuration Wizard to create a WebLogic domain for Content Engine. In the examples below, use the domain name *FNCEDomain*. Keep the following in mind as you configure the domain:
 - a. Set the server start mode to Production mode.
 - b. Select the appropriate Java Development Kit (JDK) for your environment.
 - (Windows) The version of the JRockit SDK specified in the Third Party Support Information section of the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
 - (AIX) IBM SDK 1.5.0 at /usr/java5
2. Use the WebLogic Administration Console to set the following:
 - a. (Optional) Create a WebLogic authentication provider. You can use the Content Engine Configuration Manager tool to create a WebLogic authentication provider later on, or you can create the provider now using the WebLogic Console. If you want to configure the LDAP later with See “Configure Content Engine instances” in the *FileNet P8 Platform Installation and Upgrade Guide* for information on using the Configuration Manager tool to create the LDAP provider.

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

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

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

NOTE If performance problems are encountered, change the Group Membership Searching parameter value to *limited*.

- b. Do one of the following to allow or prohibit logons to *FNCEDomain* by LDAP-authenticated users in the DefaultAuthenticator who are not in *FNCEDomain*'s active security realm:
 - Set the Control Flag to SUFFICIENT to allow logons by users not in the active security realm (such as users in your authentication domain).
 - Set the Control Flag to REQUIRED to prohibit logons by users not in the active security realm (such as users in your authentication domain).

If you choose REQUIRED, all the users you wish to authenticate for Content Engine must exist not only in the LDAP directory, but also exist as WebLogic users who are in the Default Authenticator provider.
 - c. If you are using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers so that the most-frequently-used provider is first in the list, and the least-frequently-used is last. Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.
 - d. Specify the following heap sizes for the JVM:
 - Initial Java heap size (-Xms): 512 MB
 - Maximum Java heap size (-Xmx): 1024 MB
 - e. (Optional) Adjust transaction-timeout value. Content Engine relies on the transaction-timeout value, whose default may be too short for some standard or administrative processes (such as adding an expansion product or upgrading to the latest version of Content Engine). Set the JTA node timeout in seconds to at least 600 seconds.
3. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
 4. Depending on your database, use one of the following procedures to install the JDBC drivers.

DB2

- a. Obtain the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](http://www.ibm.com) (<http://www.ibm.com>) by searching for “JDBC Type 4”.
- b. Add the db2jcc.jar and db2jcc_license_cu.jar files to the WebLogic classpath. Edit the file startWebLogic.cmd or startWebLogic.sh for the WebLogic domain you created. For example,


```
set CLASSPATH=%CLASSPATH%;c:\db2\jdbc\db2jcc.jar;c:\db2\jdbc\
db2jcc_license_cu.jar
```
- c. Stop and then start WebLogic Server.

Microsoft SQL Server

- a. Download and unzip Microsoft SQL Server Driver 2005 JDBC Driver, sqljdbc.jar, from Microsoft Support to a directory *jdbc_path* on your application server machine, such as:

UNIX
/opt/jars

Windows

C:\jars

- b. Perform one of the following steps, depending on your operating system type.

Windows

Edit the file startWebLogic.cmd (by default, in the directory C:\bea\user_projects\domains\bin\FNCEDomain) for the WebLogic domain you created. Insert the following two lines immediately after the first occurrence of the line `CLASSPATH=...`

```
set JDBC_PATH=jdbc_path\sqljdbc_1.0\enu\sqljdbc.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%
```

AIX

Add the following line to the file setDomainEnv.sh file:

```
JAVA_OPTIONS="$JAVA_OPTIONS -
Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"
```

UNIX

Edit the file startWebLogic.sh by inserting the following two lines immediately after the first occurrence of the line `CLASSPATH=...`

```
JDBC_PATH=jdbc_path/sqljdbc_1.0/enu/sqljdbc.jar
CLASSPATH=$JDBC_PATH:$CLASSPATH
```

- c. Stop and then start WebLogic Server.

Oracle

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

NOTE If you intend to install AddOns (extensions to IBM FileNet P8 core components), and your Content Engine database will be Oracle, your Oracle JDBC Driver file requirements may be more restrictive. For the required version and patch number, see the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

- c. From the Oracle web site, apply the patch Oracle Patch `Ojdbc##.jar`.
- d. Edit the file startWebLogic.cmd or startWebLogic.sh for the WebLogic domain you created. Add the following line immediately after the first line that starts with `set CLASSPATH`.

Windows

```
set JDBC_PATH=<jdbc_path>\ojdbc##.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%;
```

UNIX

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


- e. Stop and then start WebLogic Server.
5. Give the Configuration Manager user the following permissions:
- Read, write, and execute permission on the domain directory `../users_projects/domains/your_domain`.
 - Read and execute permission on the `../common/bin` directory.

Configure JBoss for Content Engine

NOTE This task assumes you have already installed JBoss Application Server on the machine where you are going to deploy Content Engine.

To configure JBoss for Content Engine

1. Navigate to the JBoss directory JBOSS_DIST/server, which contains configuration file sets.
NOTE If you are installing into a JBoss cluster, use the JBoss directory JBOSS_DIST/all.
2. Create a new configuration file set by copying the default configuration file set to a new directory (called server1 in this procedure) within the /server or /all directory.
3. Edit the run.conf configuration file, located at JBOSS_DIST/bin, as follows:
 - a. Add a line to specify the path to the JDK on the machine where JBoss is installed, as shown in the following example:

```
JAVA_HOME="<path_to_Java_JDK>"
```
 - b. In the JAVA_OPTS line, change the -Xms and -Xmx values from

```
-Xms128m -Xmx512m
```

to

```
-Xms512m -Xmx1024m
```
 - c. Save your edits.
4. Open the file login-config.xml for editing. This file is typically located at ../server/*myserver*/conf, where *myserver* is the name of the JBoss server instance.
 - a. In the <!DOCTYPE declaration, change

```
"http://www.jboss.org/j2ee/dtd/security_config.dtd"
```

to

```
"<jboss_install_dir>/docs/dtd/security_config.dtd"
```

where *<jboss_install_dir>* is the directory where JBoss is installed.
 - b. Save your edit.
5. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19](#).
6. Install JDBC drivers on the JBoss machine, as follows:
 - a. Obtain the JDBC drivers, depending on your database type.

DB2

Find the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](http://www.ibm.com) (<http://www.ibm.com>) by searching for “JDBC Type 4”.

Microsoft SQL Server

Find the Microsoft SQL Server Driver 2005 JDBC Driver, sqljdbc.jar, at Microsoft Support.

Oracle

Access the [Oracle JDBC Driver Downloads](http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) web site (http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) and find the JDBC driver file that matches the version of the JDK on the JBoss machine.

- b. Place the file JDBC driver file from [Step a](#) in CLASSPATH by copying it to the directory JBOSS_DIST/server/server_name/lib.

NOTE If you are installing into a JBoss cluster, use the JBoss directory JBOSS_DIST/all/server_name/lib.

7. Increase the database transaction timeout. Edit the jboss-service.xml file, and set TransactionTimeout to at least 600, as in this example:

```
<mbean code="org.jboss.tm.TransactionManagerService"
  name="jboss:service=TransactionManager"
  xmbean-dd="resource:xmdesc/TransactionManagerService-xmbean.xml">
  <attribute name="TransactionTimeout">600</attribute>
```

8. If you are deploying multiple instances of Content Engine of the same server, do the following for each additional instance:
 - a. Copy the configuration file set that you just created and modified in [Step 2](#) through [Step 6](#) from the /server/server1 directory to a new directory. Use a separate directory for each instance.
 - b. Assign unique port numbers to each instance. Refer to your JBoss documentation for details.
9. If it isn't already running, start JBoss as follows, and leave the command window open:

UNIX

```
./run.sh -c server1
```

Windows

```
run.bat -c server1
```

Assign directory permissions

Give the Configuration Manager user read and write permission on the server directory where the Content Engine instance will be installed.

Configure JBoss server clusters

JBoss servers can be grouped together into a cluster for performance or to provide high availability. This guide provides on minimal instructions for setting up a JBoss cluster. Refer to the *IBM FileNet P8 Platform High Availability Technical Notice* for details on how to set up your IBM FileNet P8 system using clusters, farms, and other high availability software and hardware. To

download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Because JBoss clusters do not have an administrative server, you will choose a single JBoss server on which to install and configure the first instance of Content Engine, then copy the necessary files to the rest of the nodes in the cluster. See [“Deploy Content Engine to additional JBoss servers in a cluster” on page 50](#) for details of which files to copy to other nodes.

Configure WebSphere 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 for Application Engine or Workplace XT

1. Verify that the application server is set to use JSESSIONID as default cookie name.

To avoid forcing end users to log in individually 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 after you install Application Engine or Workplace XT. For IBM specific recommendations, see the IBM FileNet P8 Platform Performance Tuning Guide. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.](#)

Configure WebLogic for Application Engine or Workplace XT

You must install WebLogic 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 Application Engine or Workplace XT and each instance of the Content Engine must run in its own JVM. For assistance in sizing your system, contact your service representative.

To configure WebLogic for Application Engine or Workplace XT

1. Verify that the application server is set to use JSESSIONID as default cookie name.

To avoid forcing end users to log in individually applets such as Process Designer, Search Designer, and Process Simulator, configure the application server to use JSESSIONID as cookie name, and not use application-unique cookie names. Using JSESSIONID is typically the default setting for the supported application servers. Both Application Engine and Workplace XT use cookie names for passing session information between Application Engine or Workplace XT and the client browser.

2. Create a WebLogic domain before installing and deploying Application Engine or Workplace XT. Refer to your BEA documentation for detailed instructions.

NOTE You will perform further configuration on WebLogic after you perform the installation.

Configure JBoss for Application Engine or Workplace XT

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

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

To configure JBoss for Application Engine or Workplace XT

1. Verify that the application server is set to use JSESSIONID as default cookie name.

To avoid forcing end users to log in individually applets such as Process Designer, Search Designer, and Process Simulator, configure the application server to use JSESSIONID as cookie name, and not use application-unique cookie names. Using JSESSIONID is typically the default setting for the supported application servers. Both Application Engine and Workplace XT use cookie names for passing session information between Application Engine or Workplace XT and the client browser.

NOTE You will perform further configuration on JBoss after you perform the installation.

Configure Process Engine clients for ORB

To configure 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:

- WebSphere with the IBM JVM:
No changes are required.
- JBoss Installed by unzipping binaries with the IBM or the Sun JVM:
No changes are required.
- JBoss Application Server Installed via JEMS JBoss Enterprise Middleware System Installer with IBM or Sun JVM:
No changes are required if you installed the JBoss application server by installing the JBoss Enterprise Middleware Software with the “Default” option.
- WebLogic with the Sun JVM:
No changes are required.
- WebLogic with the JRockit JVM:
Replace the default WebLogic ORB with the Sun ORB by adding the following code to your application server startup file:

```
set JAVA_OPTIONS=%JAVA_OPTIONS% -  
Dorg.omg.CORBA.ORBClass=com.sun.corba.se.impl.orb.ORBImpl  
  
set JAVA_OPTIONS=%JAVA_OPTIONS% -  
Dorg.omg.CORBA.ORBSingletonClass=com.sun.corba.se.impl.orb.ORBSingleton
```


Plan and prepare for IBM FileNet P8 upgrade

This section contains the following major topics:

- [“Upgrade planning considerations” on page 151](#)
- [“IT Administrator upgrade tasks” on page 157](#)
- [“Security Administrator upgrade tasks” on page 180](#)
- [“Database Administrator upgrade tasks” on page 196](#)
- [“Application Server Administrator upgrade tasks” on page 202](#)
- [“FileNet P8 Administrator upgrade tasks” on page 214](#)

Plan the upgrade

This section describes considerations and information you should use when planning your IBM FileNet P8 system upgrade. Review this section thoroughly before you start to upgrade IBM FileNet P8 components or required third-party software.

Before you begin to upgrade the IBM FileNet P8 Platform, review the following information:

- [“Upgrade overview” on page 147](#)
- [“Upgrade planning considerations” on page 151](#)
- [“Definition of upgrade roles” on page 154](#)

Upgrade overview

This topic provides a brief overview of the methods you can use to upgrade to components in the 4.5 versions of the Content Manager (CM) and Business Process Manager (BPM) products.

This 4.5 versions of CM and BPM, which contains the core IBM FileNet P8 Platform components (Content Engine 4.5, Content Search Engine 4.5, Process Engine 4.5, and Application Engine 4.0.2) support upgrades from either the 3.5 or 4.0 version series.

This support also covers the associated expansion components, such as those currently packaged as part of the CM and BPM products, for example, Workplace XT, Rendition Engine, Content Federation Services, Process Analyzer, and Process Simulator. Other associated expansion products, such as Records Manager, also support these dual upgrade starting points. See [“General requirements for all IBM FileNet P8 systems” on page 151](#) for a list of product-component versions from which you can start an upgrade.

Standard upgrade from version 3.5 or 4.0

A *standard upgrade* sequence assumes that you will complete the upgrade of all installed components in one session, for example, over the course of a weekend. The flow of procedural task topics in the upgrade section of this guide reflects a sequence that supports a standard upgrade from either the 3.5 or 4.0 product versions.

Staged upgrade from version 4.0 only

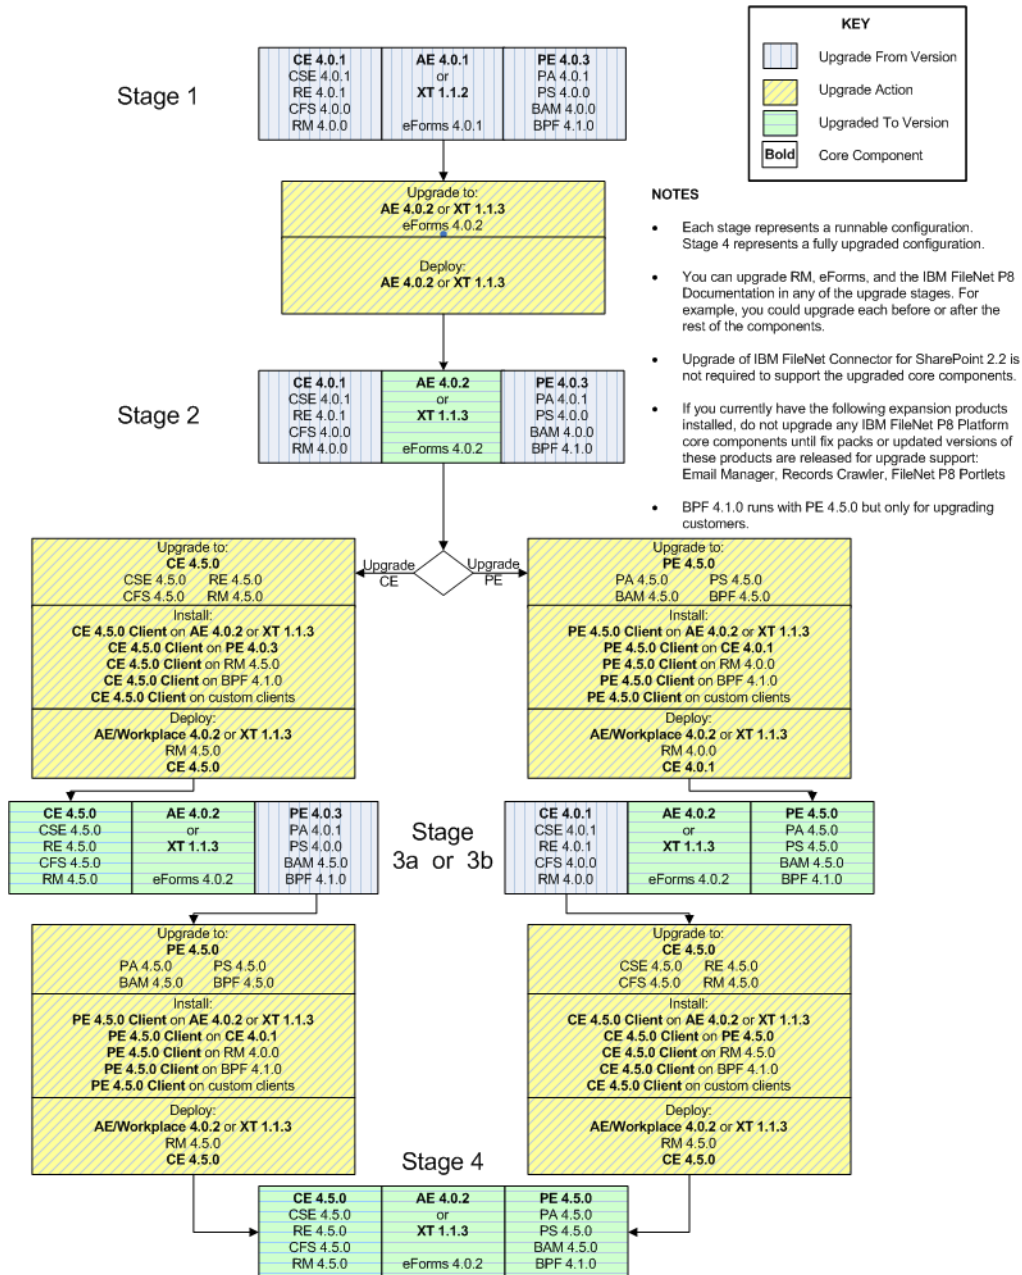
Customers already running IBM FileNet P8 Platform version 4.0 components (including any supporting expansion products) can choose to stage their upgrades over an extended period of time. A *staged upgrade* is one that you complete in a series of independent tasks, or sets of tasks, that you complete in phases. Each phase, or stage, is a potential stopping point where you have a runnable configuration. This means that you can upgrade one component and still have a working configuration without upgrading all components during the same upgrade session. Upgrade stages do not necessarily follow a single sequential flow. You can start at any number of points and can upgrade core components in an order that fits your scheduling or resource requirements.

The following flow charts show the various start points for such staged upgrade scenarios. Study these flow charts carefully to understand the relationship of particular components to one another. Unless otherwise noted, components shown in a given graphical box must be upgraded together. For example, when you upgrade Process Engine, you must also upgrade Process Analyzer, Process Simulator, and Business Activity Monitor.

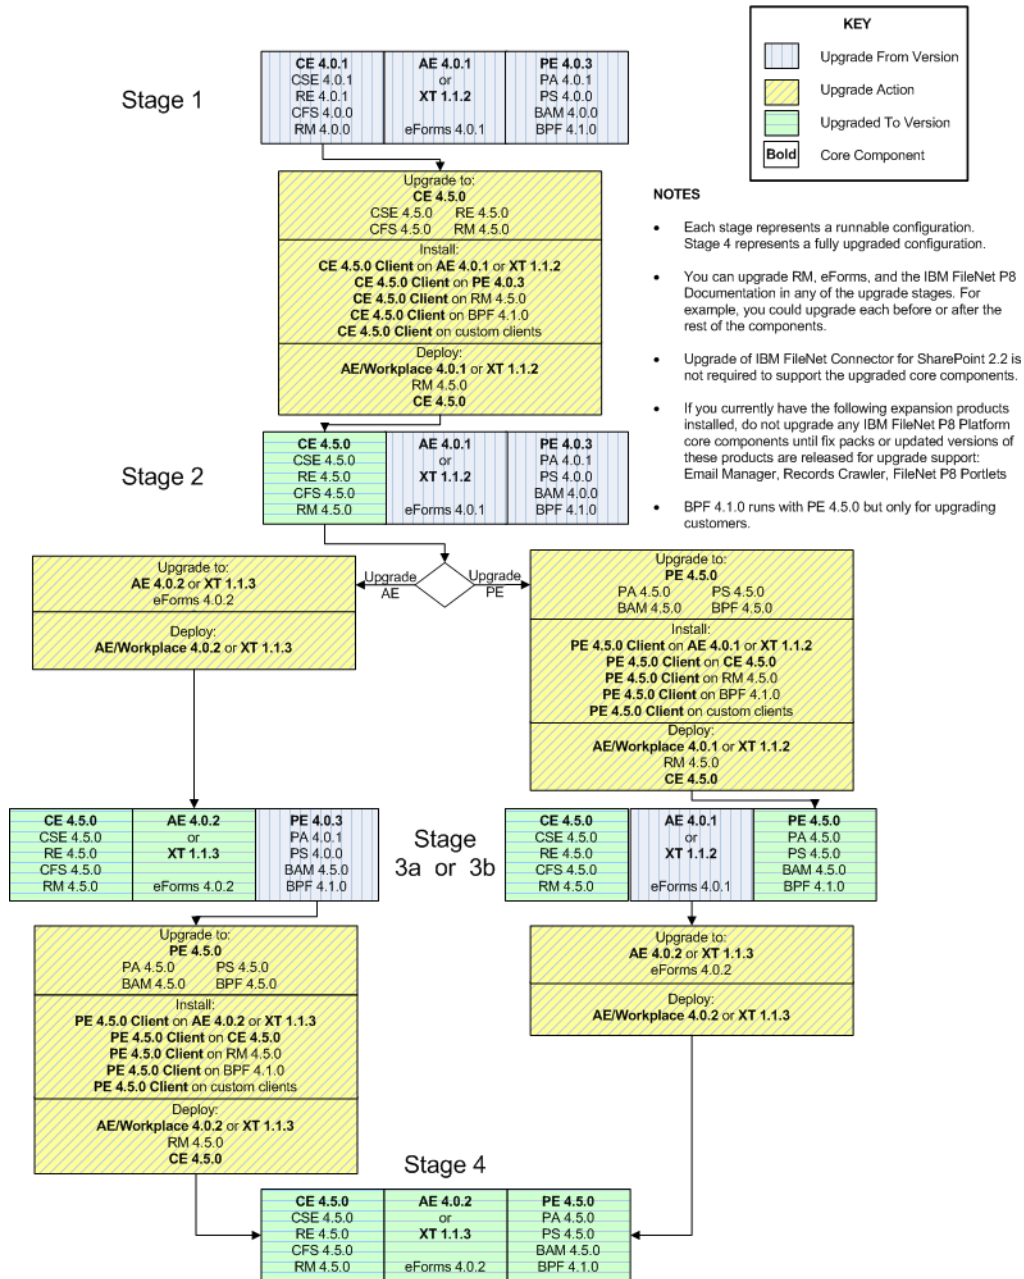
CAUTION Pay special attention to the bulleted notes on the flow chart because special options and restrictions do exist for some expansion products.

NOTE Staged upgrades can impact how and when you apply Content Engine and Process Engine Client files to various component servers and client machines that require them. This guide points out when such installation activities are warranted and directs you to the appropriate procedures, which can be positioned elsewhere in the standard-upgrade sequence.

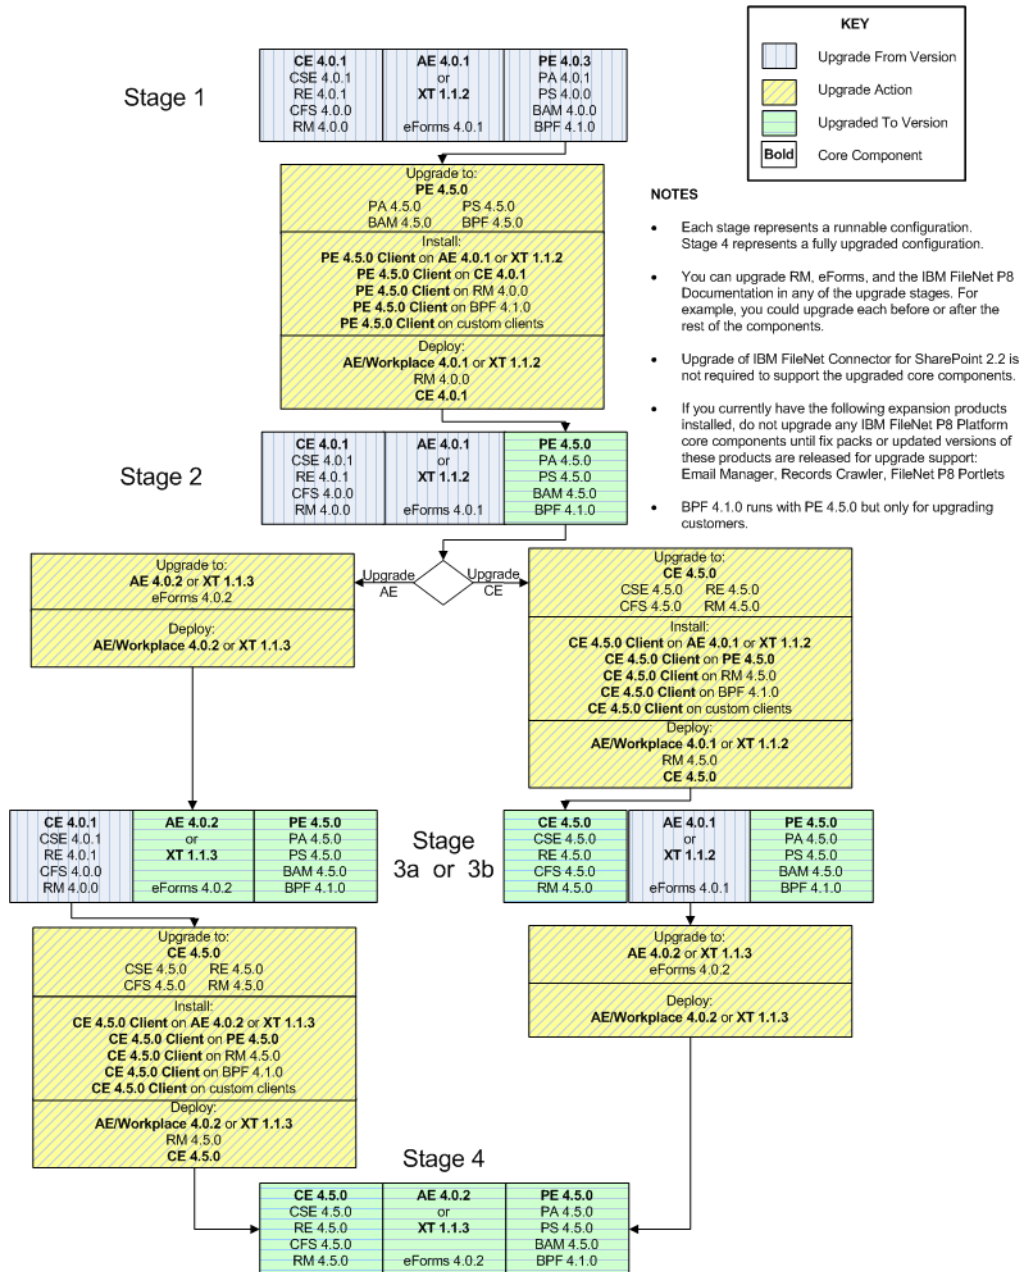
Staged upgrade scenario 1: upgrade Application Engine or Workplace XT first



Staged upgrade scenario 2: upgrade Content Engine first



Staged upgrade scenario 3: upgrade Process Engine first



Upgrade planning considerations

This section lists details that will help you prepare your environment for the upgrade of a IBM FileNet P8 system. In many cases, the items you see listed are links to more detailed information, which will help you plan a system upgrade. Review this section thoroughly before you start to upgrade IBM FileNet P8 components or required vendor software.

General requirements for all IBM FileNet P8 systems

- **It is best practice to plan and test the upgrade on a designated test system first.** Verify the upgrade is successful by running functionality and stress tests. After successful verification, perform the production upgrade.
- **Apply the required minimum level of IBM FileNet P8 software updates for the release you are currently running before you upgrade.** Before you begin your upgrade to IBM FileNet P8 Platform 4.5.x, use the information in the following table to verify that each component is using the minimum fix pack level or higher. This information applies to standard upgrades from version 3.5 or 4.0. For information related to staged upgrades, review the IBM FileNet Compatibility Matrix on the [IBM Information Management support page on www.ibm.com](http://www.ibm.com).

Component	3.5.x Fix Pack	3.5.x Software Build	4.0.x Fix Pack	4.0.x Software Build	4.5 Release	4.5 Software Build
Content Engine	CE-3.5.2-002	kl195.025	CE 4.0.1-002	dap435.032	CE 4.5	dap440
Process Engine	PE-3.5.2-001 PE 3.5.3-002	pe185.006 and pui185.009 pe197.xxx and pui197.xxx	PE 4.0.2-001 PE 4.0.3	pe425.007 and pui425.016 pe430.060a and pui430.052	PE 4.5	pe450 pui450
Application Engine	AE-3.5.1-003	per185.027.a01	AE-4.0.1-002	per410.041d	AE 4.0.2	per450
IBM FileNet P8 eForms	P8eF-3.5.1	raptor220.004	eForms 4.0.1	raptor305A.006	eForms 4.0.2	raptor330
Records Manager	3.7.0	rec190.007	4.0.0	rec200.056	RM 4.5	rec210
Workplace XT	1.1.2	orion210	1.1.2	orion210	Workplace XT 1.1.3	orion300

Component	3.5.x Fix Pack	3.5.x Software Build	4.0.x Fix Pack	4.0.x Software Build	4.5 Release	4.5 Software Build
CFS	No direct upgrade from 3.5 to CFS 4.5	hel195	4.0.0-001	hel435.010	CFS 4.5	hel440
Process Analyzer	PA-3.5.2	pa190.015b	PA-4.0.1.001	pa435.014	PA 4.5	ps450
Process Simulator	PS-3.5.2-005	ps185.023	PS-4.0.0.006	ps420.027c	PS 4.5	ps450
BAM	BAM-3.6-001.001	bam195.002	BAM-3.6-001.001	bam195.002	BAM 4.5	bam450
Rendition Engine	RE-3.5.1	re400	RE-4.0.1	re400	RE 4.5	re440
Content Search Engine			CSE 4.0.1	cse435	CSE 4.5	cse440
Business Process Framework	BPF 3.6		BPF 4.1	bpf405.008		

Run the version tool, as documented in the *IBM FileNet P8 Platform 3.5.x or 4.0.x Version Tools Technical Notice* to determine the software build. Use the table above to map the software build number to the minimum software update.

If any FileNet P8 component is below the minimum software level documented here, update that component accordingly.

Applying a fix pack for one component might require fix packs for other components. Prior to applying any fix pack, review the IBM FileNet Compatibility Matrix on the [IBM Information Management support page on www.ibm.com](#) to assure compatibility between all installed components.

IBM FileNet P8 expansion product considerations

Content Federation Services

Content Federation Services (CFS) 4.0.x is not compatible with Content Engine 4.5.x. You should shut down CFS Exporter before beginning the upgrade to Content Engine 4.5.x, and leave it shut

down until you finish the upgrade to CFS 4.5. Therefore you should plan to upgrade CFS as soon as possible after upgrading Content Engine.

For more information, see the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* topic “Upgrade overview.”

Records Manager

- **(3.7 to 4.5 only) You must update the Records Manager (RM) data model on each file plan object store (FPOS) before you attempt to upgrade associated Content Engine data.** The Upgrader Tool will fail if run on an FPOS unless it has been updated to the 4.5.x data model. This limitation only applies to FPOSs; you can update ROS object stores without first updating the data model. For more information, see [“Upgrade Content Engine Data” on page 538](#) and the *IBM FileNet Records Manager Installation and Upgrade Guide* task “Prepare to Upgrade Object Stores.”

NOTES

- To complete and confirm your IBM FileNet P8 Platform upgrade you must update at least one object store. If your configuration contains only Records Manager object stores, you must either update at least one of them or create a non-Records Manager object store prior to upgrading and use that object store to confirm the upgrade.
- If your site preferences reside on a Records Manager object store, you must update that object store to complete the IBM FileNet P8 Platform upgrade. Alternatively, you can move the site preferences to another object store before the upgrade.
- **Do not install Enterprise Manager 4.5.x on any machine running the 3.5.x version until the Records Manager 4.5.x upgrade is complete.** This includes updating all Records Manager object stores to version 4.5.x. Installing Enterprise Manager 4.5.x will cause unexpected behavior with the Records Manager Data Upgrade Tool.

Vendor software considerations

The 3.5.x, 4.0.x, and 4.5.x versions of IBM FileNet P8 components support common levels of independent software vendor (ISV) products. Exceptions are listed in the *IBM FileNet P8 Hardware and Software Requirements*. To download this document from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Definition of upgrade roles

The tasks in this *Plan and Prepare Your Environment for IBM FileNet P8* guide as well as the rows in the [“Installation and upgrade worksheet” on page 216](#) are organized by administrative roles, listed below.

Your organization may 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 installers during initial setup.
- Runs the Configuration Manager tool during initial setup, followed by launching Enterprise Manager.
- Runs IBM FileNet 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. See [“Installation and upgrade worksheet” on page 216](#).

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

Information technology administrator

- Responsible for the networking and operating systems issues 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.

For tasks assigned to the ITA, see [“IT Administrator upgrade tasks” on page 157](#).

Security administrator

- Responsible for configuring the directory servers required by IBM FileNet P8, including Content Engine, Application Engine, CFS Federation Administration application.
- Creates and maintains directory server user and group accounts.
- Decides on configuration parameters required to connect.
- 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.

For tasks assigned to the SA, see [“Security Administrator upgrade tasks” on page 180](#).

Database administrator

- Creates, configures, maintains database installations and database/tablespaces.
- Responsible for creating database accounts needed by FileNet P8.
- Might have responsibilities regarding the JDBC datasources.
- 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.

For tasks assigned to the DBA, see [“Database Administrator upgrade tasks” on page 196](#).

Application server administrator

- Responsible for providing the application servers required by 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.

For tasks assigned to the ASA, see [“Application Server Administrator upgrade tasks” on page 202](#).

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 using the *gcd_admin* account or an *object_store_admin* account is considered a 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.

For tasks assigned to the P8A, see [“FileNet P8 Administrator upgrade tasks” on page 214](#).

Email Administrator

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

NOTE The Email Administrator role is not used in an upgrade, but is included in this topic for completeness.

Perform the required upgrade preparation tasks

The tasks in this section are divided by administrator role. For information about assigning and defining these roles, see [“Definition of upgrade roles” on page 154](#).

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 [“Use the Installation and Upgrade Worksheet” on page 26](#) for details.

Some of the upgrade procedures in the following subsections are to be performed only if your 3.5.x is your current version of IBM FileNet P8; others only if 4.0.x is your current version. If not otherwise indicated, perform each procedure irrespective of the current version.

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

- [“IT Administrator upgrade tasks” on page 157](#)
- [“Security Administrator upgrade tasks” on page 180](#)
- [“Database Administrator upgrade tasks” on page 196](#)
- [“Application Server Administrator upgrade tasks” on page 202](#)
- [“FileNet P8 Administrator upgrade tasks” on page 214](#)

IT Administrator upgrade tasks

As the Information Technology (IT) Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform upgrade:

- Review all rows assigned to the IT Administrator for Upgrade in the [“Installation and upgrade worksheet” on page 216](#). Provide values for any rows appropriate to your installation that you have not yet completed. For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- Configure the operating systems to prepare for component upgrade:
 - [“Configure UNIX” on page 158](#)
 - [“Configure Microsoft Windows” on page 172](#)
 - [“Configure network” on page 176](#)
- [“Assign directory permissions for Content Engine upgrade for 4.0.x to 4.5 on UNIX” on page 177](#)

Configure UNIX

The topics in this task describe how to configure UNIX on the servers that will comprise your FileNet P8 system.

Configure UNIX for FileNet P8 servers (all components)

To configure UNIX FileNet P8 servers

1. **Ensure hosts file contents.** On each UNIX-based IBM 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 [“To configure the `/etc/hosts` file” on page 162](#) for additional information.
2. Consult with the application server, database, and P8 administrators to determine port requirements for all the servers in your environment. For details, see [“IBM FileNet P8 ports” on page 265](#).

Configure Content Engine servers (all UNIX)

To configure UNIX FileNet CE servers

- **Content Engine running on a UNIX-based application server.** 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/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 will need to specify when creating file storage areas on UNIX file servers.

NOTE This `umask` setting is required for the user running Content Engine Setup but need not be in the `.profile` file of the user.

- **If you intend to move to UNIX for Content Engine when you upgrade from 3.5.x to 4.5.x, consider the following.**

NOTE If you are upgrading from 4.0.x and want to move from Windows to UNIX, you must contact your IBM FileNet representative and arrange a Professional Services engagement.

- You must install Content Engine server 4.5.x on UNIX as an initial step in the upgrade process. To migrate after you have completed the upgrade and begun using your system, you will have to contact your IBM FileNet representative and arrange a Professional Services engagement.
- You must temporarily leave your existing 3.5.2 file stores (and associated index stores) on Windows to upgrade them using the CE 3.5.2 to 4.5.x Upgrader Tool. This means that if you install Content Engine server on UNIX, you must have an NFS gateway in place (for

example, Windows R2 Gateway or Samba) to enable communication between the new UNIX Content Engine server and Windows file storage areas and index areas.

Once they are upgraded, you can physically move the shared directories for the file storage areas and index areas to UNIX, but you must be sure to establish comparable security settings. For details on these settings, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Storage Area Security](#).

Configure Autonomy K2 Servers (UNIX)

To configure HP-UX for CSE

To install Content Search Engine on HP-UX, manually configure the kernel with following parameters before you begin the Autonomy K2 Master Administration Server installation:

Value	Setting
maxdsiz	1.9 Gbytes (0x7B033000)
maxfiles	2048 Kbytes
maxfiles_lim	2048 Kbytes
maxssiz	160 Mbytes (0xA000000)
max_thread_proc	1024
maxswapchunks	8192
maxtsiz	1 Gbyte (0x40000000)
maxuprc	512
maxusers	128
nkthread	1024
nproc	517

To configure RedHat Enterprise Linux 5.1 for CSE

For RedHat Enterprise Linux 5.1 x86_64, install the following libraries from the original system media:

- *DVD-ROM_mount_point/Server/compat-libstdc++-33-3.2.3-61.i386.rpm*
- *DVD-ROM_mount_point/Server/compat-libstdc++-33-3.2.3-61.x86_64.rpm*

To set the Java_home variable for all UNIX types

Set the Java_Home variable for the supported Java SE Development Kit(JDK).

```
JAVA_HOME=/usr/jdk install path
export JAVA_HOME
```

NOTE Enter the Java_home variable in the .profile file to set this variable each time the user logs in.

To append the appropriate environment variables

Set the following variables, depending on your operating system:

AIX (default install path)

- `PATH=$PATH:/opt/verity/k2/_rs6k43/bin`
`export PATH`
- `LIBPATH=$LIBPATH:/opt/verity/k2/_rs6k43/bin`
`export LIBPATH`

HP-UX (default install path)

- `PATH=$PATH:/opt/verity/k2/_hpux/bin`
`export PATH`
- `SHLIB_PATH=$SHLIB_PATH:/opt/verity/k2/_hpux/bin`
`export SHLIB_PATH`

Linux (default install path)

- `PATH=$PATH:/opt/verity/k2/_ilnx21/bin`
`export PATH`
- `LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/verity/k2/_ilnx21/bin`
`export LD_LIBRARY_PATH`

Solaris (default install path)

- `PATH=$PATH:/opt/verity/k2/_ssol26/bin`
`export PATH`
- `LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/verity/k2/_ssol26/bin`
`export LD_LIBRARY_PATH`

Configure Process Engine Servers (all UNIX)

To configure UNIX FileNet PE servers

1. Perform the following prerequisite tasks in any order:
 - Ensure minimum /tmp size. The /tmp directory must have 510 MB free.

- Save the following files for the root user

```
.cshrc
.Xdefaults
.Xresources
.dbxinit
.dtprofile
.env
.login
.mwmrc
.xinitrc
.profile
```

- Process Engine requires the presence of several partitions. 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.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fnsw	/fnsw file system	2GB	fnsw	fnusr	775
local	/fnsw/local file system	1GB	fnsw	fnusr	775
fn_sec_db0 (raw)	n/a logical volume	64MB	fnsw	fnusr	664
fn_sec_rl0 (raw)	n/a logical volume	64MB	fnsw	fnusr	664

NOTES

- (AIX 6.1 only) Permissions must be set correctly on both the /fnsw and /fnsw/local mount points and the file systems before mounting the file systems.
- Solaris volume management software might use port 32776. This is the default for the Process Engine Communication Port (IOR port).
- These minimum sizes have changed from versions 3.5.x and 4.0.x. For the fn_sec volumes, expand the sizes to the minimum 64MB and the Process Engine upgrade wizard will automatically handle associated configuration changes.

- Allocate a minimum of 500 MB of additional disk space to the /fnsd disk volume. This space is required for upgrades and is in addition to the minimum space requirements called for in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

To configure the /etc/hosts file

Information must be entered into either the server’s DNS table or the hosts file related to Process Engine IP address, server name and NCH domain name. For non-farmed configurations this information can be in either the DNS table or the hosts file on the server. For farmed configurations this must be entered into the host file. In a farmed environment, entries must exist for every Process Engine server in the farm.

Entries must be 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 nch_domain-organization-nch-server load_balancer_name

where:

IP addr is the IP address of the Process Engine server.

hostname is the corresponding host name.

nch_domain-organization is the NCH domain and organization name, as provided to the Process Engine installation program.

load_balancer_name is the name of the load balancer in a farmed configuration.

When entering the domain-organization name, follow these rules:

- Eliminate all characters except ASCII alphanumeric characters and underscores.
- Convert all upper case characters to lower case.
- Insert a hyphen between the domain and organization names.
- Append “-nch-server” as a literal.

For example, a Process Engine has a domain “ace-1” and organization “FileNet”. Its IP address is 123.45.6.78. For this system, the hosts file entry is:

123.45.6.78 ace-1 ace1-filenet-nch-server

NOTE The hyphen in the *nch_domain-organization* name has been removed and the F and N in the “FileNet” organization name have been converted to lower case.

To configure the /etc/hosts file

Information must be entered into either the server’s DNS table or the hosts file related to Process Engine IP address, server name and NCH domain name. For non-farmed configurations this information can be in either the DNS table or the host file on the server. For farmed configurations this must be entered into the hosts file. In a farmed environment, entries must exist for every Process Engine server in the farm.

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

IP addr hostname nch_domain-organization-nch-server load_balancer_name

where:

IP addr is the IP address of the Process Engine server.

hostname is the corresponding host name.

nch_domain-organization is the NCH domain and organization name.

load_balancer_name is the name of the load balancer in the farmed configuration.

When entering the domain-organization name, follow these rules:

- Eliminate all characters except ASCII alphanumeric characters and hyphens.
- Convert all upper case characters to lower case.
- Insert a hyphen between the domain and organization names.
- Append “nch-server” as a literal.

For example, a Process Engine has a domain “ace_1” and organization “FileNet”. Its IP address is 123.45.6.78. For this system, the hosts file entry is:

123.45.6.78 ace-1 ace1-filenet-nch-server

NOTE The underscore character has been removed from the “ace_1” domain name and the F and N in the “FileNet” organization name have been converted to lower case.

To expand partitions for the PE's SEC databases on UNIX

- **(3.5.x to 4.5.x upgrades only) If necessary, expand the following raw partitions that are used for the SEC databases:**
 - fn_sec_rl0: expand raw partition to 64MB
 - fn_sec_db0: expand raw partition to 64MB

Configure Process Engine servers (AIX)

To configure AIX-based FileNet PE servers

1. Perform the following prerequisite tasks in any order:
 - The Kernel must be set to 64-bit mode.
 - The swap space must be set to 1.5 - 2 times RAM.
 - The Maximum Number of Processes allowed per user must be set to at least 400.
 - The Maximum Kbytes of real memory allowed for MBUFS must be set 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.

- The Maximum Number of FIXED licenses (Num) must be set to a minimum of 16.
- The following filesets must be installed and committed:
 - 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).

To modify /etc/rc.dt and /etc/tunables/nextboot for AIX 5.3 and 6.1.

1. As the root user, execute the following commands or set them by editing the /etc/tunables/nextboot.

```
/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
```

2. Add the following statements at the beginning for /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
```

3. Restart the server (shutdown -Fr) for these settings to take effect. Executing the commands at the command line is not sufficient. The changes must be generated via the "nextboot" to avoid bind failures..
4. Check the values by executing:

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

To correct a required link in AIX 6.1

AIX 6.1 installs the file /usr/lib/libMrm.a in a directory that is different from the one required by the IS 4.1.0 mini-installer. As a result, the Process Engine installation will fail when running lic_admin.

To prevent this failure, use the following workaround after you install AIX 6.1, but before you install IS4.1.0:

1. Log in as a user with root privileges.
2. Remove any "filenet-*" entries in the /etc/services file.
3. Download and install either SP3 or APAR: IZ13179 on AIX 6.1.
4. Enter the following command:

```
ln -s /usr/lpp/x11/lib/R1/libMrm.a /usr/lib/libMrm.a
```

To install a required Oracle patch for AIX 6.1

AIX 6.1 requires the Oracle 10gR2 6613550 patch to fix a problem with rootpre.sh.

1. Download patch number 6613550 from the Oracle support web site.
2. As a user with root privileges, run the script.
3. As Oracle user, launch the Oracle Universal Installer (runInstaller).

Configure Process Engine servers (HP-UX)

The following operating system prerequisites apply to HP-UX-based FileNet PE servers.

To configure HP-UX servers for Process Engine

1. Perform the following prerequisite tasks in any order:
 - On each HP-UX 11 or HP-UX 11i machine where you will install a JVM-based IBM FileNet P8 component (such as Content Engine), or where an associated third-party JVM-based component (such as WebLogic or WebSphere) will run, increase the values of the kernel parameters max_thread_proc (maximum number of threads per process) and nkthread (maximum number of kernel threads in the system) beyond their default values, which are too small for IBM FileNet P8 applications.

Refer to the HP web page "Programmer's guide for Java 2 HP-UX configuration for Java support" for tools to determine values of these two kernel parameters that are sufficient for IBM FileNet P8.

- The physical memory must be at least 512 MB.
- The Kernel must be set to 64-bit mode.
- The swap space must be set as follows:
 - Two times RAM if RAM < 1GB
 - 1.5 times RAM if RAM between 1GB and 2GB

- Equal to RAM if between 2GB and 8GB
- .75 times RAM if > 8GB
- The /etc/nsswitch.conf should have the following entry:

```
hosts:  files [NOTFOUND=continue] dns
```

To perform symbolic links for X11 libraries

1. Log on as the root user.
2. At the prompt, execute the following:

```
cd /usr/lib
ln -s /usr/lib/libX11.3 libX11.sl
ln -s /usr/lib/libXIE.2 libXIE.sl
ln -s /usr/lib/libXext.3 libXext.sl
ln -s /usr/lib/libXhp11.3 libXhp11.sl
ln -s /usr/lib/libXi.3 libXi.sl
ln -s /usr/lib/libXm.4 libXm.sl
ln -s /usr/lib/libXp.2 libXp.sl
ln -s /usr/lib/libXt.3 libXt.sl
ln -s /usr/lib/libXtst.2 libXtst.sl
```

To configure kernel parameters

Ensure that the following parameters are set to at least the values shown unless otherwise noted. The values are appropriate for both HP PA-RISC and Integrity operating systems unless noted otherwise. These values are sufficient to install and initialize the software but system tuning will be required, specifically for the nfiles and maxfiles parameters. |

Kernel Parameter	Minimum Setting PA-RISC	Minimum Setting Integrity V2 and V3 unless noted otherwise
maxdsiz	0x10000000 or 268435456 (256MB)	0x10000000 or 268435456 (256MB)
maxfiles	512	1024
nproc	1005	1005
maxuprc	400	400
nfile	1024	2048
ninode	1085	1085
semms	2000	2000
semgni	2000	2000

To collect kernel parameter information before you upgrade the Process Engine software.

Kernel Parameter	Minimum Setting PA-RISC	Minimum Setting Integrity V2 and V3 unless noted otherwise
shmmax	0x10000000 or 268435456 (256MB)	0x20000000 or 536870912 (512MB)
shmseg	120	120
semmnu	1000	1000
semume	500	500
msgmni	2048	2048
msgseg	16384	16384 (obsolete in HP 11i V3)
msgtql	6640	6640
msgmap	msgtql + 2	msgtql + 2 (obsolete in HP 11i V3)
dbc_max_pct	1 to 10 (the value can not be greater than 30)	1 to 30 (the value can not be greater than 30) (obsolete in HP 11i V3)
dbc_min_pct	5	5 (obsolete in HP 11i V3)
timezone	Set appropriately	Set appropriately
msgmnb (DB2 only)	65535	65535
msgmax (DB2 only)	65535	65535

NOTE If you will be installing the 8.2 version of DB2 Client software, pay particular attention to the shmmax parameter. The minimum value documented for Process Engine might not be high enough to allow successful installation of the DB2 software. See the appropriate vendor documentation for the 8.2 release for recommended kernel parameter settings.

To collect kernel parameter information before you upgrade the Process Engine software.

- Before you upgrade, run the kmtune utility to collect current kernel parameter information and save the output. After the upgrade, use the HPjconfig utility to collect information on required patches and recommended kernel parameters. The HPjconfig utility's recommendations for the kernel parameters are based on analysis of historical data on the server. See the HP website for configuration details

To check and optionally modify the timezone settings

HP-UX has two timezone settings: the kernel parameter `timezone` and the environment variable `TZ`. The value of both timezone settings must match. Review and, if necessary, change these settings on all servers.

1. As the root user, enter:

```
sam
```

2. Select the Kernel Configuration option.
3. Select the Configurable Parameters option, then check the Pending Value for the `timezone` parameter. The default value is 420 minutes west of Greenwich Mean Time (GMT), which is the U.S. Mountain timezone.

Determine the number of minutes east or west of GMT for your location by multiplying the number of hours east or west of GMT by 60 minutes per hour. For example, the U.S. Pacific timezone is 8 hours west of GMT. Multiply 8×60 to get 480 minutes. If your timezone location is east of GMT, you should use a negative number. For example, Middle European Time is one hour east of GMT. Multiply -1×60 to get -60 minutes for MET (Middle European Time).

4. If the Pending Value for the `timezone` parameter is correct, proceed to step 6. To change the value, continue with Step 5.
5. To change the `timezone` kernel parameter value:
 - a. Select the `timezone` parameter by pressing the spacebar and then press Tab to go to Actions menu.
 - b. Select the Modify Configurable Parameter option from the Actions menu and press Return.
 - c. In the popup window that displays, the Specify New Formula/Value option should already be selected.
 - d. Tab to the Formula/Value field and type the new value.
 - e. Tab to OK and press Return. When the popup window disappears, you should see the new value in the Pending Value column.
 - f. Rebuild the kernel to make your change take effect:
 - i. Press the F4 key to access the menu bar.
 - ii. From the Action menu, select the Create a New Kernel option using the Arrow keys and press Return.
 - iii. Answer Yes when prompted about creating the kernel now.
 - iv. On the next screen, make sure the Move Kernel into Place and Shutdown/Reboot the System Now option is selected, tab to OK and press the Return key to reboot the system and make the new changes take affect.

6. As the root user, enter the following to check the current value of the `TZ` environment variable:

```
echo $TZ
```

7. If the current setting is not correct, enter the following to set the correct timezone:


```
/sbin/set_parms timezone
```

Choose the appropriate timezone from the menus displayed. Remember that the value must match that of the timezone kernel parameter.

If you change the current setting, you will be prompted to reboot the server.

NOTE If the HP-UX `set_parms` command is not available on your server, the timezone might be set via the SAM interface using the Kernel Parameters option in the same manner that other parameters are set. The System Administrator should consult the HP-UX operating system documentation to determine the appropriate way to set the TZ environment variable.

Configure Process Engine servers (Solaris)

The operating system prerequisites in this subsection pertain to Solaris-based FileNet PE servers.

To enable ports

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons. By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet uses several ports higher than 32768. See [“IBM FileNet P8 ports” on page 265](#) for details on which ports IBM FileNet uses.

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

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

1. To determine the current `tcp_smallest_anon_port` setting, enter the following at the command prompt:

```
ndd -get /dev/tcp tcp_smallest_anon_port
```

If the port is less than 32778, you must enable port 32777.

2. To enable port 32777 on Solaris9, use a text editor to edit the `/etc/rc2.d/S69inet` file.

Enter the following line:

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

3. To enable port 32777 on Solaris10, use a text editor to edit the `/lib/svc/method/net-init` file.

Enter the following line:

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

NOTE Put this entry in the file before the `exit 0` entry at the bottom of the file.

4. Reboot the Process Engine server to force the release of ports required by Process Engine that might be in use by the OS. Failure to reboot after these changes are made can result in port 32776 being unavailable, generating OpenSocket errors.

To verify national language character set and time settings

- The default time mask varies on UNIX depending on the LANG and LC_TIME environment settings. Verify the current LC_TIME settings by entering:

```
locale -k t_fmt
```

The result might appear similar to this:

```
t_fmt=%r
```

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

```
LC_TIME=C
export LC_TIME
```

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

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

To increase the operating system kernel limits

- Make a copy of the system file (with a new name). Log on as root, and enter a command similar to the following:

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

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

```
vi /etc/system
```

- Ensure that the following parameters are listed and are set to at least the values shown.

Solaris 9

```
set semsys:seminfo_semmap=50
set semsys:seminfo_semmni=2000
set semsys:seminfo_semmns=2000
set semsys:seminfo_semmnu=500
set semsys:seminfo_semmsl=512
set semsys:seminfo_semopm=256
```

```
set semsys:seminfo_semume=500
set semsys:seminfo_sevmx=32767
set semsys:seminfo_semaem=16384
set shmsys:shminfo_shmmax=4294967295*
set shmsys:shminfo_shmmni=0
set shmsys:shminfo_shmmni=2000
set shmsys:shminfo_shmseg=100
set msgsys:msginfo_msgmni=2048
set max_nprocs=1000
set fnsod:sod_Debug=0
set rlim_fd_max=1024
set rlim_fd_cur=256
noexec_user_stack=1
```

Solaris 10

```
set semsys:seminfo_semmni=2000
set semsys:seminfo_semmni=512
set semsys:seminfo_semmni=256
set shmsys:shminfo_shmmax=4294967295*
set shmsys:shminfo_shmmni=2000
set msgsys:msginfo_msgmni=2048
set max_nprocs=1000
set fnsod:sod_Debug=0
set rlim_fd_max=1024
set rlim_fd_cur=256
noexec_user_stack=1
```

* FileNet recommends this value be set to 4GB but do not set it higher than physical memory. It is recommended that this be set to less than 80% of physical memory.

4. Save your changes.
5. Reboot the server.

Configure Process Engine clients for ORB (all UNIX)

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 or custom applications

Therefore, if you have UNIX-based non-J2EE or custom applications, configure the Java installation on those servers with either the IBM or the Sun ORB.

Configure Application Engine or Workplace XT servers (Linux)

To configure Linux-based Application Engine or Workplace XT servers

Ensure that Linux® libraries are installed. To install Application Engine or Workplace XT on Linux, several legacy libraries are required. You must install the compat-libstdc++ packages on your RedHat system prior to beginning your install of Application Engine or Workplace XT.

Configure Microsoft Windows

Configure Windows for FileNet P8 servers

To configure Windows FileNet P8 servers

- **Upgrade to Windows 2003 before you begin the upgrade process to IBM FileNet P8 4.5.x.** See the *IBM FileNet P8 Hardware and Software Requirements* for details on any required Windows 2003 Service Packs and patches. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.](#)

NOTE IBM FileNet field personnel indicate that most customers prefer to do a fresh installation of Windows 2003 rather than upgrading the Windows 2000 software. This will require you to install Content Engine 4.5.x on its own Windows 2003 server (separate from the 3.5.2 Content Engine server).

- **(3.5.x to 4.5.x upgrades only) Ensure proper upgrade to Windows 2003 on your existing Content Engine 3.5.2 servers prior to IBM FileNet P8 upgrade.** If you do want to upgrade to Windows 2003 on your Content Engine 3.5.2 servers and intend to run them for any length of time before you upgrade to Content Engine 4.5.x, be sure to see the “Configure Content Engine and SQL Servers for Windows 2003” in the *IBM FileNet P8 Platform Installation and Upgrade Guide v 3.5.x* for important details and procedures. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.](#)

- **Ensure that you have the necessary Windows clients for Enterprise Manager 4.5.x.** Although Content Engine server runs on UNIX as well as Windows, the Enterprise Manager administrative client still runs on Windows only.

For upgrade purposes, it is a best practice to install Enterprise Manager 4.5.x on a separate client than the current 3.5.x or 4.0.x Enterprise Manager, so that you can access the version 3.5.2 or 4.0.x object stores that have not been upgraded yet. If you install Enterprise Manager 4.5.x on a client running the previous version, the older version will be upgraded.

- Consult with the application server, database, and P8 administrators to determine port requirements for all the servers in your install environment. For details, see [“IBM FileNet P8 ports” on page 265.](#)

Content Engine and Enterprise Manager

To configure Windows for Content Engine and Enterprise Manager

1. On any Windows machine where you are going to install Content Engine or Enterprise Manager (Content Engine is a Java application that can optionally be installed on Windows; Enterprise Manager must be installed on Windows), you must first install the following:
 - Microsoft .NET Framework
 - Web Services Enhancements (WSE)

Check the IBM FileNet P8 Hardware and Software Requirements for the latest version requirements of these two components. To download this guide from the IBM support page,

see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
Enterprise Manager requires no other Content Engine services or files.

Content Engine authenticating with Active Directory

To configure Windows for Content Engine

- If Windows Active Directory is your directory service, set the primary DNS server IP address on your Content Engine machine to the IP address of the machine where DNS is installed.

Process Engine on Windows

To install Process Engine using a domain user

- If Process Engine will be installed by a domain user rather than a local user on the associated server, see [“Specify IBM FileNet P8 accounts” on page 65](#) for details on creating required users and groups.

To verify TCP/IP parameter settings (Windows)

1. Log on as the Administrator user and run regedit to verify the following registry key values.

NOTE These values are decimal. The default in regedit is hexadecimal.

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\MaxUserPort => 65534 (default = 5000)

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpTimedWaitDelay => 90 (default = 240, or 4 min)

2. If necessary, add or modify a new DWORD value with the values as described above and save the changes.

To add inbound rules to the Windows 2008 firewall

Configure inbound rules in the Windows 2008 firewall to allow the following Process Engine ports access.

Port	Protocol
32768	TCP
32769	TCP
32770	UDP

To configure the /etc/hosts file

Information must be entered into either the server's DNS table or the hosts file related to Process Engine IP address, server name and NCH domain name. For non-farmed configurations this

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

Entries must be 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 nch_domain-organization-nch-server load_balancer_name

where:

IP addr is the IP address of the Process Engine server.

hostname is the corresponding host name.

nch_domain-organization is the NCH domain and organization name, as provided to the Process Engine installation program.

load_balancer_name is the name of the load balancer in a farmed configuration.

When entering the domain-organization name, follow these rules:

- Eliminate all characters except ASCII alphanumeric characters and underscores.
- Convert all upper case characters to lower case.
- Insert a hyphen between the domain and organization names.
- Append “-nch-server” as a literal.

For example, a Process Engine has a domain “ace-1” and organization “FileNet”. Its IP address is 123.45.6.78. For this system, the hosts file entry is:

123.45.6.78 ace-1 ace1-filenet-nch-server

NOTE The hyphen in the *nch_domain-organization* name has been removed and the F and N in the “FileNet” organization name have been converted to lower case.

To configure Process Engine clients for ORB (Windows)

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

Therefore, if you have Windows-based non-J2EE and custom applications, configure the Java installation on those servers with either the IBM or the Sun ORB.

Content Search Engine on Windows

To configure security for Content Search Engine servers

Set *K2 Operating System User* as an administrator on each Windows Content Search Engine Server. See [“Specify IBM FileNet P8 accounts” on page 65](#) for details on creating required users and groups.

NOTE For Content Engine 3.5.x to 4.5 upgrades, where CE 4.5 will be on UNIX, Content Search Engine 4.5 must also be installed on a UNIX server. Ensure that the accounts required include the UNIX operating system accounts.

Configure network

The requirements in this Task apply to the network on which FileNet P8 servers are running.

To configure 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 FileNet P8 so that they can all communicate with one another.
- Ensure NetBIOS over TCP/IP is enabled on Windows.
- Ensure availability of required port numbers. Several port numbers are required by the various IBM FileNet P8 components. See ["IBM FileNet P8 ports" on page 265](#).
- For information about proxy/firewall configuration requirements, see the *IBM FileNet P8 Hardware and Software Requirements* for support information related to IBM FileNet P8 components and database engines. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19](#).

To configure time and date

Synchronize the time and date on all servers. 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 match the UTC time on the Process Engine database server, plus or minus 15 minutes.

- 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 need to 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.

Assign directory permissions for Content Engine upgrade for 4.0.x to 4.5 on UNIX

For upgrades from Content Engine 4.0.x to 4.5 on UNIX, you change the installation user.

To assign directory permissions to the user (*ce_upgrade_user*) who will upgrade Content Engine

1. Determine which user installed the current version of Content Engine.

If you don't know who installed the current version, log on to the application server as any user and inspect the properties of the Content Engine install directory (the default install directory is ContentEngine).

NOTE (AIX only) Only the root user is able to install the current (4.0) version of Content Engine. This restriction is not in effect for the new (4.5) version.

2. Designate the user, referred to as *ce_upgrade_user*, who will install the new version of Content Engine. The additional preparation, if any, that you need to perform depends on two factors:
 - Whether the user who installed the current version of Content Engine is identical to the user who will install the new version.
 - The type of operating system on the machine that hosts the application server where Content Engine is deployed.

Continue at [Step 3](#) if either of the following conditions holds:

- The two users are not identical, and Content Engine is deployed on UNIX.
- The user who will install Content Engine is non-root, and Content Engine is deployed on AIX.

If neither of the above conditions holds, skip the remainder of this procedure.

3. Log on to the application server machine as the user who installed the current version of Content Engine.
4. Navigate to the ContentEngine directory and recursively give ownership of this directory and all its files and subdirectories to the user who will install the new version of Content Engine.

For example, if Content Engine is at /opt/FileNet/ContentEngine, then to give ownership to *ce_upgrade_user*, the user who will install the new version, run the following command:

```
chown -R ce_install_user /opt/FileNet/ContentEngine
```

5. Delete all files and directories in the /tmp directory owned by the previous *ce_install_user*.
6. Logon to the server as the user who initially installed the application server and created the application server instance (profile for Websphere, domain for WebLogic, server for JBoss).
7. Navigate to the application server instance directory, and give group rights to the *ce_appserver_install_group* (who's members are the *ce_install_user* and *ce_appserver_install_user*). For example:

```
chgrp -R ce_appserver_install_group /opt/IBM/Websphere/Appserver/profiles/Appsvr01
```

8. Give *ce_appserver_install_group* read/write permissions. For example:

```
chmod -R 775 //opt/IBM/WebSphere/Appserver/profiles/Appsvr01
```

9. Log off the application server machine and log back on as the user who will install the new version of Content Engine.
10. Grant read, write, and execute permissions on the Content Engine directories to the user who will install the new version of Content Engine, as follows:

```
chmod -R +rwx /opt/FileNet/ContentEngine
```

11. Copy the InstallShield directory from the \$HOME directory of the user who installed the current version of Content Engine to the \$HOME of the user who will install the new version and grant read, write, and execute permissions to the user who will install the new version.

Process Engine

To prepare Process Engine for upgrade

- **(3.5.x to 4.5.x) Reconcile the Process Engine user security information.**

The Process Engine duplicates certain parts of the user security information in its own database. Over time, the information in the directory service might be changed or updated. When this happens, the information in the Process Engine's environment records, whether cached or permanent, can end up containing old, invalid information about the Process Engine users and groups.

WARNING During the Process Engine upgrade it is critical that this user information is correct and up to date. Before upgrading you must reconcile the cached and permanent user data environment records on the Process Engine with the possibly more-up-to-date data in an LDAP-based directory service.

For more information, see the IBM FileNet P8 help topic [FileNet P8 Administration > Process Engine Administration > Administrative tools > vwtool > Commands > environment](#).

- If you are using the Process Analyzer expansion product, several steps must be taken on the Process Analyzer, on the Process Engine database, and on Process Engine, before upgrading Process Engine. See ["Complete pre-upgrade Process Engine configuration"](#) on page 292 of the *IBM FileNet P8 Platform Installation and Upgrade Guide* for details.

Security Administrator upgrade tasks

As the Security Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform upgrade:

- Review all rows assigned to the Security Administrator (SA) for upgrades in the [“Installation and upgrade worksheet” on page 216](#). Provide values for any rows appropriate to your installation that you have not yet completed. For descriptions of the properties in the Installation Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- Review the following planning considerations:
 - [“Security upgrade planning considerations” on page 181](#)
- Create the operating system accounts required by FileNet P8. Refer to the following topic:
 - [“Specify IBM FileNet P8 Accounts for upgrades” on page 183](#)

Security upgrade planning considerations

- **(3.5.x to 4.5 upgrades only) If you intend to migrate to a different directory service, do so prior to the 4.5 upgrade.** If you are expecting to migrate to a different directory service for use with IBM FileNet P8 Platform 4.5, you must do so before you begin the upgrade from 3.5.x. See the *IBM FileNet P8 Platform 3.5.x Directory Service Migration Guide* for complete information. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.](#)

NOTE Migrating to a different directory service is not supported if you begin your upgrade from the 4.0.x release or once you are running the 4.5 release. However, this capability is likely to be provided in subsequent releases. Contact your IBM FileNet representative if you intend to switch to IBM Tivoli Directory Server, which is supported for IBM FileNet P8 4.5, but not for 3.5.x.

- **Be sure to understand updated user and group account requirements for release 4.5.** The 4.5 release requires you to designate or create several new accounts (for example, the Configuration Manager user). For details, see [“Specify IBM FileNet P8 Accounts for upgrades” on page 183.](#)
- **(3.5.x to 4.5 upgrades only) Determine whether you intend to use Secure Socket Layers (SSL).** It is a best practice to use SSL to strengthen security. Configuring SSL for 4.5 is noticeably different than for 3.5.x. For example, Content Engine 4.5 now relies on its own application server to direct calls to your chosen authentication provider. For configuration details, see [“Set up Content Engine and client transport SSL security” on page 157](#) and [“Set up Application Engine SSL security” on page 160.](#)
- **(3.5.x to 4.5 upgrades only) Gather 3.5.x authentication information and plan 4.5 authentication.**

In the 3.5.x release, Content Engine authentication and authorization information are both specified by configuring one or more authentication providers, using Active Directory, Novell eDirectory, or Sun Java System Directory Server.

- For Novell and Sun authentication providers, the root of the configured 3.5.x Directory Service contains one or many children, or naming contexts, each of which are used automatically by 3.5.x Content Engine as FileNet P8 authentication realms.
- For Windows Active Directory, the 3.5.x Content Engine uses the Windows domain (also known as *deployment domain*) it resides in and automatically gets all trusted domains, siblings, parents, children within the domain's forest. By default, the deployment domain is the *default realm*.

3.5.x authentication realms are viewable (but not editable) in Enterprise Manager's root domain property sheet > Authentication Provider tab > DefaultRealm property. The names of all realms for the current FileNet P8 domain appear as values in the drop-down list for this property.

NOTE If you are using the 3.5.x Content Engine attributes DefaultRealm and RestrictToDefaultRealm to modify 3.5.x authentication behavior, then you must carefully map out which 3.5.x authentication realms are active and essential to your upgraded 4.5 Content Engine authentication scheme. The DefaultRealm and RestrictToDefaultRealm attributes are

not supported by Content Engine 4.5, because of the introduction of the application server, which uses Java Authentication and Authorization Service (JAAS) for authentication.

By contrast, in Content Engine 4.5, configuring authentication and authorization are two separate steps, both of which must be completed on the Content Engine 4.5 server environment before upgrading object stores from 3.5.x. Authentication is configured through the application server's administration console. Authorization is configured by creating one or more Directory Configuration objects in Content Engine 4.5 for the FileNet P8 domain. (The 3.5.x authentication information is not carried forward by the CE 3.5.2 to 4.5 Upgrader tool.) Therefore, the following are prerequisites to upgrading a 3.5.x domain:

- Install and configure one of the supported J2EE application servers, using the application server's install tools.
- Install a Content Engine 4.5 server as an application into the application server, using the Content Engine installer.
- Configure the authentication environment through Configuration Manager.
- Create a Content Engine 4.5 domain, using the Enterprise Manager's Add Domain Configuration wizard.
- Create one or more Directory Configurations, using the Enterprise Manager's Create a Directory Configuration wizard. One Directory Configuration is required for each distinct 3.5.x authentication realm.

There are several cases that might exist for this last bullet. In all cases, Directory Configurations must be created that preserve access to all users who have used the object stores being upgraded. Access to all users' SIDs and group information is necessary for authorization to occur.

- In the typical case, you will create one 4.5 Directory Configuration object for each realm used in 3.5.x, specifying the same directory server (LDAP) host and port information that was used in Content Engine 3.5.x.
 - However, in some cases customers will choose to do the following:
 - Configure the Content Engine 4.5 server to obtain directory configuration information from a different directory server than was used in Content Engine 3.5.x (for instance by using a replica of the 3.5.x directory service which contains the same user or group information).
- OR
- Expand the users of the Content Engine 4.5 domain being upgraded by adding additional realms. If this is the case, then the set of Directory Configurations that must be created for the new 4.5 domain could differ from the set of realms that were configured in 3.5.x. Just make sure that the 4.5 Directory Configurations contain the same user and group principals, with the exact same SIDs, as those used in Content Engine 3.5.x.

Remember that users in the newly added realms will have no authorization to access any Content Engine 4.5 objects until they are explicitly granted authorization to do so, typically by using Enterprise Manager in ways described in the *Help for Content Engine Administration*.

Specify IBM FileNet P8 Accounts for upgrades

The following procedures direct you to create or designate the accounts needed to upgrade 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 [FileNet P8 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 FileNet P8 user interface, such as a setup 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. An example of this is *ce_install_user*, the account you log in as to run the Content Engine setup program.

- By a name that looks like a variable but is not formatted in italics. Examples are the Process Engine's required accounts `f_maint` and `f_sw`, which, because they are not italicized, are not to be replaced by accounts of your choosing.

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 for that reference and find the account table that defines it.

Accounts for Content Engine upgrade

Task to be performed by: Database Administrator

To create Content Engine database accounts for upgrades from version 3.5.x

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

User/Group	Description
Database user name: DB2 <i>ce_db_user</i>	<p>The database owner accounts that Content Engine uses to access DB2. Create a new account for the GCD tablespace. Your existing version 3.5.x object store user accounts do not require changes for the upgrade.</p> <p><i>ce_db_user</i> can be an operating system user on the database server. In the case of a remote database, no equivalent users are needed on the Content Engine server.</p> <p>Grant each <i>ce_db_user</i> at least the following DB2 database access permissions:</p> <ul style="list-style-type: none"> • Connect to the database • Create tables in the tablespace (CREATETAB) • Use the tablespace (USE OF) for User and User Temp tablespaces

User/Group	Description
<p>Database user name: Microsoft SQL Server <i>ce_db_user</i></p>	<p>The database owner accounts that Content Engine uses to access SQL Server, depending on whether you use one account for all Content Engine databases, or use one (for example, <i>ce_db_user</i>) for the GCD database and different accounts for each object store (for example, <i>ce_db_user1</i>, <i>ce_db_user2</i>, and so on).</p> <p><i>ce_db_user</i> can be a local account or a Windows domain account. It does not have to be an account in the configured directory service.</p> <p>For each object store you are going to upgrade, grant the associated 3.5.x database user at least the following additional roles and access permissions:</p> <p>SQL Server roles</p> <ul style="list-style-type: none"> • System Administrators • Security Administrators • Database Creators <p>Grant each <i>ce_db_user</i> at least the following database access permissions:</p> <ul style="list-style-type: none"> • public • db_owner <p>Add these accounts to SQL Server's master database and grant the public role to each. When you perform the procedure "To configure the JDBC Distributed Transaction Components" on page 200 these accounts will also be granted the SqlDBCXAUser role.</p>
<p>Database user name: Oracle <i>ce_db_user</i></p>	<p>The tablespace owner accounts that Content Engine uses to access Oracle. Create a new account for the GCD tablespace. Your existing object store user accounts do not require changes for the upgrade.</p> <p><i>ce_db_user</i> does not have to be an LDAP account.</p> <p>Grant each <i>ce_db_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> • CREATE SESSION • CREATE TABLE • CREATE SEQUENCE (object store tablespaces only)

Task to be performed by: Application Server Administrator

To create Content Engine application server accounts for upgrade

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

User/Group	Description
<p>Application server administrator user name</p> <p>Administrative Console User</p> <p><i>ce_appserver_admin</i></p> <p><i>ce_appserver_console_admin</i></p>	<p>An application server administrative account or accounts that can log on to the application server administration console.</p> <p>WebSphere: If your site uses a FederatedLDAP registry, <i>ce_appserver_console_admin</i> must be a unique user across all federated realms. If not, <i>ce_appserver_admin</i> and <i>ce_appserver_console_admin</i> can be the same account.</p> <p>WebLogic: <i>ce_appserver_admin</i> and <i>ce_appserver_console_admin</i> must be different accounts.</p> <p>(JBoss does not require an administrative account.)</p> <p>Login to the application server as <i>ce_appserver_admin</i> to perform application tasks such as the following:</p> <ul style="list-style-type: none"> • Create directory service providers within the application server. • Create JDBC Providers and/or data sources for db connectivity. • Deploy the Content Engine application. • Stop and restart servers and cluster members (WebSphere). • Stop and restart managed servers (WebLogic).

User/Group	Description
Application Server Installation Administrator: WebSphere WebLogic JBoss <i>ce_appserver_installer</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_installer</i> at least the following permissions:</p> <ul style="list-style-type: none"> For Windows, <i>ce_appserver_installer</i> must be a member of the local administrators group. For UNIX, <i>ce_appserver_installer</i> must have read, write, and execute permissions to the Content Engine installation directory. <p>Log in as <i>ce_appserver_installer</i> to do the following:</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. This includes creating UNIX mounts to external NFS devices such as SnapLock shares, Autonomy K2 Content Search Engine collection shares, and NFS->NTFS mounts to access NTFS shares for Content Engine 3.5.2 upgrades when migrating to Content Engine 4.5 on UNIX. <p><i>ce_appserver_installer</i> must be a member of the "ce_appserver_install_group" on page 188.</p>

User/Group	Description
<p>Application Server Installation Group</p> <p><i>ce_appserver_install_group</i></p>	<p>Create an operating system group account and add to it the following accounts:</p> <ul style="list-style-type: none"> • “ce_appserver_admin” on page 186 • “ce_install_user” on page 191 user (Windows) or “ce_install_user” on page 191 (UNIX) . <p>Use the user accounts in <i>ce_appserver_install_group</i> to do the following:</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.

Task to be performed by: IT Administrator

To create Content Engine operating system accounts for upgrade

1. Create new or designate existing installation accounts for Content Engine, as shown in the following table:

User/Group <i>unique_name</i>	Description
Configuration Manager user: WebSphere WebLogic JBoss <i>config_mgr_user</i>	<p>The operating system account you will use to run the Configuration Manager.</p> <p>(Windows) Using Active Directory tools, add <i>config_mgr_user</i> to either the Power Users group or the Local Administrators group.</p> <p>After Content Engine is installed you must grant additional permissions to <i>config_mgr_user</i>:</p> <ul style="list-style-type: none"> • Grant it permission to execute the Configuration Manager executable file: configmgr.bat (Windows) or configmgr.sh (UNIX). • Grant it permission to write to the directory where Configuration Manager will create the configuration XML files. This directory is one of the following: <ul style="list-style-type: none"> – the directory you specify using the optional -path parameter when you run the tool – the default directory, <CE_install_path>/tools/configurationmanager/tasks, if you don't specify a path parameter • Grant it permission to read and write the contents of the directory specified by -path parameter, or the default directory. <p>You will be instructed to grant additional permissions to <i>config_mgr_user</i> at several times in the <i>Planning and Preparation Guide</i>:</p> <ul style="list-style-type: none"> • “To set permissions for the Configuration Manager user” on page 19 (WebSphere) • “Assign directory permissions” on page 139 (JBoss) • “Give the Configuration Manager user the following permissions:” on page 137 (WebLogic)

Task to be performed by: Security Administrator

To create Content Engine directory server accounts

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

User/Group	Description
<p>Bootstrap user name</p> <p><i>ce_bootstrap_admin</i></p>	<p>An directory service and application server account that is stored in the CEMPBoot.properties 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>WARNING If you are deploying Content Engine on an application server (WebSphere 6.1.x or higher, WebLogic, or JBoss) with federated user repositories and with multiple realms in your 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 FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups for information on how <i>ce_bootstrap_admin</i> is the creator of the new P8 domain.</p> <p>This user must also have read/write access to the root directories for the P8 3.5.x file stores.</p> <p>If version 3.5.x file stores are on a SnapLock fixed content device, then this user must also have read/write access to the device.</p>
<p>Content Engine Upgrader tool user</p> <p><i>ce_upgrade_user</i></p>	<p>The Windows operating system user who runs the Upgrader tool needs the following permissions:</p> <ul style="list-style-type: none"> • Read access to the Content Engine 3.5.x directory (default location is C:\Program Files\FileNet\Content Engine). • Read access to the 3.5.x version of sysinit.dat. • Read/write access to the root directories for the file storage areas

User/Group	Description
<p>Content Engine Setup: Windows</p> <p><i>ce_install_user</i></p>	<p>For upgrading from version 3.5, you can use the existing Content Engine install user, or designate a new user.</p> <p>An operating system account you will use to log on to a machine to launch Content Engine install wizard.</p> <p><i>ce_installer</i> can be the same user as the “ce_appserver_admin” on page 186.</p> <p>Use Windows administrative tools to add <i>ce_installer</i> to the Local Administrators group and to the “ce_appserver_install_group” on page 188.</p>
<p>Content Engine Setup: UNIX</p> <p><i>ce_install_user</i></p>	<p>For upgrading from version 3.5, you can use the existing Content Engine install user, or designate a new user.</p> <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 /tmp directory. • Membership in the “ce_appserver_install_group” on page 188. <p><i>ce_install_user</i> can be the same user as the “ce_appserver_admin” on page 186.</p>

User/Group	Description
GCD Administrator <i>gcd_admin</i>	<p>A directory service account that has Full Control access to the Content Engine's domain object.</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>Logon 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. • Carry out administrative tasks for the FileNet P8 domain. <p>To perform upgrades, at least one FileNet P8 4.5 GCD administrator must:</p> <ul style="list-style-type: none"> • Also be an object store administrator (<i>object_store_admin</i>) for all object stores to be upgraded. Use P8 3.5.x FileNet Enterprise Manager to ensure that the <i>gcd_admin</i> also has Full Control to each P8 3.5.2 object store. • Have Full Control for the FileNet P8 3.5.x and FileNet P8 4.5 domains. <p>For more information, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups.</p>
Object Store Administrator <i>object_store_admin</i> <i>object_store_admin_group</i>	<p>You do not need to create new object store administrators for an upgrade. However, you can create new object store administrators as needed after completing the upgrade.</p> <p>The <i>object_store_admin</i> is 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>

User/Group	Description
<p>Directory service user: Active Directory</p> <p>(Referred to as “Directory service bind user name” in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<p>An Active Directory user account that Content Engine uses to connect to Active Directory.</p> <p><i>ce_service_user</i> performs the following roles:</p> <ul style="list-style-type: none"> Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace. Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object once a user has been authenticated. <p>Using Active Directory tools, grant <i>ce_service_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> Member of the Pre-Windows 2000 Compatible Access Group in each desired domain in the Active Directory forest. <p>Provide the fully qualified distinguished name of <i>ce_service_user</i> as the Directory service bind user name while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>
<p>Directory service user: Windows Active Directory Application Mode (ADAM)</p> <p>(Referred to as Directory service bind user name in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<p>An ADAM user account that Content Engine uses to connect to a single Microsoft ADAM partition.</p> <p><i>ce_service_user</i> performs the following roles:</p> <ul style="list-style-type: none"> Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace. Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object once a user has been authenticated. <p>Using ADAM administrative tools, grant <i>ce_service_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> Ability to see the other users in the partition. (For a procedure, see the entry for the ADAM directory service user in FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups.) <p>Provide the fully qualified distinguished name of <i>ce_service_user</i> as the <i>LDAPBindDN</i> while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>

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

User/Group	Description
<p>Directory service user: IBM Tivoli Directory Server</p> <p>(Referred to as Directory service bind user name in Configuration Manager.)</p> <p><i>ce_service_user</i></p>	<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 FileNet P8 object once 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 <i>LDAPBindDN</i> while running the Configuration Manager tool and also when you run Enterprise Manager's Directory Configuration Wizard.</p>

Database Administrator upgrade tasks

As the Database Server Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform:

- Review all rows assigned to the Database Server Administrator in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. For descriptions of the properties in the Installation and Upgrade Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- (SQL Server only). Upgrade SQL Server 2000 to SQL Server 2005. Refer to your Microsoft documentation.
- Review the following planning considerations:
 - [“Database Administrator planning considerations” on page 197](#)
- Create the GCD database for the upgrade. Refer to the following topic:
 - [“Create the GCD data source for Content Engine \(upgrades from 3.5.x\)” on page 199](#)
- Update the CFS database user for CFS fixed content devices as follows.

If your version 3.5.x or 4.0.x CFS fixed content devices use MS SQL Server databases, then perform the following steps for each database.

1. For each CFS database, add the following roles to the CFS user:

- System Administrators
- Security Administrators
- Server Administrators
- Database Creators

2. For each CFS database, verify the following access permissions for the CFS database user:

- public
- db_owner

Database Administrator planning considerations

Review the following database planning considerations before you begin preparing for the upgrade.

General

- **Update to the appropriate database patches before you upgrade IBM FileNet P8 components.** For minimum patch requirements, see *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
- **Update database versions after you upgrade IBM FileNet P8 components.** If you intend to upgrade the version of the database, complete the upgrade of all IBM FileNet P8 components and verify a fully functional system before upgrading the database software.
- **(3.5 to 4.5 upgrades only) Plan for upgrade implications regarding Content Engine databases.** If you are upgrading to version 4.5, you will continue to use existing object store databases. However, if you are upgrading from version 3.5.x, you must create a new database for the 4.5 global configuration data (GCD), which the Content Engine Upgrader Tool migrates from its 3.5.x file-based format (sysinit.dat).
- **(3.5 to 4.5 upgrades only) Consider the following for Business Process Manager components (PE, PA, PS):**
 - The Process Engine database upgrade automatically renames the existing VWLogxxx_yyy tables to VWLogxxx_yyy_Archive and creates the new VWLogxxx_yyy table.
 - The Process Engine Log SeqNumber starts from 21000.
 - Process Analyzer gets its data from the Process Engine database. The 3.5.x and 4.0.0 database schemas changed both on Process Engine and Process Analyzer. The 3.5.x data from Process Engine must be transmitted to Process Analyzer 3.5.x before you upgrade either of these components to 4.5.0. For details, see

Oracle

- **Ensure you have applied appropriate Oracle patches to Oracle clients as well as servers.** Be sure that clients remote from Oracle database servers have patches that are comparable to the database server. Oracle clients include any machines remote from the Oracle database server where Process Engine software is installed. You can download all the required Oracle database server patches from the Oracle website and install them.
- **Consider increasing the number of Oracle sessions and lowering the session timeout settings.** Upgrading a Content Engine object store can use a large number of database sessions. The number of sessions required depends on the number of custom class definitions in the object store being upgraded. If the maximum number of sessions is exceeded during an object store upgrade, the upgrade will fail. To recover from such an upgrade failure, you will have to increase the number of sessions, decrease the session timeout settings, and then perform the following steps:
 - a. In the Content Engine Upgrader Tool, note the steps that completed successfully during the object store upgrade, that is, those items that have green check marks next to them.
 - b. Shut down the Content Engine Upgrader Tool.

- c. Shut down and then restart Oracle to release all database sessions and processes.
- d. Restart the Content Engine Upgrader Tool.
- e. Resume the object store upgrade and make sure to reset only the failed upgrade steps.

CAUTION To avoid a potential corruption of the object store, ensure that you have not chosen to rerun any upgrade steps that were already completed in the initial run of the Upgrader Tool.

- **Determine when to execute SQL scripts.** An upgrade SQL script must be executed. This script can be executed manually, before starting Process Engine Setup, or from Process Engine Setup. See [“Process Engine SQL scripts” on page 218](#) for details about the scripts, including information on execution modes and associated security requirements.

Microsoft SQL Server

(3.5 to 4.5 upgrades, Process Engine only) Create an ODBC data source. The 4.5 version of Process Engine uses an ODBC data source to connect to the database. The database administrator should create the data source before trying to upgrade the Process Engine software. See [“To create the Process Engine ODBC data source and test the connection \(SQL Server only\)” on page 293](#) for details on creating the data source and testing the database connection.

Create the GCD data source for Content Engine (upgrades from 3.5.x)

Prepare DB2

You must create a tablespace for the Content Engine GCD.

To create the DB2 tablespaces

Create a tablespace for the GCD with the following attributes:

IBM FileNet Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
GCD_ts (for the GCD database)		256		32 (required)
user temporary ts (for CE)		40		32 (required)
system temporary ts (for CE)		40		32 (required)

NOTE Record the tablespace name in the [“Installation and upgrade worksheet”](#) on page 216.

Prepare Oracle Server

You must create a tablespace for the Content Engine GCD.

To create tablespaces for the GCD

Using Oracle Enterprise Manager or SQL*Plus, create a user, password, and default tablespace in the Oracle database for the GCD that Content Engine will access. Grant CONNECT and RESOURCE roles to the user. These two roles combine to include the minimal privileges required by Content Engine: CREATE SESSION, CREATE TABLE, and CREATE SEQUENCE.

NOTE Because these two roles include other privileges as well, IBM recommends that you design your own roles if you prefer to grant only the minimal privileges required by the GCD.

Grant the following additional permission to the user:

- select on pending_trans\$
- select on dba_2pc_pending
- select on dba_pending_transactions
- execute on dbms_system

WARNING The Oracle user you create for the permanent and temporary tablespaces of the GCD must be unique. That is, the Oracle user for the GCD must not be the same as that of the user for any object store. Otherwise, the objects you intend to add only to the GCD will show up in all object stores that share the same Oracle user.

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

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

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

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

NOTE Record the tablespace name in the “[Installation and upgrade worksheet](#)” on page 216.

Prepare SQL Server

You must create a database for the Content Engine GCD.

To create a SQL Server database for the GCD

Create a SQL Server database for GCD database, which is required for Content Engine installation. Create the database with an initial size of 100MB, minimum. Make note of the database name as it will be required later when installing Content Engine software.

NOTE Record the database name in the “[Installation and upgrade worksheet](#)” on page 216.

To configure the JDBC Distributed Transaction Components

Execute these steps on every SQL Server that will have a Content Engine database.

1. Download the Microsoft SQL Server 2005 JDBC Driver that is referenced in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see “[Access IBM FileNet documentation, compatibility matrices, and fix packs](#)” on page 19.
2. Copy the sqljdbc_xa.dll from the JDBC installation directory to the Program Files\Microsoft SQL Server\80\Tools\Binn directory if there is only a default instance or the Program Files\Microsoft SQL Server\MSSQL\$instance nameBinn if you are using a named instance. If you are on a 32-bit

processor, use the sqljdbc_xa.dll file in the x86 folder. If you are on a 64-bit processor, use the sqljdbc_xa.dll file in the x64 folder.

3. Log on as a database administrator and execute the database script xa_install.sql on every SQL Server instance that will participate in distributed transactions. This script installs sqljdbc_xa.dll as an extended stored procedure and creates the SqlJDBCXAUser role in the Master database.

CAUTION Use SQL Server database credentials, not Windows credentials, to log on. Windows Integrated Logon to SQL Server is not supported with IBM FileNet P8.

4. Navigate to **Component Services > My Computer > Properties > , MSDTC > Security Configuration**, and make sure that XA transactions are enabled.

Application Server Administrator upgrade tasks

As the Application Server Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform:

- Review all rows assigned to the Application Server Administrator in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. For descriptions of the properties in the Installation and Upgrade Worksheet, see [“Installation and upgrade glossary” on page 228](#).
- Create the accounts required by FileNet P8, which are listed for the Application Server Administrator in the following task:
 - [“Specify IBM FileNet P8 Accounts for upgrades” on page 183](#)

Upgrades from 3.5.x

- Review the following planning considerations:
 - [“Application server planning considerations \(upgrades from 3.5.x\)” on page 203](#)
- Configure your application server for Content Engine according to your application server type. See the following task:
 - [“Configure the application server for Content Engine \(upgrades from 3.5.x\)” on page 205](#)

Application server planning considerations (upgrades from 3.5.x)

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication](#) for reference information about support for EJB and Web Services transports.

All platforms

Note the requirement for J2EE application servers. 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 (IBM WebSphere, BEA WebLogic, or JBoss) and their version numbers are identical for both components. Also, the applications must use Enterprise Java™ Bean (EJB) transport.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication](#) for reference information about support for EJB and Web Services transports.

Note that the Java Virtual Machine determines the maximum number of object stores. 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.

Note the impact of deploying Content Engine and other applications on the same machine. 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, as configurations will not match what has been qualified by IBM FileNet 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.

Note the impact of deploying Content Engine on multiple application server nodes. If you intend to deploy Content Engine to multiple server nodes in your production environment, the type of environment will determine how you install Content Engine, as described in “WebSphere and WebLogic” on page 33 and “JBoss” on page 33.

WebSphere and WebLogic - Content Engine Multi-Server Deployment

Note the impact of deploying centrally managed servers in farms or clusters. 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 (WebSphere) or the Administrator node (WebLogic). To install Content Engine in such an environment, see the IBM FileNet P8 Platform High Availability Technical Notice.

Note the impact of deploying centrally managed servers not in farms or clusters. In an environment in which multiple Content Engine instances may be geographically dispersed, and where each instance may have its own local Application Engine or Workplace XT server, you will

initially install Content Engine on the Deployment Manager node (WebSphere) or the Administrator node (WebLogic). You will then perform post-deployment steps using the administrative console of the application server to deploy Content Engine to other managed servers (see “Deploy Content Engine into a Managed Environment” on page 207).

JBoss Content Engine Multi-Server Deployment

Note the impact of deploying non-managed servers that are farmed or clustered. JBoss does not have a central management server. In an environment of multiple physical servers that are farmed or clustered, you will install Content Engine initially on one physical server (using the "all" instance), and then copy pertinent files and directories to the other servers. To install Content Engine in such an environment, see the IBM FileNet P8 Platform High Availability Technical Notice.

Note the impact of deploying non-managed servers that are not farmed or clustered. In an environment in which standalone application servers are individually managed and not farmed or clustered, and where the servers may be geographically dispersed, you will install Content Engine on each server. Each Content Engine instance will point to the same directory service and GCD database. The data sources you create for each instance will also point to the same object store databases.

Application server considerations

(3.5 to 4.5.x upgrades only) Determine whether the number of object stores to be upgraded will require you to run a 64-bit Java Virtual Machine (JVM) on your Content Engine application servers. If the 3.5.x FileNet P8 domain you are upgrading contains more than 50 object stores, it is a best practice to install Content Engine on application servers running a 64-bit, rather than 32-bit, JVM. Otherwise, you might experience significant Content Engine performance issues.

Configure the application server for Content Engine (upgrades from 3.5.x)

Set up the application server for Content Engine. See one of the following topics:

- “Configure WebSphere for Content Engine” on page 205
- “Configure WebLogic for Content Engine” on page 207
- “Configure JBoss for Content Engine” on page 211

Configure WebSphere for Content Engine

NOTE This task assumes you have already installed WebSphere Application Server on the machine where you are going to install and deploy Content Engine.

Content Engine is deployed into a profile in WebSphere Application Server. The profile serves as an application server environment for Content Engine. A default profile (named *AppSrv01* on WebSphere 6.1) is part of the initial WebSphere Application Server installation. You can deploy Content Engine into this profile, or create another profile for this purpose.

Perform the following procedure to specify the WebSphere environment variables for Content Engine and to set permissions on the profile directory for the user who will run Configuration Manager to deploy Content Engine. Among the variables you will specify is the JDBC driver corresponding to the type of database (DB2, Oracle, or SQL Server) where the Global Configuration Data (GCD) database will reside.

To specify the WebSphere environment variables

1. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see “[Access IBM FileNet documentation, compatibility matrices, and fix packs](#)” on page 19.
2. Install JDBC drivers on the WebSphere machine, as follows:
 - a. Obtain the JDBC drivers, depending on your database type.

DB2

Find the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](http://www.ibm.com) (<http://www.ibm.com>) by searching for “JDBC Type 4”.

Microsoft SQL Server

Find the Microsoft SQL Server Driver 2005 JDBC Driver, `sqljdbc.jar`, at Microsoft Support.

Oracle

Access the [Oracle JDBC Driver Downloads](http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) web site (http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) and find the JDBC driver file that matches the version of the JDK on the WebSphere machine.

- b. Copy the JDBC driver file from [Step a](#) to one of the following locations:

UNIX

/opt/jars

CAUTION Do not copy the file to ...WebSphere/AppServer/lib/ext.

Windows

C:\jars

CAUTION Do not copy the file to ...WebSphere\AppServer\lib\ext.

3. If you already have a profile for Content Engine, continue at [Step 4](#); or run the command script at one of the following (default) locations to create a new profile.

NOTE Make a note of your profile name, as you will need to specify it when you run Configuration Manager.

AIX

/usr/IBM/WebSphere/AppServer/wasprofile.sh

Other UNIX

/opt/IBM/WebSphere/AppServer/wasprofile.sh

Windows

C:\Program Files\IBM\WebSphere\AppServer\bin\wasprofile.bat

4. Start the WebSphere administrative console and log on to your profile as *ce_appserver_console_admin*, the Administrator Console User. For details on required accounts and related permissions, see [“Accounts for Content Engine” on page 66](#)
5. (WebSphere 6.1) Navigate to **Environment > WebSphere Variables** and perform these substeps to specify the JDBC driver path:
 - a. Choose the Cell scope entry from the *All scopes* drop-down list.
 - b. Click **New** to create a new WebSphere variable whose name is one of the JDBC environment variables shown in the following table, depending on your database type.
 - c. Set the value of the variable to the JDBC driver path you specified in [Step 2](#) and save your change to the master configuration.
 - d. Choose the Node scope entry from the *All scopes* drop-down list.
 - e. In the table of substitution variables, click the item in the Name column that corresponds to the JDBC environment variable in the following table for your database type.

Database	JDBC Environment Variable
MS SQL Server	WebSphere 6.x MSSQLSERVER_JDBC_DRIVER_PATH
	WebSphere 7.x MICROSOFT_JDBC_DRIVER_PATH
Oracle	ORACLE_JDBC_DRIVER_PATH
DB2	DB2UNIVERSAL_JDBC_DRIVER_PATH

- f. Set the value of the item to the JDBC driver path you specified in [Step 2](#), and save your change to the master configuration.
6. Navigate to **Servers > Application servers > server1 > Java and Process Management > Process Definition > Java Virtual Machine** and perform the following substeps:
 - a. Set the values for initial and maximum heap sizes (in megabytes), as follows, where *server1* is the name of the server in which you will deploy Content Engine:

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

- b. Save your changes to the master configuration.
7. If any of your object stores will be of a database type that differs from those whose JDBC environment variables you have already specified in this procedure, return to [Step 1](#); otherwise, continue at [Step 8](#).
8. (Optional) Navigate to **Servers > Application Servers > server1** and perform the following substeps to set the transaction timeout value:
 - a. Click the **Runtime** tab and then click **Transaction Service**.
 - b. Change the Total transaction lifetime timeout parameter value to at least 600 (seconds).

CAUTION If the timeout value is not large enough, some administrative processes (such as upgrading Content Engine from version 3.5.x or adding an expansion product) may fail.

To set permissions for the Configuration Manager user

1. Set the Configuration Manager user (*config_mgr_user*) permissions on the profile directory (and all its subdirectories) in which Content Engine will be deployed, as follows:

UNIX

Read, write, and execute permissions

Windows

Read & Execute, and Write permission

Configure WebLogic for Content Engine

NOTE This task assumes you have already installed WebLogic Server on the machine where you are going to install and deploy Content Engine.

To configure WebLogic

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

1. Use the WebLogic Configuration Wizard to create a WebLogic domain for Content Engine. In the examples below, use the domain name *FNCEDomain*. Keep the following in mind as you configure the domain:
 - a. Set the server start mode to Production mode.
 - b. Select the appropriate Java Development Kit (JDK) for your environment.
 - (Windows) The version of the JRockit SDK specified in the Third Party Support Information section of the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
 - (AIX) IBM SDK 1.5.0 at /usr/java5
2. Use the WebLogic Administration Console to set the following:
 - a. (Optional) Create a WebLogic authentication provider. You can use the Content Engine Configuration Manager tool to create a WebLogic authentication provider later on, or you can create the provider now using the WebLogic Console. If you want to configure the LDAP later with See “Configure Content Engine instances” in the *FileNet P8 Platform Installation and Upgrade Guide* for information on using the Configuration Manager tool to create the LDAP provider.

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

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

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

NOTE If performance problems are encountered, change the Group Membership Searching parameter value to *limited*.

- b. Do one of the following to allow or prohibit logons to *FNCEDomain* by LDAP-authenticated users in the DefaultAuthenticator who are not in *FNCEDomain*'s active security realm:
 - Set the Control Flag to SUFFICIENT to allow logons by users not in the active security realm (such as users in your authentication domain).

- Set the Control Flag to REQUIRED to prohibit logons by users not in the active security realm (such as users in your authentication domain).

If you choose REQUIRED, all the users you wish to authenticate for Content Engine must exist not only in the LDAP directory, but also exist as WebLogic users who are in the Default Authenticator provider.

- c. If you are using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers so that the most-frequently-used provider is first in the list, and the least-frequently-used is last. Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.
 - d. Specify the following heap sizes for the JVM:
 - Initial Java heap size (-Xms): 512 MB
 - Maximum Java heap size (-Xmx): 1024 MB
 - e. (Optional) Adjust transaction-timeout value. Content Engine relies on the transaction-timeout value, whose default may be too short for some standard or administrative processes (such as adding an expansion product or upgrading to the latest version of Content Engine). Set the JTA node timeout in seconds to at least 600 seconds.
3. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).
 4. Depending on your database, use one of the following procedures to install the JDBC drivers.

DB2

- a. Obtain the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](#) (<http://www.ibm.com>) by searching for “JDBC Type 4”.
- b. Add the db2jcc.jar and db2jcc_license_cu.jar files to the WebLogic classpath. Edit the file startWebLogic.cmd or startWebLogic.sh for the WebLogic domain you created. For example,

```
set CLASSPATH=%CLASSPATH%;c:\db2\jdbc\db2jcc.jar;c:\db2\jdbc\
db2jcc_license_cu.jar
```
- c. Stop and then start WebLogic Server.

Microsoft SQL Server

- a. Download and unzip Microsoft SQL Server Driver 2005 JDBC Driver, sqljdbc.jar, from Microsoft Support to a directory *jdbc_path* on your application server machine, such as:

UNIX

/opt/jars

Windows

C:\jars

- b. Perform one of the following steps, depending on your operating system type.

Windows

Edit the file startWebLogic.cmd (by default, in the directory C:\bea\user_projects\domains\bin\FNCEDomain) for the WebLogic domain you created. Insert the following two lines immediately after the first occurrence of the line `CLASSPATH=...`

```
set JDBC_PATH=jdbc_path\sqljdbc_1.0\enu\sqljdbc.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%
```

AIX

Add the following line to the file setDomainEnv.sh file:

```
JAVA_OPTIONS="$JAVA_OPTIONS -
Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0"
```

UNIX

Edit the file startWebLogic.sh by inserting the following two lines immediately after the first occurrence of the line `CLASSPATH=...`

```
JDBC_PATH=jdbc_path/sqljdbc_1.0/enu/sqljdbc.jar
CLASSPATH=$JDBC_PATH:$CLASSPATH
```

- c. Stop and then start WebLogic Server.

Oracle

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

NOTE If you intend to install AddOns (extensions to IBM FileNet P8 core components), and your Content Engine database will be Oracle, your Oracle JDBC Driver file requirements may be more restrictive. For the required version and patch number, see the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

- c. From the Oracle web site, apply the patch Oracle Patch `Ojdbc##.jar`.
- d. Edit the file startWebLogic.cmd or startWebLogic.sh for the WebLogic domain you created. Add the following line immediately after the first line that starts with `set CLASSPATH`.

Windows

```
set JDBC_PATH=<jdbc_path>\ojdbc##.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%;
```

UNIX

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

- e. Stop and then start WebLogic Server.

5. Give the Configuration Manager user the following permissions:
 - Read, write, and execute permission on the domain directory `../users_projects/domains/your_domain`.
 - Read and execute permission on the `../common/bin` directory.

Configure JBoss for Content Engine

NOTE This task assumes you have already installed JBoss Application Server on the machine where you are going to deploy Content Engine.

To configure JBoss for Content Engine

1. Navigate to the JBoss directory `JBOSS_DIST/server`, which contains configuration file sets.

NOTE If you are installing into a JBoss cluster, use the JBoss directory `JBOSS_DIST/all`.
2. Create a new configuration file set by copying the default configuration file set to a new directory (called `server1` in this procedure) within the `/server` or `/all` directory.
3. Edit the `run.conf` configuration file, located at `JBOSS_DIST/bin`, as follows:
 - a. Add a line to specify the path to the JDK on the machine where JBoss is installed, as shown in the following example:

```
JAVA_HOME="<path_to_Java_JDK>"
```
 - b. In the `JAVA_OPTS` line, change the `-Xms` and `-Xmx` values from

```
-Xms128m -Xmx512m
```

to

```
-Xms512m -Xmx1024m
```
 - c. Save your edits.
4. Open the file `login-config.xml` for editing. This file is typically located at `.../server/myserver/conf`, where *myserver* is the name of the JBoss server instance.
 - a. In the `<!DOCTYPE` declaration, change

```
"http://www.jboss.org/j2ee/dtd/security_config.dtd"
```

to

```
"<jboss_install_dir>/docs/dtd/security_config.dtd"
```

where `<jboss_install_dir>` is the directory where JBoss is installed.
 - b. Save your edit.
5. 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 GCD or for an object store you will be creating later. To download this guide from the IBM support page, see ["Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19](#).

6. Install JDBC drivers on the JBoss machine, as follows:

- a. Obtain the JDBC drivers, depending on your database type.

DB2

Find the latest version of the Redistributable DB2 JDBC Driver Type 4 driver from the [IBM web site](http://www.ibm.com) (<http://www.ibm.com>) by searching for “JDBC Type 4”.

Microsoft SQL Server

Find the Microsoft SQL Server Driver 2005 JDBC Driver, sqljdbc.jar, at Microsoft Support.

Oracle

Access the [Oracle JDBC Driver Downloads](http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) web site (http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html) and find the JDBC driver file that matches the version of the JDK on the JBoss machine.

- b. Place the file JDBC driver file from [Step a](#) in CLASSPATH by copying it to the directory JBOSS_DIST/server/server_name/lib.

NOTE If you are installing into a JBoss cluster, use the JBoss directory JBOSS_DIST/all/server_name/lib.

7. Increase the database transaction timeout. Edit the jboss-service.xml file, and set TransactionTimeout to at least 600, as in this example:

```
<mbean code="org.jboss.tm.TransactionManagerService"
  name="jboss:service=TransactionManager"
  xmbean-dd="resource:xdmdesc/TransactionManagerService-xmbean.xml">
  <attribute name="TransactionTimeout">600</attribute>
```

8. If you are deploying multiple instances of Content Engine of the same server, do the following for each additional instance:

- a. Copy the configuration file set that you just created and modified in [Step 2](#) through [Step 6](#) from the /server/server1 directory to a new directory. Use a separate directory for each instance.
- b. Assign unique port numbers to each instance. Refer to your JBoss documentation for details.

9. If it isn't already running, start JBoss as follows, and leave the command window open:

UNIX

```
./run.sh -c server1
```

Windows

```
run.bat -c server1
```

Assign directory permissions

Give the Configuration Manager user read and write permission on the server directory where the Content Engine instance will be installed.

Configure JBoss server clusters

JBoss servers can be grouped together into a cluster for performance or to provide high availability. This guide provides on minimal instructions for setting up a JBoss cluster. Refer to the *IBM FileNet P8 Platform High Availability Technical Notice* for details on how to set up your IBM FileNet P8 system using clusters, farms, and other high availability software and hardware. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Because JBoss clusters do not have an administrative server, you will choose a single JBoss server on which to install and configure the first instance of Content Engine, then copy the necessary files to the rest of the nodes in the cluster. See [“Deploy Content Engine to additional JBoss servers in a cluster” on page 50](#) for details of which files to copy to other nodes.

FileNet P8 Administrator upgrade tasks

As the FileNet P8 Administrator, perform the following tasks to prepare your environment for IBM FileNet P8 Platform upgrade:

- Review all rows assigned to the FileNet P8 Administrator in the [“Installation and upgrade worksheet” on page 216](#). While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. For descriptions of the properties in the Installation and Upgrade Worksheet, see [“Installation and upgrade glossary” on page 228](#).

Appendixes

This appendix section contains the following major topics:

- [“Installation and upgrade worksheet” on page 216](#)
- [“Installation and upgrade glossary” on page 228](#)
- [“Installing P8 Platform in a non-English environment” on page 248](#)
- [“Process Engine SQL scripts” on page 218](#)
- [“IBM FileNet P8 ports” on page 265](#)

Installation and upgrade worksheet

The *Installation and Upgrade Worksheet* is a spreadsheet that contains properties, parameters, and field names for installation, upgrade, and initial configuration of FileNet P8 Platform components and add-on products. Administrators who are preparing the environment for installation or upgrade of IBM FileNet P8 Platform can use the worksheet during their preparation tasks to record the appropriate values and provide them to the Installation Administrator who will run the installation or upgrade.

You can access the worksheet through the following link:

[Installation and Upgrade Worksheet](#)

or on the IBM product documentation page. To download this and other IBM FileNet product documentation from the IBM web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Worksheet organization

Review the following considerations about the organization of the worksheet:

- 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 the rows usefully.
- For descriptions of the non-obvious properties and parameters in the worksheet, see the [“Installation and upgrade glossary” on page 228](#).
- The “Role” column uses the following acronyms for Administrator roles:
 - ITA: Information Technology Administrator
 - ASA: Application Server Administrator
 - DBA: Database Administrator
 - SA: Security Administrator
 - P8A: FileNet P8 Administrator

Optimizing the worksheet

The worksheet provides all possible properties, parameters, and field names for installation, upgrade, and initial configuration of P8 Platform components as well as add-on products. Only a small subset of these rows will apply to a particular installation or administrator role as part of that installation. For example, a database administrator who is only working with DB2 can receive a smaller spreadsheet with rows that have been sorted for the Admin Role of DBA, and the independent software vendor of DB2.

You may receive a simplified or more focused variation of this spreadsheet that is appropriate to your environment or role. Work with your install team and your IBM services representative to make the best use of this tool.

HINT 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 installation properties you must specify for a given installer or configuration program:

- Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select the program you are interested in (for example, PE installer).
- Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

Process Engine SQL scripts

Process Engine installation requires that several SQL scripts be run for Oracle and SQL Server databases for new installations. Also, one SQL script for Oracle databases is needed for upgrades. The scripts can be run in the following ways:

- Run the scripts manually before you start the Process Engine installation program
- Let the installation program prompt you for the Oracle SYS or SQL server administrator (sa) password and then run the scripts automatically
- Run the scripts silently using operating system authentication

Use operating system authentication only in a trusted environment or when Process Engine is configured with a local database.

If the scripts are executed from the Process Engine installation program, the ability to connect to the database is validated during installation. If the connection fails, the user can correct the errors and proceed with the installation. If the scripts are executed manually, the database connections are not validated until the end of the installation.

If the scripts are executed manually for a SQL Server database, the SQL Server Client software does not have to be installed on the Process Engine server.

If the scripts are copied to the database server and executed manually by the DBA, the following conditions are true:

- There is no need to provide access to the database / sysadmin password for the person who executes the Process Engine installation program.
- The default run-time and maintenance users and their passwords can be modified.

CAUTION The run-time and maintenance user names and passwords must be entered during Process Engine installation and the names must match those defined here.

SQL Scripts for SQL Server

The following scripts are for new installations only.

Script Name	Action
CreatePEinstallSP_1.sql	Creates PE_createDbUsers stored procedure in the Process Engine database.
CreatePEinstallSP_2.sql	Creates a stored procedure that is called during Process Engine installs.
CreatePEinstallSP_3.sql	Creates fn_error stored procedure.

SQL Scripts for Oracle

The **pe_upgrade_scripts.sql** is the only script needed for upgrades. All other scripts are needed for new installations only.

Script Name	Action
pe_install_scripts.sql	<p>A wrapper script that executes the following scripts: pe_filenet_site.sql, pe_create_stored_procedures.sql, and pe_grant_sp_permissions.sql.</p> <p>The pe_install_scripts.sql script generates an Oracle output spool file. After the script completes, the Process Engine installation program scans the output file for errors.</p>
pe_filenet_site.sql	<p>Creates the database users for IBM FileNet P8. See “Changes Made by pe_filenet_site.sql and pe_oracle_users_defaults.sql” on page 220 for details.</p>
pe_create_stored_procedures.sql	<p>Creates several stored procedures. See “Changes Made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql” on page 221 for details.</p>
pe_grant_sp_permissions.sql	<p>Executes grants and creates synonyms for stored procedures. See “Changes Made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql” on page 221 for details.</p>
pe_oracle_users_defaults.sql	<p>Sets user tablespace defaults and privileges for IBM FileNet P8 users. See “Changes Made by pe_filenet_site.sql and pe_oracle_users_defaults.sql” on page 220 for details.</p>
pe_upgrade_scripts.sql	<p>A wrapper script that executes the following scripts: pe_create_stored_procedures.sql pe_grant_sp_permissions.sql.</p> <p>The pe_upgrade_scripts.sql script generates an Oracle output spool file. After the script completes, the Process Engine installation program scans the output file for errors.</p>

All the scripts are located in the root of the Process Engine installation directory. Process Engine installation cannot complete until all these scripts have run successfully. If an error message indicates that any of the scripts did not run, you must resolve the errors before you proceed. See the *IBM FileNet P8 Platform Troubleshooting Guide* for more information. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs”](#) on page 19.

Changes Made by pe_filenet_site.sql and pe_oracle_users_defaults.sql

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

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

Passwords can be manually changed after installation with the Xdbconnect tool, as described in each “Install Process Engine” and “Upgrade Process Engine” topics in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

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

Changes Made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql

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

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

Oracle and SQL Server scripts run with the following options:

run_time_user

This option value is one of the following:

f_sw - if you use the default database users when you run the Process Engine installation program.

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

maintenance_user

This option value is one of the following:

f_maint - if you use the default database users when you run the Process Engine installation program.

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

data_tablespace (Oracle only)

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

index_tablespace (Oracle only)

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

temp_tablespace (Oracle only)

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

PE_database_name (SQL Server only)

The database name that will be entered during Process Engine installation.

DSN (SQL Server only)

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

To edit SQL scripts and change default passwords (SQL Server)

1. Open the following script file with a text editor:

CreatePEinstallSP_1.sql

2. Locate and modify the following line to change the default password for the run_time_user, where 'filenet' is the default password. Passwords must be enclosed in single quotes.

```
set @passwd1 = 'filenet'
```

3. Locate and modify the following line to change the default password for the maintenance user, where 'change\$this_obnoxious_passwd' is the default password.

```
set @passwd2 = 'change$this_obnoxious_passwd'
```

To run SQL scripts manually for a new installation (SQL Server)

1. Copy CreatePEinstallSP_1.sql, CreatePEinstallSP_2.sql, and CreatePEinstallSP_3.sql from the Process Engine software package to the local disk on the SQL Server database server.
2. Run CreatePEinstallSP_3.sql first, then CreatePEinstallSP_2.sql, and then CreatePEinstallSP_1.sql.

```
osql -E -D DSN -d PE_database_name -i CreatePEinstallSP_3.sql -n -o output3.log
```

```
osql -E -D DSN -d master -i CreatePEinstallSP_2.sql -n -o output2.log
```

```
osql -E -D DSN -d PE_database_name -i CreatePEinstallSP_1.sql -n -o output1.log
```

Instead of -E you can use the following:

```
-U sa -P sa password
```

Successful execution will record nothing in the output3.log and output2.log files. On a SQL Server 2000 configuration, the output1.log file may contain the following error, which can be ignored:

Cannot add rows to sysdepends for the current stored procedure because it depends on the missing object 'sys.sp_validname'. The stored procedure will still be created.

3. Create a text file named input.txt to identify the run-time and maintenance users. The file should contain the following:

```
PE_createDbUsers 'run_time_user', 'maintenance_user', 'PE_database_name'
```

For example:

```
PE_createDbUsers 'f_sw', 'f_maint', 'VWdb'
```

where:

run_time_user is the default *f_sw* user account

maintenance_user is the default *f_maint* user account

PE_database_name is VWdb

In this example:

- the *run_time_user* is the default *f_sw*
- the *maintenance_user* is the default *f_maint*
- the database is VWdb

NOTE If names other than these defaults are used, these alias names must match the names provided to the Process Engine installation program.

4. Execute the following procedure:

```
osql -E -D DSN -d simDB -i input.txt
```

Instead of -E you can use the following:

```
-U sa -P sa password
```

To edit SQL scripts and change default passwords (Oracle)

1. Open the following script file with a text editor:

```
pe_filenet_site.sql
```

2. Locate and edit the default run-time password (filenet) in the following command:

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

3. Locate and edit the default maintenance password (change\$this_obnoxious_passwrđ) in the following command:

```
grant dba to &2 identified by change$this_obnoxious_passwrđ;
```

4. Open the following script file with a text editor:

`pe_install_scripts.sql`

5. Locate and edit the default run-time password (filenet) in the following command:

`connect &1/filenet`

To run SQL scripts manually for a new installation (Oracle)

1. From the Process Engine software package, copy the scripts to the database server.
2. Start SQL Plus. For example, type the following command:

`sqlplus "sys/password as sysdba"`

3. At the SQL prompt, enter:

`@pe_install_scripts.sql run_time_user maintenance_user data_tablespace
index_tablespace temp_tablespace`

For example:

`@pe_install_scripts.sql f_sw f_maint vwdata_ts vwindex_ts vwtemp_ts`

To run SQL scripts manually for an upgrade (Oracle)

1. From the Process Engine software package, copy the scripts to the database server.
2. Start SQL Plus. For example, type the following command:

`sqlplus "sys/password as sysdba"`

3. At the SQL prompt, enter:

`@pe_upgrade_scripts.sql run_time_user maintenance_user`

For example:

`@pe_upgrade_scripts.sql f_sw f_maint`

NOTE These scripts that are run via `pe_upgrade_scripts.sql` must be executed for all upgrades.

Use Cases for Running Scripts and Setting Passwords

Following are a number of use cases to describe Process Engine installation and configuration variables and how to set them in each case.

Case 1: With Oracle or SQL Server databases, the DBA wants to run SQL scripts manually before running the Process Engine installer. All default users and passwords will be used.

1. (Oracle only) Turn off Oracle password complexity.
2. Run the scripts without making any changes to passwords, setting the `run_time_user` to `f_sw` and the `maintenance_user` to `f_maint`.
3. Run Process Engine installer and use default users (do not set aliases).

4. Leave the f_sw and f_maint password fields blank (using the defaults assigned with the scripts ran).
5. (Oracle only) Turn Oracle password complexity back on.
6. Reset the f_sw and f_maint passwords by running Xdbconnect. This changes both the encrypted version of the password and the password in the database. Xdbconnect works only if the passwords in the encrypted file and the database match after installation.

Case 2: With Oracle or SQL Server databases, the DBA wants to run SQL scripts manually before running the Process Engine installer but non-default users are set. Default passwords for these users are used.

1. (Oracle only) Turn off Oracle password complexity.
2. Run the scripts without making any changes to passwords, setting the run_time_user and maintenance_users to user names defined by the customer.
3. Run Process Engine installer and indicate that aliases will be configured. The user names set when the scripts ran must be indicated as the aliases for f_sw and f_maint. None of the fields for alias names can be left blank in the Process Engine installer screen, but default user names can be entered.
4. Leave the f_sw and f_maint password fields blank (using the defaults assigned with the scripts ran).
5. (Oracle only) Turn Oracle password complexity back on.
6. Reset the f_sw and f_maint passwords by running Xdbconnect. This changes both the encrypted version of the password and the password in the database. Xdbconnect works only if the passwords in the encrypted file and the database match after installation.

Case 3: The DBA wants to run SQL scripts for SQL Server or Oracle manually before running the Process Engine installer, but non-defaults users and passwords are set.

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

Case 4: SQL Server or Oracle scripts are executed from the Process Engine installation program. Either the default f_sw and f_maint users or their aliases are selected. Default passwords for these users will be used.

1. (Oracle only) Turn off Oracle password complexity.
2. Run the Process Engine installation program and indicate either default or alias users.
3. In the Process Engine installer screen that prompts for f_sw and f_maint passwords, leave the fields blank.
4. (Oracle only) Turn Oracle password complexity back on.
5. Reset the f_sw and f_maint passwords by running Xdbconnect. This changes both the encrypted version of the password and the password in the database. Xdbconnect works only if the passwords in the encrypted file and the database match after installation.

Case 5: With a DB2 database, the f_sw and f_maint users or their aliases must already exist as operating system users. There are no P8 default passwords.

1. Run the Process Engine installation program and indicate that either the default users or aliases will be configured. None of the fields for alias names can be left blank in the Process Engine installer screen, but default user names can be entered.
2. During the Process Engine installation, set the f_sw and f_maint passwords to the match the passwords set when the users were created.
3. You need not change passwords with Xdbconnect.

The following table summarizes the use cases and required actions.

	Case 1	Case 2	Case 3	Case 4	Case 5
Oracle Password Complexity turned off before running the Process Engine installation program	x	x		x	n/a
Default users (f_sw and f_maint) set in SQL scripts	x				n/a
Aliases defined during Process Engine installation		x	x	either defaults or aliases assigned	x (unless the default user names were assigned to the operating system users already defined)
Default passwords set in SQL scripts	x	x		x	n/a
f_sw and f_maint password field values in Process Engine installation	leave blank	leave blank	set to values in scripts	leave blank	set to value defined for operating system user
Passwords reset for f_sw and f_maint after Process Engine installation	x	x		x	

Installation and upgrade glossary

This glossary contains the descriptions of the major properties and parameters found in the various IBM FileNet P8 software installers and configuration programs (such as Configuration Manager). You can use this glossary in conjunction with the *Installation and Upgrade Worksheet* to understand what values and settings you must supply when installing, upgrading, and configuring your IBM FileNet P8 systems.

Property or Parameter (In user interface, XML file, or script)	Description
(Remote database) Location of shared Temp directory	A shared Temp directory on the PA server with full access rights to the pa_administrator.
.NET API COM Compatibility Layer (CCL) Server URL	If you use custom COM applications that require the CCL, provide a URL. For example, http://localhost:9080/wsi/FNCEWS40MTOM/ .
Accept User Tokens?	Bootstrap preference for Application Engine or Workplace XT. This setting enables user token authentication.
Administrative console user name	A user account that has access to log on to the application server administration console. If your site uses a FederatedLDAP registry, this account name must be a unique user across all federated realms.
Administrative User and Group aliasing method	Indicate whether default Process Engine administrative user and group names will be used or aliases will be created for these names. If aliases will be configured, actual names will be prompted for from the Process Engine installation program. NOTE For a DB2 for z/OS database, you must choose to assign aliases.
Analysis Services Database Name	The name of the Analysis Service database.
Analysis Services Instance Name	The instance name of the Analysis Service database.
Analysis Services Instance Name	The Analysis Services Instance Name. Only used when installing on SQL Server 2005 with a non-default Analysis Service.
Analysis Services Server Name	The server name of the Analysis Service database.
Application Integration Settings - Prompt to add email?	Determines if a user is prompted to add e-mail to an object store when the user sends e-mail from Outlook if the user has Application Integration installed.

Property or Parameter (In user interface, XML file, or script)	Description
Application server administrator password	The password for the application server administrative user account. The Configuration Manager GUI encrypts the password as you type. If you edit the configuration file manually, you can enter an encrypted password or a plain-text password. See “Encrypt passwords” on page 453 for the procedure to encrypt a password before pasting the password into a file.
Application server administrator user name	The directory service user account that has been assigned the application server administrative role. This account is used to manage the application server domain or profile, to configure the data sources and connection pools for the GCD, and to deploy the Content Engine application. If Administrative Security is already enabled, use an existing administrative user account.
Application server cell	The name of the WebSphere cell where Content Engine is or will be deployed.
Application server domain name	The name of the Weblogic domain where Content Engine is or will be deployed.
Application server host	The machine name or the IP address of the local host.
Application server name	The name of the WebSphere JVM server where Content Engine is or will be deployed.
Application server node	The name of the WebSphere node where Content Engine is or will be deployed.
Banner Image: Image width & height	Bootstrap preference for Application Engine. Specifies size of the banner image to display in the application window.
Banner Image: path to file	Bootstrap preference for Application Engine. Specifies a path to the banner image file to display in the application window.
Bootstrap user name	The name of a directory server user that accesses the Global Configuration Data (GCD) data sources. Use only the short name of the bind user defined by the LDAP user attribute. For example, administrator.

Property or Parameter (In user interface, XML file, or script)	Description
Bootstrap user password	<p>The password for the directory server user that accesses the Global Configuration Data (GCD) data sources.</p> <p>The Configuration Manager GUI encrypts the password as you type. If you edit the configuration file manually, you can enter an encrypted password or a plain-text password. See “Encrypt passwords” on page 453 for the procedure to encrypt a password before pasting the password into a file.</p>
Bootstrapped EAR directory	<p>The name of a subdirectory that will store the EAR file that contains the Content Engine bootstrap information. The bootstrap information is needed for creating the Global Configuration Data (GCD) and for starting Content Engine. Specify the directory relative to the ce_install_path/lib directory. For example, to specify /opt/FileNet/ContentEngine/lib/bootstrap, set the value to bootstrap.</p>
Bootstrapped EAR path	<p>The fully qualified path to the bootstrapped Content Engine EAR file that was created by the configure bootstrap task. For example for WebSphere use /opt/FileNet/ContentEngine/lib/bootstrap/Engine-ws.ear or c:\Program Files\FileNet\ContentEngine\lib\bootstrap\Engine-ws.ear.</p>
CE Application Server	The application server used by CE.
CE Application Server version	The application server version used by CE.
CE Service Password	The password for the CE service user.
CE Service Username	The CE service username.
Choose authentication method	<p>For Application Engine installation, select the authentication method for use at your site.</p> <p>Application-Managed Authentication uses authentication specific to the application and does not share credentials.</p> <p>Container-Managed Authentication provides the ability to use single sign-on (SSO) capabilities to share credentials between Application Engine and custom applications.</p> <p>When you select Container-Managed Authentication, the installer installs a sample log-in application, and modifies the web.xml file to support SSO. You will need to perform additional configuration for SSO after the installer is finished.</p>
Common files directory	Directory for configuration files that will be shared with FileNet P8 components other than PA or PS.

Property or Parameter (In user interface, XML file, or script)	Description
Company name	Customer Information: The name of your company.
Content Engine API - Transport Method	The transport method for the Content Engine to which PA will connect. Select WSI
Content Engine application name	The Content Engine application name as it appears or will appear in the application server (for example, in an administration console). The application name is subject to application server naming constraints. For WebSphere, each application in a cell must have unique name.
Content Engine Client Software URI	The Content Engine server name and HTTP port number that Process Analyzer will connect to.
Content Engine Client software URL	The URL for the Content Engine Web Services client API. This URL will contain the WcmApiConfig.properties file, which is required for applications to communicate with Content Engine Server, regardless of whether they use the EJB or Web Services transport method.
Content Engine Download URL	The download URL is used for internal processes during EJB transport activities.
Content Engine EAR path	The fully qualified path to the Content Engine EAR file that was installed by the Content Engine setup program. For example, /opt/FileNet/ContentEngine/lib/Engine-ws.ear or c:\Program Files\FileNet\ContentEngine\lib\Engine-ws.ear, where Engine-ws.ear is the WebSphere EAR file name.
Content Engine Upload URL	The upload URL is used for internal processes during EJB transport activities.
Content Engine URI	The Content Engine server name and HTTP port number that Process Analyzer will connect to.

Property or Parameter (In user interface, XML file, or script)	Description
Content Engine URL	<p>The Content Engine application server to which Process Engine or Application Engine will connect.</p> <p>Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example:</p> <p>WebSphere http://hqcemp2:9080/wsi/FNCEWS40DIME</p> <p>WebLogic http://hqcemp1:7001/wsi/FNCEWS40DIME</p> <p>JBoss http://hqcemp3:8080/wsi/FNCEWS40DIME</p> <p>Do not modify the remainder of the string from the default values.</p>
Create maximum strength keys	For Application Engine or Workplace XT installations, specifies whether you choose to have the maximum encryption for user token authentication.
Create war file during install	During the installation of the ISRA servlet, you can decide whether to create a war file for deployment on WebSphere application server.
Database class	<p>The name of the JDBC driver class that is used to establish a connection with the database.</p> <p>For DB2, use com.ibm.db2.jcc.DB2Driver.</p> <p>For Microsoft SQL Server, use com.microsoft.sqlserver.jdbc.SQLServerDriver.</p> <p>For Oracle, use oracle.jdbc.OracleDriver.</p>
Database host	The PA Database server.
Database instance name	<p>Content Engine</p> <p>The name of the DB2 instance for z/OS for the Global Configuration Data (GCD) or an object store database.</p> <p>Process Analyzer</p> <p>The name of the database instance to be used by PA.</p>

Property or Parameter (In user interface, XML file, or script)	Description
Database location	Indication of whether the PA database will be local or remote.
Database name	<p>Content Engine</p> <p>The name of the Content Engine Global Configuration Data (GCD) or an object store database.</p> <p>For DB2 this value is the database name.</p> <p>For Oracle, this value is the SID of the database containing the GCD or object store tablespaces.</p> <p>For SQL Server, this value is the name of the GCD or an object store database.</p> <p>Process Analyzer</p> <p>The name of the PA database you want to create.</p>
Database Name for Process Analyzer Analysis Services	The name of the Analysis Service database.
Database password	<p>For DB2 for Linux, UNIX, and Windows, this is the password for the operating system user of the Global Configuration Data (GCD) or object store tablespace.</p> <p>For DB2 for z/OS, this is the password for the operating system user associated with the Global Configuration Data (GCD) or object store.</p> <p>For Oracle this is the password for the GCD or object store tablespace owner.</p> <p>For SQL Server, this is the password for the user with administrative rights to the GCD or object store database.</p> <p>The Configuration Manager GUI encrypts the password as you type. If you edit the configuration file manually, you can enter an encrypted password or a plain-text password. See “Encrypt passwords” on page 453 for the procedure to encrypt a password before pasting the password into a file.</p>

Property or Parameter (In user interface, XML file, or script)	Description
Database port number	<p>Content Engine</p> <p>The port number used by the database instance in which you create (SQL Server) databases or (Oracle or DB2) tablespaces for the Global Configuration Data (GCD) and object stores.</p> <p>Process Analyzer</p> <p>The database port number for the PA database.</p> <p>Process Simulator</p> <p>The database port number for the PA database.</p>
Database server name	The host name of the machine where the database software is installed.
Database service name	The SID of the Oracle database containing the GCD or object store tablespaces.
Database type	The database vendor type. DB2, SQL Server, and Oracle are supported.
Database user name	<p>Content Engine</p> <p>For DB2, this is the name of the Global Configuration Data (GCD) or object store tablespace user.</p> <p>For Oracle this is the GCD or object store tablespace owner.</p> <p>For SQL Server, this is the name of the user with administrative rights to the GCD or object store database.</p> <p>Process Simulator</p> <p>The PA Database administrator user (<i>pa_db_administrator</i>).</p>
Date/Time Mask	The date/time mask to be entered in Process Task Manager configuration. Default value is mm/dd/yyyy hh:tt:ss.
DB2 blob tablespace name	The default blob tablespace name for Process Engine. See “Prepare DB2 Server for Linux, UNIX and Windows” on page 113 for more information.
DB2 data tablespace name	The default data tablespace name for Process Engine. See “Prepare DB2 Server for Linux, UNIX and Windows” on page 113 for more information.

Property or Parameter (In user interface, XML file, or script)	Description
DB2 database alias	The database alias for the Process Engine database. See “Prepare DB2 Server for Linux, UNIX and Windows” on page 113 for more information.
DB2 database name	The name of the DB2 database identified for Process Engine use. See “Prepare DB2 Server for z/OS” on page 109 for more information.
DB2 index tablespace name	The default index tablespace name for Process Engine. See “Prepare DB2 Server for Linux, UNIX and Windows” on page 113 for more information.
DB2 instance alias name	The name of the instance alias for the DB2 for z/OS instance identified for Process Engine use. This is the alias created on the Process Engine, which is the DB2 client. See “Configure DB2 Client” on page 121 for additional information.
DB2 local instance owner's name	The DB2 Client instance owner's name. This is the instance on the Process Engine server. This is applicable to both DB2 for Linux, UNIX, and Windows and DB2 for z/OS and applies to Process Engine on UNIX platforms. See “Configure DB2 Client” on page 121 for more information.
Default Storage group (z/OS)	The default storage group to be used when DB2 issues index creation SQL statements.
Deployment type	The Content Engine application deployment type. For WebSphere, the choices are Standard, Cluster, and Network Deployment. For JBoss and WebLogic, the choices are Standard and Cluster.
Destination folder	The folder where the installation program creates a Lipient directory and installs the Records Manager software.
Directory name	<p>Process Analyzer</p> <p>The destination drive and directory for Process Analyzer Engine or Client software.</p> <p>Process Simulator</p> <p>The installation location for the PS executable files where the \FNSW directory will be created.</p>
Directory Name (for ISRA Servlet)	The directory where you plan to install the ISRA servlet.

Property or Parameter (In user interface, XML file, or script)	Description
Directory service bind user name	The fully qualified distinguished name of the LDAP bind user for authenticating to the directory server.
Directory service bind user password	<p>The password used by the LDAP bind user to authenticate to the directory server.</p> <p>The Configuration Manager GUI encrypts the password as you type. If you edit the configuration file manually, you can enter an encrypted password or a plain-text password. See “Encrypt passwords” on page 453 for the procedure to encrypt a password before pasting the password into a file.</p>
Directory service port number	The port number configured on the directory server host for communicating with the directory server.
Directory service server host name	The directory server host name. The short or full name or IP address of the system that hosts the directory service.
Do not use SSL certificates for server communication	<p>This setting specifies whether SSL certificates will be used for server communication. Select this check box to turn off SSL certificate use.</p> <p>NOTE Selecting this check box will change your WebSphere settings for communicating with other servers, such as Application Engine. If you have already configured WebSphere to use certificates, clear this check box. The default is to turn off certificate use.</p> <p>If you select the default, Do not use certificates for server communication, then four WebSphere settings are changed, as follows.</p> <p>CSlv2 inbound authentication > Client certificates authentication = “Never”</p> <p>CSlv2 outbound authentication > Client certificates authentication = “Never”</p> <p>CSlv2 inbound transport > Transport = “TCP/IP”</p> <p>CSlv2 outbound transport > Transport = “TCP/IP”</p> <p>Refer to your WebSphere Application Server documentation for more information.</p>
Documentation server URL	The URL for the application server where the IBM FileNet P8 documentation is deployed.

Property or Parameter (In user interface, XML file, or script)	Description
Download directory path	In Application Engine or Workplace XT installations, specifies the directory used by Workplace or Workplace XT to store temporary copies of files downloaded from the application.
Execution mode for SQL scripts	<p>Several SQL scripts must be executed for Process Engine SQL Server and Oracle installations and for upgrades with Oracle databases.</p> <p>Indicate how the SQL scripts should be executed. The choices are:</p> <ul style="list-style-type: none"> • I have already run the pre-install scripts manually. • I want to run the scripts in an xterm window. (UNIX) • I want to run the scripts silently using operating system authentication. • I want to run the scripts with a prompted password. (Windows) <p>Execution from an xterm window or by prompting for a password will require entry of the SQL Server sa password or the Oracle SYS password. See “Process Engine SQL scripts” on page 218 for more information on execution modes for SQL scripts.</p>
f_maint password	At the Process Engine installation program prompt, leave this blank to use the default password "change\$this_obnoxious_passwr" or enter a new password. In the PE_silent_install.txt file, enter the encrypted version of the password. For Oracle and SQL Server databases, the f_maint user is a database user and this password must match the password used in the SQL scripts. For DB2 databases, f_maint is an operating system user and there are no SQL scripts. See “Accounts for Process Engine (Windows)” on page 79 or “Accounts for Process Engine (UNIX)” on page 80 for more information.
f_maint user name	The alias to be assigned for the Process Engine maintenance user for Process Engine. See For Oracle and SQL Server databases this is a database user. For all DB2 databases this is an operating system user. For DB2 for z/OS, underscore characters are not valid. See “Accounts for Process Engine (Windows)” on page 79 or “Accounts for Process Engine (UNIX)” on page 80 for more information.

Property or Parameter (In user interface, XML file, or script)	Description
f_sw password	At the Process Engine installation program prompt, leave this blank to use the default password "filenet" or enter a new password. In the PE_silent_install.txt file, enter the encrypted version of the password. For Oracle and SQL Server databases, the f_sw user is a database user and this password must match the password used in the SQL scripts. For DB2 databases, f_sw is an operating system user and there are no SQL scripts. See “Accounts for Process Engine (Windows)” on page 79 or “Accounts for Process Engine (UNIX)” on page 80 for more information.
f_sw user name	The alias to be assigned for the Process Engine runtime user. For Oracle and SQL Server databases this is a database user. For all DB2 databases this is an operating system user. For DB2 for z/OS, underscore characters are not valid. See “Accounts for Process Engine (Windows)” on page 79 or “Accounts for Process Engine (UNIX)” on page 80 for more information.
Federated Repository virtual realm name	The name of the Federated Repository virtual realm. This value is case sensitive. The default name is defaultWIMFileBasedRealm.
FNADMIN group name	The alias to be used in place of the default FNADMIN group name for Process Engine. See “Specify IBM FileNet P8 accounts” on page 65 for additional information.
FNOP group name	The alias to be used in place of the default FNOP group name for Process Engine. See “Specify IBM FileNet P8 accounts” on page 65 for additional information.
FNSW user name	The alias to be used in place of the default FNSW user name for Process Engine. See “Specify IBM FileNet P8 accounts” on page 65 for additional information.
FNUSR group name	The alias to be used in place of the default FNUSR group name for Process Engine. See “Specify IBM FileNet P8 accounts” on page 65 for additional information.
Full pathname for the device files for the fn_SEC_DB0 volume	Enter the full pathname for the device files for the fn_SEC_DB0 volume. See “Configure Process Engine Servers (all UNIX)” on page 38 .
Full pathname for the device files for the fn_SEC_RL0 volume	Enter the full pathname for the device files for the fn_SEC_RL0 volume. See “Configure Process Engine Servers (all UNIX)” on page 38 .

Property or Parameter (In user interface, XML file, or script)	Description
Generate User Tokens?	Bootstrap preference for Application Engine or Workplace XT. Determines if user tokens are generated. Default is true.
GlobalJNDIName	XML upgrade status file property. See JDBC XA data source name.
Group attribute	An attribute in a directory server entry that identifies the group.
Group filter	The filter used by the bind user when searching for groups in a directory server.
HTTP port on the SSO proxy host	Specifies a port on the single sign-on (SSO) proxy server to be added to the web.xml file when configuring SSO on the application server for Application Engine or Workplace XT.
Index Server	Content Search Engine service configured on the Autonomy K2 Master Administration Server.
Install location for common files	<p>The installation location for a common set of configuration files that can be shared by Process Engine and other FileNet P8 products.</p> <p>The default location on Windows is C:\Program Files\FileNet\Common Files.</p> <p>The default location on UNIX is /opt/FileNet/CommonFiles</p>
Install location for data files	The installation location for Process Engine data files is configurable, with the default being C:\. This is where the FNSW_LOC directory will reside. This must be a shared drive in a cluster configuration.
Install location for program files	The installation location for Process Engine program files is configurable, with the default being C:\. This is where the FNSW directory will reside. This must be a local drive in a cluster configuration.
ISRA Interface Servlet URL	<p>If your site integrates with Image Services, enter the Web address (URL) to the server where the ISRA Interface Servlet application was installed, in the format <code>http://ApplicationEngineISRAServlet servername:non-SSL-port/ApplicationEngineISRAServlet/ISRA</code>.</p> <p>NOTE The ISRA Interface Servlet URL should always use a non-SSL port even if Workplace or Workplace XT is configured for SSL.</p>

Property or Parameter (In user interface, XML file, or script)	Description
Java Server HTTP port	Bootstrap preference for Application Engine or Workplace XT- the Java Server HTTP Port for SSL.
JDBC data source name	The JNDI name of the non-XA JDBC data source associated with the Global Configuration Data (GCD) tablespace or database. In the XML upgrade status file, the property name is LocalJNDIName.
JDBC Driver Classpath	The Microsoft SQL Server JDBC 2005 1.2 driver classpath.
JDBC Driver Name	The Microsoft SQL Server JDBC 2005 1.2 driver name.
JDBC XA data source name	The JNDI name of the XA JDBC data source associated with the Global Configuration Data (GCD) tablespace or database. In the XML upgrade status file, the property name is GlobalJNDIName.
K2 Administrative User	Additional user/s, defined in the Autonomy K2 Dashboard, with authorization to log into the Autonomy K2 Dashboard and apply modifications to the configuration or perform maintenance functions.
K2 Broker	Content Search Engine service configured on the Autonomy K2 Master Administration Server.
K2 Operating System User	Autonomy K2 services run as this user and K2 Administration Servers require this user during installation to allow synchronization with the Master Administration Server. User must also have Administrator privileges on the Master Administration Server.
K2 Security Group	Autonomy K2 group used to secure collections. Used in the Verity Domain Configuration in Enterprise Manager.
K2 Security User	A user account used when logging on to the Autonomy K2 Dashboard.
K2 Server	Content Search Engine service configured on the Autonomy K2 Master Administration Server.
K2 Ticket Server	Content Search Engine service configured on the Autonomy K2 Master Administration Server.
Kerberos support	This setting specifies whether Kerberos authentication is used. When set to true, Kerberos authentication is used. When set to false, Kerberos authentication is not used.
License file	The location of the RE license file: IBMFileNetRenditionEngineLicense.xml.

Property or Parameter (In user interface, XML file, or script)	Description
LocalJNDIName	XML upgrade status file property. See JDBC data source name.
Master key	<p>A word or phrase of any length that is used to encrypt sensitive Global Configuration Data (GCD) entries. The Master key is used to configure the GCD settings for the Content Engine bootstrap settings. Store the Master key in a secure location, as it is not retrievable. You will have to specify it later any time you access the GCD with applications built with non-FileNet P8 APIs.</p> <p>The Configuration Manager GUI encrypts the Master key as you type. If you edit the configuration file manually, you can enter an encrypted Master key or a plain-text Master key. See “Encrypt passwords” on page 453 for the procedure to encrypt a password before pasting the password into a file.</p>
NCH clearinghouse domain name	<p>The NCH domain name for the Process Engine. The value will be required for Process Engine installation as well as entry into the Process Engine server's host file or DNS table.</p> <p>In Process Engine installation, enter <i><domain name>:<organization></i>, where:</p> <ul style="list-style-type: none"> • The maximum length of your <i><domain name></i> entry does not exceed 19 characters. • The maximum length of your <i><organization></i> entry does not exceed 19 characters. • Both your <i><domain name></i> and <i><organization></i> entries contain only alphanumeric characters and underscores. <p>A typical convention is to enter <i><your PE machine name>:<your company name></i>. If the machine name or company name include hyphens, replace them with underscores in your entry.</p> <p>For information on entering the NCH name into the hosts file, see “Configure UNIX” on page 36 or “Configure Microsoft Windows” on page 49.</p>
Nickname	A user-friendly name that allows a P8 administrator to connect to one of many P8 domains. The name can be the same as an existing P8 domain name.

Property or Parameter (In user interface, XML file, or script)	Description
Number of keys	Specifies the number of encryption keys for user token authentication. Generally, security increases with the number of keys used.
Object store display name	<p>The name of the object store displayed in Enterprise Manager. The name can be the same as or different from the associated database (SQL Server) or tablespace (DB2 or Oracle).</p> <p>See the IBM FileNet P8 help topic System Administration > Content Engine Administration > Content Engine Wizard Help > Create Object Store Create an Object Store > Name and describe object store for more details.</p>
Object Store Location for Site Preferences file:	Bootstrap preference for Application Engine and Workplace XT. Specifies the object store where the site preferences file is stored.
ODBC data source name	The ODBC data source name defined for the connection to the Process Engine SQL Server database. See “To create the Process Engine ODBC data source and test the connection” on page 94 for additional information.
OLAP Instance Name	The Analysis Services Instance Name.
Oracle database information: Password	The <i>re_db_user</i> (VISTAUSER) password.
Oracle database information: Service	The Oracle service name to use with Rendition Engine.
Oracle database information: User	The name of the RE Oracle database user (default: VISTAUSER).
Oracle DBA OS group name	Required for a UNIX-based Process Engine. See “To create Oracle accounts for Process Engine (UNIX)” on page 83 for more information.
Oracle default DATA tablespace name	The default Oracle data tablespace to be used by Process Engine. See “Prepare Oracle Server” on page 97 for additional information.
Oracle default INDEX tablespace name	The default Oracle index tablespace to be used by Process Engine. See “Prepare Oracle Server” on page 97 for additional information.
Oracle default TEMP tablespace name	The default Oracle temp tablespace to be used by Process Engine. See “Prepare Oracle Server” on page 97 for additional information.

Property or Parameter (In user interface, XML file, or script)	Description
Oracle global database name	The Oracle Global Database Name (as identified in the tnsnames.ora file). Applies for Process Engine with remote databases only.
Oracle HOME directory	The Oracle Home path you enter for Process Engine installation refers to the local Oracle installation directory. Applies for local and remote databases.
Oracle RAC delay	Amount of delay for Oracle RAC.
Oracle RAC retries	Number of retries for Oracle RAC.
Oracle SID	The Oracle SID entered for Process Engine installation applies for local databases only.
Oracle SYS password	The SYS password, only prompted for if you indicate during Process Engine installation that you want to run SQL scripts from a xterm window (UNIX) or with a prompted password (Windows). See “Process Engine SQL scripts” on page 218 for more information on execution modes for SQL scripts.
Oracle user name	Required for a UNIX-based Process Engine. See “To create Oracle accounts for Process Engine (UNIX)” on page 83 for more information.
Overwrite existing repository	This setting specifies whether to overwrite an existing LDAP repository. When set to true, the existing repository is overwritten. This setting is valid only with the StandaloneLDAP repository type.
PE Database Host	The name of the server where the PE database is installed. In the case of a clustered database server, the virtual server name.
PE Database Name	The name of the PE database.
PE Database Port	The port used by the PE database.
PE Database Type	The type of database that PE is using.
PE Database user for PA (pe_db_user_for_pa)	The P8 user name used to connect to the PE database.
PE Database user for PA (pe_db_user_for_pa)	The P8 user name used to connect to the PE database.
PE Database user for PA (pe_db_user_for_pa)	The P8 user name used to connect to the PE database.

Property or Parameter (In user interface, XML file, or script)	Description
PE JDBC Driver Classpath	The location of the JDBC JAR files used by the JDBC driver.
PE JDBC Driver Name	The name of the JDBC driver used to access the PE database.
Process Analyzer Analysis Server	The name of the Process Analyzer Analysis Server where the PA OLAP database is located.
Process Analyzer Host	The Process Analyzer Engine server.
Process Analyzer Port	The Process Analyzer Engine server port.
Process Engine Communication Port	The communication port used by Process Engine. The value entered in Process Task Manager configuration must match the value specified as the communication port when configuring isolated regions. The default value is 32776.
Process events from specific regions and event logs	PA configuration setting to process events from one or more regions.
PSConsole	Access role that determines which users can access Simulation Console. By default, this access role has no members until you add members.
PSDesigner	Access role that determines which users can access Simulation Designer. By default, this access role has no members until you add members.
Rendition Engine database type	The type of database used by RE: Oracle or SQL Server.
Rendition Engine user information: Domain	The domain your RE server is a member of or the server name of a standalone machine.
Rendition Engine user information: Password	The password for the Rendition Engine Administrator (<i>FNRE_Admin</i>).
Rendition Engine user information: User	The administrative RE user (<i>FNRE_Admin</i>). By default the user account with which you logged in to the RE server.
Scheduling options for the cube processing and pruning event	Scheduling Cube Processing and Prune Events actions on the Process Analyzer (PA) server. You can manually process the cubes or prune events if you want to do so at an unscheduled time.
Script	The Configuration Manager script to be used to accomplish a particular configuration task.

Property or Parameter (In user interface, XML file, or script)	Description
Service username	The value entered in Process Task Manager to identify the user that Process Engine uses when connecting to the Content Engine server. This user must belong to the Process Engine Administrator group.
Set as current active user registry	This setting specifies whether to set the LDAP repository as the active user registry. Valid values are: true / false.
SQL Login user	The name of the RE SQL Server database user (default: <i>SQL_login</i>).
SQL Network Name	The SQL network name, for Process Engine in clustered configurations only.
SQL Server Database Information: Catalog (Database)	The database name to be created for Rendition Engine use.
SQL Server database information: Machine	The machine name where the SQL Server database for RE will be created. If you plan to use a named instance rather than the default instance, you must include the instance name: db server name\instance name
SQL Server database information: Password	The <i>re_db_user</i> (SQL_login) user's password.
SQL Server database information: User	The name of the RE SQL Server database user (default: SQL_login).
SQL Server database name	The SQL Server database name to be used by Process Engine, as defined in “Prepare Microsoft SQL Server” on page 89 .
SQL Server filegroup name	The SQL Server filegroup to be used by Process Engine, as defined in “Prepare Microsoft SQL Server” on page 89 .
SQL Server JDBC 2005 1.2 driver classname	The Microsoft SQL Server JDBC 2005 1.2 driver classpath.
SQL Server JDBC 2005 1.2 driver classpath	The Microsoft SQL Server JDBC 2005 1.2 driver classname.
SQL server password	The SQL Server sa password. This will only be required if the Process Engine SQL scripts are to be executed from the Process Engine installation program. See “Process Engine SQL scripts” on page 218 for more information on execution modes for SQL scripts.

Property or Parameter (In user interface, XML file, or script)	Description
SSO proxy host server name	Specifies single sign-on (SSO) proxy host server name to be added to the web.xml file when configuring SSO authentication on the application server for Application Engine or Workplace XT.
SSO proxy host URL	Specifies the URL to be added to the web.xml file when configuring single sign-on SSO authentication on the application server for Application Engine or Workplace XT.
startup script MEM_ARGS settings	For Application Engine and Workplace XT, specify this value during configuration on the WebLogic server. Consult the WebLogic documentation for an appropriate value for your environment.
Storage group option (z/OS)	The storage group option that Content Engine SQL commands use with USING STOGROUP clauses.
SysAdmin password	The password for the SysAdmin user. This value is set by default when Process Engine is installed. It is recommended that the default password be changed after installation. The SysAdmin password will be needed during Process Engine installations and upgrades in order to run the Xapex and Xdbconnect tools to change passwords for the f_sw and f_maint users.
Task enabled	This setting specifies whether the selected task will be executed so that you can specify the settings for a task without executing the task. When set to true, the task will be executed, and the settings will be applied. When set to false, the task will not be executed.
Temp Directory (UNC)	The path to the shared Temp directory required if Process Simulator uses a Simulation Process Analyzer database that is on a remote server.
Temporary directory	The fully qualified path to a temporary directory to be used by the configuration task. For example, /opt/FileNet/ContentEngine/tools/configurationmanager/tmp or c:\Program Files\FileNet\ContentEngine\tools\configurationmanager\temp.
Token Timeout Interval (1-15 minutes)	Bootstrap preference for Application Engine and Workplace XT. Specifies the integer value in minutes for the timeout interval. The token expires if there is no active connection between the applications after this amount of time.

Property or Parameter (In user interface, XML file, or script)	Description
Transport method (Application Engine and Workplace XT)	The transport method for the Content Engine to which Application Engine or Workplace XT will connect. Select EJB.
Transport method for CE API configuration	The transport method for the Content Engine to which Process Engine will connect. For Process Engine installation, select WSI.
Upload directory path	In Application Engine or Workplace XT installations, specifies the directory used by Workplace or Workplace XT to store temporary copies of files uploaded to the application
User base distinguished name	The fully qualified distinguished name of the container for starting a user search in a directory server. For example, cn=users,dc=mydomain.
User filter	The filter used by the bind user when searching for users in a directory server.
User name	Customer Information: Your name.
User name attribute	The attribute by which a user logs on to the directory server. For example, user:sAMAccountName.
UserProjectsPath	The location of the projects path for Weblogic. For example, c:\bea\user_projects\domains.
Users or groups to add to Application Engine Administrators group	Bootstrap preference for Application Engine and Workplace XT. Specifies which existing users or groups should be Application Engine administrators.
Verity Domain	Domain in which the Autonomy K2 services are running.
Web Services HTTP port	The Content Engine Web Services HTTP port number. The Content Engine Web Service (CEWS) is an industry standards-conformant SOAP interface to the FileNet P8 Content Engine. It allows applications to access most of the functionality available through the Content Engine APIs.
WebSphere LDAP repository type	The LDAP repository type. Valid values are: StandaloneLDAP / FederatedLDAP.

Installing P8 Platform in a non-English environment

In order 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” on page 249](#)
- [“Security Administrator” on page 251](#)
- [“Database Administrator” on page 252](#)
- [“IT Administrator” on page 257](#)
- [“FileNet P8 Administrator” on page 259](#)

Application Server Administrator

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

WebSphere and WebLogic

FileNet P8 requires the following character encoding settings:

To configure character encoding on WebSphere

1. Navigate to **Servers > Application Servers > server node**.
2. In the Container Settings areas, navigate to **Container Services > ORB Service > Custom Properties** and click **New**.
3. Enter `com.ibm.CORBA.ORBCharEncoding` in the Name field and `0x05010001` in the Value field and click **Apply**.
4. Save your configuration changes.
5. Navigate to **Servers > Application Servers > server node**.
6. Under Server Infrastructure, navigate to **Java and Process Management > Process Definition**.
7. Under Additional Properties, navigate to **Java Virtual Machine > Custom Properties** and click **New**.
8. Enter `com.ibm.websphere.security.BasicAuthEncoding` into Name field, and `UTF8` into the Value field and click **Apply**.
9. Save your configuration changes.

To configure character encoding on WebLogic

1. Navigate to **my_domain > Servers > my_servers**.
2. Select the Protocols tab, and then select the IIOP tab.
3. Under the Advanced Options, set the Default and Wide Char Codes to UTF-8.
4. Save your configuration changes, and restart WebLogic.

WebLogic

(WebLogic 9.x only) FileNet P8 requires the BEA CR298435 fix when WebLogic 9.x is the application server.

To obtain the fix, request the CR298435_920.jar file from BEA support, or download the WebLogic 9.2 MP1 maintenance pack from the BEA maintenance pack download Web site.

To install the BEA CR298435_920.jar fix on Windows

1. Modify setDomainEnv.cmd CLASSPATH to point to the location of the CR298435_920.jar file.

To install the BEA CR298435_920.jar fix on UNIX

1. Modify your setDomainEnv.sh to reflect the absolute location of the patch jar

For example:

```
CLASSPATH="/home/opt/  
CR298435_920.jar${CLASSPATHSEP}${PRE_CLASSPATH}${CLASSPATHSEP}${WEBLOGIC_CLASSPATH}  
${CLASSPATHSEP}${POST_CLASSPATH}${CLASSPATHSEP}${WLP_POST_CLASSPATH}"  
export CLASSPATH  
  
CLASSPATH="/home/opt/  
CR298435_920.jar${CLASSPATHSEP}${PRE_CLASSPATH}${CLASSPATHSEP}${WEBLOGIC_CLASSPATH}  
${CLASSPATHSEP}${POST_CLASSPATH}${CLASSPATHSEP}${WLP_POST_CLASSPATH}"  
export CLASSPATH
```

Security Administrator

Extended characters and user names

Note the following considerations for localized FileNet P8 accounts:

- With Microsoft Active Directory, Process Engine supports:
 - Extended characters in user names and passwords for all Latin1, Latin2, and Arabic languages
 - Extended characters in passwords for double-byte languages (Users names are not supported due to Microsoft limitations.)
- 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. However, support for a Windows non-ASCII user name is limited to the Process Engine database code page. For example, if the Process Engine database code page is Japanese, Process Engine can support only Japanese and English user names. If the Process Engine database code page is Latin1, Process Engine can support all Latin1 user names, such as English, French, German, Spanish, and so forth.

Database Administrator

Install Microsoft SQL Server

During installation, the Microsoft SQL Server installer program detects the Windows locale and sets the Microsoft SQL Server language setting accordingly. Use the locale selected by the installation program throughout the entire Microsoft SQL Server installation. Microsoft does not recommend changing the selected locale unless you have to match the locale to the collation of another instance of Microsoft SQL Server or to the Windows locale of another computer.

NOTE Localized versions of Microsoft SQL Server are available in French, German, Spanish, Italian, Japanese, Korean, and Simplified and Traditional Chinese.

Process Engine does not support Unicode UTF-8 encoded characters when using Microsoft SQL Server. Process Engine stores characters according to the Process Engine Window regional setting.

Using a binary sort for Process Engine is faster when configuring Microsoft SQL Server. However, IBM does support other collation settings. (A specific list of supported sort orders is not available. The collation setting depends on which sort order you need to use in your production environment.)

Queried information returns in the selected sort order. The sort order can be either binary or based on a combination of sort options, such as language, case-insensitive (ci), case-sensitive (cs), accent-insensitive (ai), and accent-sensitive (as).

The collation designator is a collation name that is based on a Windows locale. When you specify a collation designator, you should specify a sort order as well.

For example:

French - Latin 1

SQL Collations: Dictionary order, case-sensitive, for use with 1252 Character Set (or any case-sensitive MS-SQL collation).

Install Oracle server

Before installing Oracle, verify that the operating system locale is set appropriately. For more information, see <http://www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96529/ch3.htm#49757>.

Content Engine

A single instance of Content Engine on an Oracle server in a Windows environment can support multiple languages by:

- Selecting the Unicode AL32UTF8 database character set when configuring the Oracle database.
- Installing the appropriate Windows language packs.

Process Engine

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

Some examples of Process Engine supported database character sets are:

- AL32UTF8 – Supports all languages that Process Engine currently supports.
- JA16SJIS – Supports only Japanese and English.
- WE8ISO8859P15 – Supports West European with Euro currency support, but not Japanese.

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

Install Oracle client

Consider the following language information when installing Oracle Client software.

NLS_LANG environment variable - Charset

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

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

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

NOTE The NLS_LANG environment variable setting must match the UNIX locale or the Windows code page of the Process Engine server if the Oracle database is remote. You do not need to set this variable if Oracle was installed locally with the correct locale or regional setting prior to Process Engine installation.

The format of the NLS_LANG environment variable is:

NLS_LANG = language_territory.charset

The following table provides a brief explanation about the environment variable arguments.

NOTE For more information, see the Oracle documentation.

Arguments	Description
Language	Specifies the language used for Oracle messages, day names, and month names. Each supported language has a unique name, for example, American, French, or Japanese. The language argument has default values for the territory and character set arguments, so either (or both) of the territory or the character set arguments can be omitted. If a language is not specified, the value defaults to American.
Territory	Specifies the default date format and decimal character used for numbers. Each supported territory has a unique name, for example, America, France and Japan. Because each language provides a default territory, a value is not required.
Charset	Specifies the character set used by the client application (normally that of the user's terminal). Each supported character set has a unique acronym, for example, US7ASCII, WE8ISO8859P1, or JA16EUC. Because each language provides a default character set, a value is not required.

To set the NLS_LANG variable on Windows

1. Determine the Oracle database character set that best matches the Process Engine Windows server code page.
2. Open the Registry Editor by typing `regedit` at the command prompt.
3. Navigate to and update `HKEY_LOCAL_MACHINE > SOFTWARE > ORACLE_HOME > NLS_LANG= language_territory.charset`.
4. Verify the character set matches the Process Engine Windows operating system code page.
5. Click **OK** to close the window and exit the Registry Editor.
6. For more information, see [http:// www.microsoft.com/globaldev/reference/WinCP.mspx](http://www.microsoft.com/globaldev/reference/WinCP.mspx).

UNIX

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

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

To set the NLS_LANG environment variables

1. Modify the fnsw user.profile file in /home/fnsw:

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

Where:

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

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

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

Table 1:

OS	Locale/ Windows Code Page	Remote Oracle Character Set	NLS_LANG for PE Server fnsw User	Comments
AIX	LANG= JA_JP.UTF-8	AL32UTF8	NLS_LANG= Japanese_Japan. UTF8	Because the AIX locale is Unicode, use the UTF-8 character set.
Sun Solaris	LANG= ja_JP.eucJP	JA16SJIS	NLS_LANG= Japanese_Japan. JA16EUC	Because Solaris data will be in eucJP, the Oracle JA16EUC character set is the best match.
HP-UX	LANG= ja_JP.eucJP	AL32UTF8	NLS_LANG= Japanese_Japan. JA16EUC	Because HP-UX data will be in eucJP, the Oracle JA16EUC character set is the best match.
Windows	Japanese (Shift-JIS)	AL32UTF8	NLS_LANG= Japanese_Japan. JA16SJIS	Because Shift-JIS is the code page for Windows, the Oracle JA16SJIS character set is the best match.

Table 1:

OS	Locale/ Windows Code Page	Remote Oracle Character Set	NLS_LANG for PE Server fnsw User	Comments
Windows	Japanese (Shift-JIS)	JA16SJIS	NLS_LANG= Japanese_Japan. J A16SJIS	Because Shift-JIS is the code page for Windows, The Oracle JA16SJIS character set is the best match.
Windows	Japanese (Shift-JIS)	JA16EUC	NLS_LANG= Japanese_Japan. JA16SJIS	Although the database character set is JA16EUC, the Oracle character set in NLS_LANG is JA16SJIS since Shift-JIS is the code page for Windows.

Install the DB2 server

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. It 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

Operating system considerations

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

NOTE 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

Process Engine has limited language support on Windows. Process Engine supports only those languages that are supported by the Windows code page and a corresponding database character set.

For example, if the active Windows code page and the corresponding database character set support Latin 1 languages, PE will support such languages as English, French, Spanish, and Italian (and other Latin 1 languages) because these languages share a common character code page. However, Japanese cannot be supported in this scenario because the Japanese language uses a different character code page. Conversely, if the active Windows code page and the corresponding database character set support Japanese, Japanese (and English) will be the only languages supported by PE.

NOTES

- All Windows character code pages support English in addition to their primary language(s).
- In UNIX environments, PE 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.

Microsoft Windows

Use the localized Windows version when available. If the localized version is not available, use the English version with the appropriate regional setting.

Use the Regional Options Control Panel to change the regional setting. For more information, see the Windows help system.

UNIX

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

Install the AIX® language packs. For more information, see the AIX installation documentation.

To configure support for other languages in X Windows UNIX system

1. Add language fonts for your UNIX operating system if necessary to display an X-Windows desktop in a specific language UI. Follow your UNIX operating system administration guide to install other language fonts.
2. 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.

To set the LANG, LC_TIME, and LC_MESSAGES environment variables

1. 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.

NOTES

- 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"
```

FileNet P8 Administrator

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.

CAUTION 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.

Configure 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.

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

Limitations

This section documents non-English installation limitations.

WebLogic

WebLogic 8.1 is not supported on a Turkish operating system.

English and EnglishX locale considerations

It is important to choose the correct locale when choosing between English and EnglishX. Each locale has different performance and stemming query characteristics:

- Indexing in an English locale is twice as fast as indexing in an EnglishX locale. Some variance on the InZone modifier is attributed to the way the two locales use a stemming table. Stemming identifies the root form of words, so that all forms of the word can be searched. This extends to words with irregular tenses, such as the word "eat" that is stemmed to include "ate."

NOTE Using an EnglishX locale with the Autonomy K2 categorization product for indexing is probably more complete in its linguistic analysis.

For more information, see the Autonomy K2 documentation.

- An English locale uses a static stemming table; whereas, an EnglishX locale dynamically creates a stemming table based on the words. As a result, an EnglishX locale does a linguistic analysis and can recognize a different set of words as similar. Both locales might also have different sets of similar words.

NOTE Your locale and style folder must match. For example, if you select English as your locale, your style folder cannot be EnglishX; it must also be English.

GB18030 Support for Windows PE

PE installed on Windows does not support GB18030 Simplified Chinese characters.

AE, Workplace XT, or PE Task Manager for UNIX

If AE, Workplace XT, or PE is installed on a UNIX platform, Process Task Manager (vwtaskman), takes too long to launch when executed under a UTF-8 locale using remote CDE shells such as xWindows (exceed, xManager, etc). Verify 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 CE server, double-byte documents cannot be published. This is a known 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 PE 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
```

IBM FileNet P8 database character sets

IBM FileNet P8 Database components character sets

The table below lists the character sets supported for the IBM FileNet P8 Oracle and SQL Server database components.

Database type	Character sets supported in the regular character set (via varchar and char data types)	Character sets supported in the national character set (via nvarchar, nchar, and graphics data types)	Character set used by Content Engine	Character set used by Process Engine
Oracle 10g and 11g	All standard Oracle character sets AL32UTF8 Unicode character sets	Not used by IBM FileNet P8	Regular character set AL32UTF8 recommended	Regular character set
MS SQL Server	Only non-Unicode code pages are supported here	UCS-2 (Generically called Unicode on SQL Server)	Unicode/national character set only	Regular character set only

Additional Oracle character set information

The Oracle database character set you choose depends on whether the database is shared or dedicated.

- The recommended character set for a database dedicated to Content Engine or shared with Process Engine is AL32UTF8 (Unicode 3.1 UTF-8 universal character set).

- If the database is dedicated to Process Engine, choose from among the supported character sets listed in the following table.

Description	Character set
Unicode 3.1 UTF-8 universal character set	AL32UTF8
Western European	(UNIX) WE8ISO8859P1 or WE8ISO8859P15 (Windows) WE8MSWIN1252
Eastern European	EE8ISO8859P2
South European	SE8ISO8859P3
Northern & Northeastern Europe	NEE8ISO8859P4
Latin/Cyrillic	CL8ISO8859P5
Latin/Arabic	AR8ISO8859P6
Latin/Greek	FL8ISO8859P7
Latin/Hebrew	IWISO8859P8
Western European & Turkish	WEISO8859P9
North European	NE8ISO8859P10
ASCII 7-bit American	US7ASCII

Be aware of the special NLS_LANG settings for character sets and locale on Oracle Client machines. Consider the following:

- You must ensure that the NLS_LANG Oracle environment variable on an Oracle Client machine matches the character set/locale of the client operating system.
- Under Windows, by default the Oracle Client installer sets the NLS_LANG value in the Windows registry to match the locale of the Oracle Client machine's operating system. For Process Engine you do not need to override the registry value with the user environment variable. The default NLS_LANG value is adequate for either a Unicode (for example, AL32UTF8) or non-Unicode (for example, WE8MSWIN1252) database character set.
- Under UNIX, the Oracle Client installer does not automatically set the NLS_LANG (as it does under Windows). For this reason, after you have installed Process Engine, you must manually set the locale and character set value on each UNIX Oracle Client machine from which users will access Process Engine.
- If the database uses a non-Unicode character set (e.g., WE8MSWIN1252), you can make the NLS_LANG value on the Oracle Client machine match the character set of the database by choosing the appropriate database character set during database creation. For example, when the client operating system locale is American ANSI 1252, you can designate the WE8MSWIN1252 character set when you create the database. The Windows Oracle Client installer will set NLS_LANG to be AMERICAN_AMERICA.WE8MSWIN1252 because the client

operating system locale is American ANSI 1252. Under UNIX, you would manually set this value.

- You set the NLS_LANG value manually on Oracle Client machines as follows:
 - (UNIX) Add NLS_LANG to the shell environment login files for each user who will be logging on to the machine to run IBM FileNet P8 software.
 - (Windows) Set or modify the value of the NLS_LANG key using System Properties in the Control Panel for each user who will be logging on to the machine to run IBM FileNet P8 software.

IBM FileNet P8 ports

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

- [“Content Engine ports” on page 265](#)
- [“Process Engine ports” on page 266](#)
- [“Application Engine ports” on page 268](#)
- [“Autonomy K2 Search Engine ports” on page 270](#)
- [“Rendition Engine and Content Engine ports for Lipient” on page 271](#)
- [“Process Analyzer ports” on page 273](#)
- [“Process Simulator ports” on page 273](#)
- [“System Manager ports” on page 274](#)
- [“Content Federation Services for Image Services ports” on page 274](#)

Table 2: 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 via 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 via EJB.
WebSphere WSI	9080	This port is on the WebSphere application server for Content Engine, for communication with Content Engine by clients via WSI.
WebLogic	7001	This port is on the WebLogic application server for Content Engine, for communication with Content Engine by clients.
JBoss EJB	1099	This port is on the JBoss application server for Content Engine, for communication with Content Engine by clients via EJB.

Table 2: Content Engine ports

Port name	Default number	Notes
JBoss WSI	8080	This port is on the JBoss application server for Content Engine, for communication with Content Engine by clients via WSI.
JBoss IIOP	3528	This port is on the JBoss application server for Content Engine, for communication with Content Engine by clients via IIOP.
DB2	50000	This port is on the DB2 database server, for communication with the database by Content Engine.
Microsoft SQL Server	1433	This port is on the Microsoft SQL Server database server, for communication with the database by Content Engine.
Oracle	1521	This port is on the Oracle database server, for communication with the database by Content Engine.

Table 3: 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.
Microsoft SQL Server	1433	This port is on the database server, for connection between Process Engine and Microsoft SQL Server.
Oracle	1521	This port is on the database server, for connection between Process Engine and Oracle.
Oracle Services for Microsoft Transaction Server (Windows)	2030	This port is on the Oracle (Windows) database server.
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.
COR	32769 (TCP)	Courier service. COR_Listen listens on this port for incoming RPC requests.

Table 3: Process Engine ports

Port name	Default number	Notes
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)	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 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.

NOTE If the port number assigned to Component Manager conflicts with the port number required by another application or service that runs on the Process Engine or the Application Engine server, Process Task Manager will not start and the necessary vtaskman.xml will not be automatically created. If this happens, make a copy of the sample vtaskman.xml.sample file located on the Process Engine or Application Engine. On Process Engines, the file is located in the /fns/bn directory. On Application Engines, the file is located in Drive:\Program Files\FileNet\AE\Router on Windows and in /opt/FileNet/AE/Router on UNIX. Open vtaskman.xml.sample with a text editor, change the port element value to an available port number, and save the file to vtaskman.xml in the same directory.

Table 4: 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 via SSL.
WebLogic	7001	This port is on the WebLogic application server for clients to connect to Application Engine.
WebLogic SSL	7002	This port is on the WebLogic application server for clients to connect to Application Engine via SSL.
JBoss	8080	This port is on the JBoss application server for clients to connect to Application Engine.

Table 4: Application Engine ports

Port name	Default number	Notes
JBoss SSL	8443	This port is on the JBoss application server for clients to connect to Application Engine via 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 WP 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.
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.
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.

Table 5: Autonomy K2 Search Engine ports

Port name	Default number	Notes
K2 Administrative Server	9950 (default)	
K2 Dashboard/ Business Console	9990 (default)	
K2 Index Server	9960 - 9979 (recommended range)	
K2 Broker Server	9900 - 9909 (recommended range)	
K2 Server	9920 - 9949 (recommended range)	
K2 Ticket Server	9910 - 9919 (recommended range)	

Table 6: Rendition Engine and Content Engine ports for Liquent

Port name	Default number	Notes
Liquent input port	2867 (COM Repository only)	This port allows for distributed processing of jobs on the Rendition Engine servers when there are more than one RE 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 execute 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.
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 server when the other Rendition Engine server has finished processing the job. One case of this is when the 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 Engine server 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).

Table 6: Rendition Engine and Content Engine ports for Liquent

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.
Microsoft SQL Server	1433	When the selected Rendition Engine database is Microsoft SQL Server, then this port is used by the Rendition Engine servers and the Rendition Engine clients (Content Engine Publishing server) to access the Rendition Engine database.
Oracle	1521	When the selected Rendition Engine database is Oracle, then this port is used by the Rendition Engine servers and the Rendition Engine clients (Content Engine Publishing server) to access the Rendition Engine database.

NOTE 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 RE server, one must create a "Java Repository" in the RE/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).

Table 7: Process Analyzer ports

Port name	Default number	Notes
Registry port	32771	This is the default RMI port on the PA server. The Process Task Manager application on the PA server communicates with the PA server process on this port. The Process Simulator also communicates with the PA server on this port. The Process Simulator can be installed on the PA server or be remote.
PA database port		This port is on the PA database server. The PA server communicates with the PA database on this port. The PS also communicates with the PA database on this port.
DB2	50000	
Microsoft SQL Server	1433	
Oracle	1521	
PA database port		This port is on the PE database server. The PA server communicates with the PE database on this port.
DB2	50000	
Microsoft SQL Server	1433	
Oracle	1521	

Table 8: Process Simulator ports

Port name	Default number	Notes
Process Analyzer	32771	This port is on the PA server (Registry Port). PS communicates with PA on this port.
Registry port	32771	This port is on the PS server. The Process Task Manager application on the PS server communicates with the PS server process on this port. The AE also communicates with the PS server process on this port.
PA database port	1433	This port is on the PA database server and is used by the PS to communicate with the PA database.
Return	0 (anonymous)	This port is on the PS server and is used to communicate with the AE server. By default an anonymous port number is used. However, if the PS server resides behind a firewall it will be necessary to specify an explicit port by entering a value other than 0.

Table 9: System Manager ports

Port name	Default number	Notes
Listener (first)	32775	This is the primary “pilot port” for connection to the System Manager server. This port also registers the other secondary allocated ports and communicates those numbers to the Dashboard. If needed, an administrator can use the PchConfig.properties file to override the OS-defined property and define a specific range of ports to use.
Listener (subsequent)	OS defined	If the first listener port is allocated, the OS will allocate additional ports for managers to connect to listeners on the System Manager server. If needed, an administrator can use the PchConfig.properties file to override the OS-defined property and define a specific range of ports to use.

Table 10: Content Federation Services for Image Services ports

Port name	Default number	Notes
tms	32768 (TCP)	‘tms’ is the 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.
cor	32769 (TCP)	‘cor’ is the 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)	‘fn_snmpd’ is the 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.

Table 10: Content Federation Services for Image Services ports

Port name	Default number	Notes
snmp trap	162 (UDP)	This is a well-known OS trap daemon port for listening to trap messages. All Image Services trap messages received by fn_trapd daemon are eventually routed to this port.
fn_trapd	35225 (UDP)	'fn_trapd' is the Image Services trap daemon, which listens for notifications of the end of Image Services background processes running on the server and sys_logs information.
Native default SNMP port (HP only)	8000 (UDP)	All non-Image Services based SNMP requests are routed to this port for native SNMP processing.
IBM FileNet specific SNMP port (HP and Solaris only)	8001 (UDP)	All Image Services based SNMP requests are routed to this port for fn_snmpd daemon processing.
Migration notify	anonymous	ds_notify and pri_notify processes send TPI notifications to signal completion of a task.

NOTE On UNIX platforms, Image Services port assignments are made in the /etc/services file. For more information, see the *IBM FileNet P8 Content Federation Services for Image Services Guidelines*. To download this guide from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19](#).

Index

A

accounts, create
 Application Engine 85
 Content Engine 66, 184
 Content Search Engine 87
 Process Engine 79, 80
 Application Engine
 administration 24
 authentication
 IBM Tivoli Directory Server 64
 Novell eDirectory 64
 Sun Java System Directory Server 63
 Windows Active Directory 62
 Windows AD LDS 63

C

character sets 264
 configurations, sample 24
 baseline 25
 baseline with options 26
 demo 28
 developer 26, 27
 Content Search Engine
 create accounts 87

D

data sources
 ODBC for Process Engine 94
 database
 dedicated (DB2) 109, 113
 dedicated (Oracle) 97
 dedicated (SQL Server) 89
 GCD (DB2) 119, 199
 GCD (Oracle) 92, 100, 199, 200
 GCD (SQL Server) 91
 shared (DB2) 109, 113
 shared (Oracle) 97
 shared (SQL Server) 89
 storage area 53
 supported character sets 262
 DB2
 dedicated database 109, 113
 install for Content Engine 109
 install for Process Engine 109
 JDBC drivers, WebLogic 135, 209
 JDBC drivers, WebSphere 131, 139, 205, 212
 shared database 109, 113
 Decru 53
 default time mask, Solaris Process Engine 46, 170

DNS forwarder, enable 62
 domain
 Verity (K2) 87
 WebLogic 142

E

encrypted NTFS devices 53
 Encryption products for storage 53
 Enterprise Manager
 run on terminal services from developer workstations 27
 extensible authentication ports 267

F

file storage area
 defined 53
 file store, see file storage area
 Fix Packs and Test Fixes
 minium level required 151
 fixed storage area 53

H

host file 39, 50, 162, 173

I

IBM Encryption Expert 53
 iPlanet, see Sun Java System Directory Server

J

JBoss
 configure for Application Engine 143
 configure for Content Engine 138, 211
 JDBC drivers 131, 138, 205, 212
 WebLogic, install for DB2 135, 209
 WebLogic, install for Oracle 136, 210
 WebLogic, install for SQL Server 135, 209
 WebSphere, install for DB2 131, 139, 205, 212

K

Kerberos 60

L

LANG settings, Solaris Process Engine install 46, 170

N

NLS, see character sets

O

Object Request Broker (ORB)
 application server-based PE clients [144](#)
 UNIX-based PE Clients [48](#), [171](#)
 Windows-based PE Clients [51](#), [174](#)
 object store
 DB2 database [119](#)

Oracle

 client install [106](#)
 dedicated database [97](#)
 install for Content Engine [97](#)
 install for Process Engine [97](#)
 JDBC drivers, WebLogic [136](#), [210](#)
 shared database [97](#)

P

ports [265](#)
 Process Engine
 enable port 32771 [45](#), [169](#)
 time zone settings [44](#), [168](#)

R

realm [60](#)

S

security
 authentication [63](#), [64](#)
 install considerations [60](#), [181](#)
 specify accounts [65](#), [183](#)
 Single Sign-On, see SSO
 SQL Server
 client install [95](#)
 dedicated database [89](#)
 install for Content Engine [89](#)
 install for Process Engine [89](#)
 JDBC drivers, WebLogic [135](#), [209](#)
 shared database [89](#)
 SSO
 Kerberos [60](#)
 planning considerations [60](#)
 requirements [60](#)
 Storage areas [53](#)
 Sun Java System Directory Server [63](#)
 Sun ONE, see Sun Java System Directory Server

T

time mask, Solaris Process Engine [46](#), [170](#)

U

Upgrade, flow charts [147](#)
 Upgrade, overview [147](#)
 Upgrade, scenarios [147](#)
 Upgrade, staged [147](#)
 Upgrade, standard [147](#)
 users and groups, create
 Process Engine
 create users and groups [81](#)

V

verify TCP/IP settings [49](#), [173](#)
 Vormetric [53](#)

W

WebLogic
 configure for Application Engine [142](#), [143](#)
 configure for Content Engine [134](#)
 WebSphere
 configure for Application Engine [141](#)
 configure for Content Engine [131](#)
 Windows 2008 inbound rules [50](#), [173](#)
 Workplace
 avoid logon failures in multi-forest domains
 [135](#), [209](#)
 configuration [24](#)

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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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.

Licensees 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.

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.

Trademarks

IBM, the IBM logo, and 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 "Copyright and trademark information" at 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.



Program Number: 5724-R76, 5724-R81

Printed in USA

GC19-2554-00

