IBM FileNet P8 Platform

Version 4.5

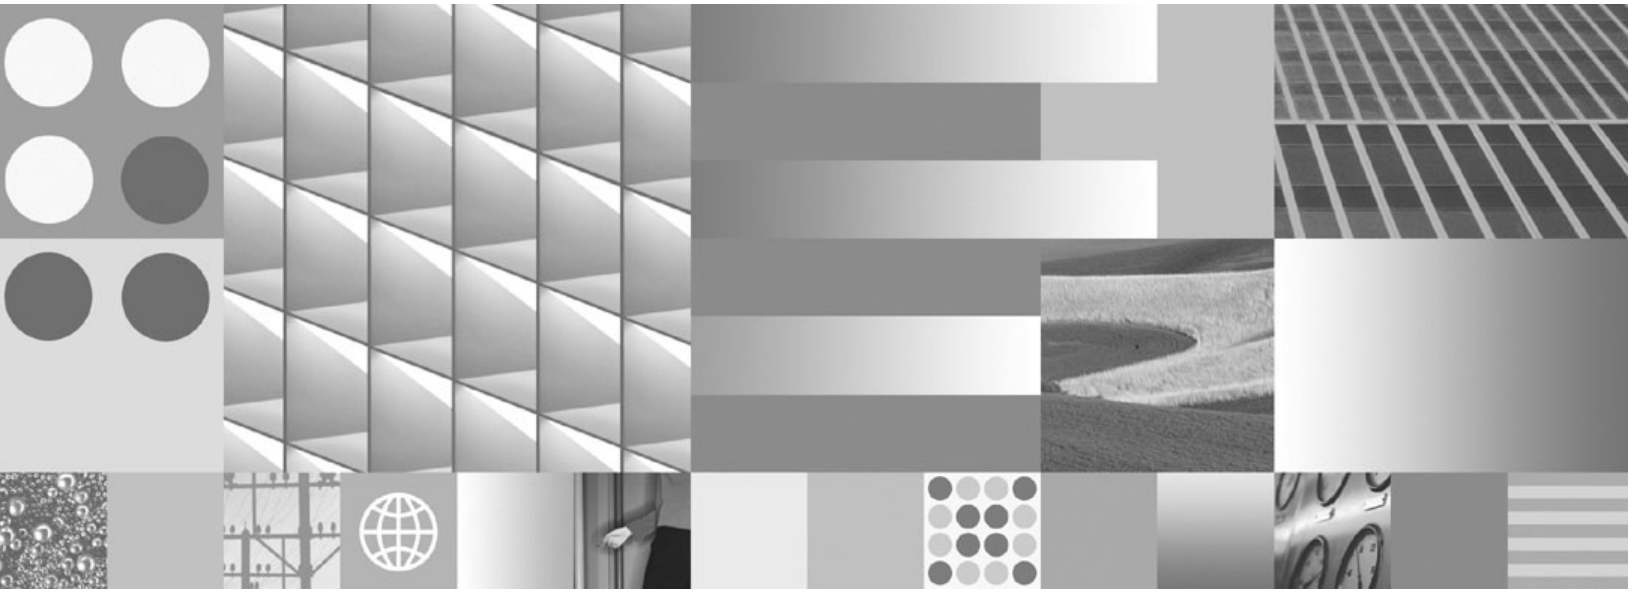

Installation and Upgrade Guide

GC19-2553-00

IBM FileNet P8 Platform

Version 4.5

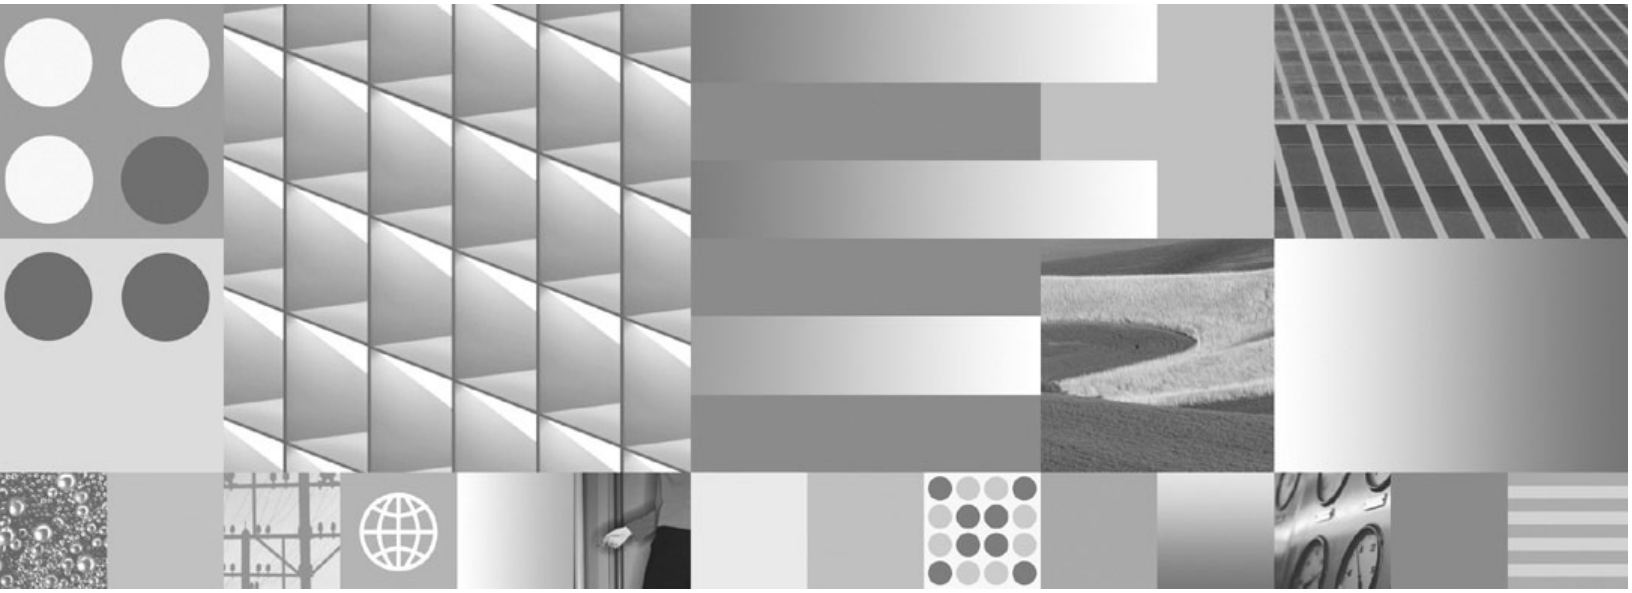

Installation and Upgrade Guide

GC19-2553-00

Note

Before using this information and the product it supports, read the information in "Notices" on page 459.

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.

Table of contents

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 installation and upgrade information about IBM FileNet P8 Platform, which includes Application Engine (AE), Content Engine (CE), and Process Engine (PE).

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.

Topics in this document apply specifically to FileNet P8 Platform installation tasks. Information about preparing the FileNet P8 Platform environment before you start these installation tasks is in the *Plan and Prepare Your Environment for IBM FileNet P8* document.

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**.
Italics	Variables that require user-supplied values	The calculation is: *number of object stores* * 16 + *number of concurrent users*.
	Document titles	You are reading the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
`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 `filed` file. Enter a value for the `new` property. Select the `senior` 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: `COPY filename`

Acronyms

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

Product Name	Acronyms
Application Engine	AE
Content Engine	CE
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:

- “Other available documentation” on page 17
- “Install IBM FileNet P8 Platform documentation” on page 21
- “Update documentation search index” on page 25

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...
Plan and Prepare Your Environment for IBM FileNet P8	To confirm the target environment is ready for a FileNet P8 installation or upgrade.
IBM FileNet Workplace XT Installation and Upgrade Guide	For information about how to install and upgrade Workplace XT in a FileNet P8 environment.
IBM FileNet P8 Hardware and Software Requirements	To confirm that the target environment has at least the minimum supported levels of software from independent software providers.
IBM FileNet P8 Compatibility Matrix	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.
IBM FileNet P8 Release Notes	To familiarize yourself with the new features and known issues associated with the FileNet P8 release.
IBM FileNet P8 Non-English Support Guide	For information about how FileNet P8 supports non-English environments.
IBM FileNet P8 High Availability Technical Notice	For information about how to set up FileNet P8 using clusters, farms and other high-availability software and hardware.
IBM FileNet P8 Performance Tuning Guide	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.
IBM FileNet P8 Troubleshooting Guide	For troubleshooting tips associated with a FileNet P8 Installation.

Document or help topic name...	Refer to this document...
IBM FileNet P8 help 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.
IBM FileNet P8 help 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...
IBM FileNet P8 System Overview	For a technical summary of the FileNet P8 architecture, including a overview of features and capabilities.
IBM FileNet P8 Disaster Recovery Technical Notice	For information about potential options and solutions involved in a FileNet P8 disaster recovery plan.
IBM FileNet P8 Process Task Manager Advanced Usage Technical Notice	For information about properties found under the **Advanced** tab in Process Task Manager.
IBM FileNet P8 Version Tools Technical Notice	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.
IBM FileNet P8 Security Help Extract	For security-related information from the FileNet P8 help system in PDF format.
IBM FileNet Rendition Engine Installation and Upgrade document	For information about how to install and upgrade Rendition Engine in a FileNet P8 environment.
IBM FileNet P8 eForms Installation and Upgrade Guide	For information about how to install and upgrade eForms in a FileNet P8 environment.
IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide	For information about how to install and upgrade IBM FileNet Connectors for SharePoint Web Parts in a FileNet P8 environment.
IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide	For information about how to install and upgrade IBM FileNet Connectors for SharePoint Document Libraries in a FileNet P8 environment.

Document name	Refer to this document...
IBM FileNet P8 Portlets for WebSphere Installation and Upgrade Guide	For information about how to install and upgrade FileNet P8 Portlets for WebSphere in a FileNet P8 environment.
IBM FileNet Process Analyzer Installation and Upgrade Guide	For information about how to install and upgrade Process Analyzer in a FileNet P8 environment.
IBM FileNet Process Simulator Installation and Upgrade Guide	For information about how to install and upgrade Process Simulator in a FileNet P8 environment.
IBM FileNet Records Manager Installation and Upgrade Guide	For information about how to install and upgrade Records Manager in a FileNet P8 environment.
IBM FileNet Business Activity Monitor Installation and Configuration Guide	For information about how to install and configure Business Activity Monitor in a FileNet P8 environment.
IBM FileNet Content Federation Services Installation and Upgrade Guide	For information about how to install and upgrade Content Federation Services in a FileNet P8 environment.
IBM FileNet Content Federation Services for Image Services Planning and Configuration Guide	For information about how to configure Image Services for document federation.
IBM FileNet Content Federation Services for Content Manager OnDemand Planning and Configuration Guide	For information about how to configure Content Manager for OnDemand for document federation.
IBM FileNet Content Management Widgets Installation Guide	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.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).

Install and configure IBM FileNet P8 Platform

This section contains the following tasks:

Install IBM FileNet P8 Platform documentation

Peform the tasks in this section to install IBM FileNet P8 documentation. See the following topics:

Task 1: Install IBM FileNet P8 Platform documentation (All)

This topic covers the installation, and (if necessary) search-related indexing of your IBM FileNet P8 Platform documentation. Because the IBM FileNet P8 Platform documentation includes a Java-based full-text search engine, it must be run as a web-based application.

If you intend to access online help from within any IBM FileNet P8 applications (for example, Workplace, Enterprise Manager, and Process Task Manager) or use the full-text search feature. you must install the documentation for IBM FileNet P8 Platform and its expansion products on a supported web application server (for example, WebSphere, WebLogic or JBoss).

Optionally, you can install and index web-posted auxiliary documentation that is available solely on the IBM web site (for example, the *IBM FileNet P8 Release Notes*), to have it included in the core IBM FileNet P8 Platform platform search functionality. For a complete list of such documentation, see "Gather auxiliary documentation" in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

NOTES

For all Web Servers

- You can collocate the IBM FileNet P8 documentation on an application server with either Application Engine or Content Engine server components installed.
- To ensure proper documentation search functionality, make sure that JavaScript™ support is enabled on each user's browser client.
- Depending on your operating system and application server levels, your options might be slightly different than those documented in the examples below.
- You will have to re-index the IBM FileNet P8 documentation any time you install (or re-install) the core IBM FileNet P8 Platform documentation and add on any expansion products or auxiliary documentation.
- You can use multipart root directories (for example, /docs/ecm_help) if your application server supports them.
- If you are upgrading your IBM FileNet P8 Platform installation and have downloaded a refresh of the documentation from the IBM Information Management support page on www.ibm.com.
- If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (chmod 755) to the **indexFiles.sh** file.

For WebSphere

- IBM FileNet P8 documentation must be deployed as a WAR file (ecm_help.war). You cannot deploy as an EAR file because in that case the IBM FileNet P8 help files remain packaged in EAR files. The internal documentation search engine will not function as expected.

For WebLogic and JBoss

- You cannot install the IBM FileNet P8 help files as WAR or EAR files because in those cases, WAR file support is not available, and the IBM FileNet P8 help files remain packaged in EAR files. The internal documentation search engine will not function as expected.

- A valid Java SDKmust be installed or available for the WebLogic and JBoss application servers.
- Unpack the ecm_help.war file or use the ecm_help flat file structure to install the IBM FileNet P8 documentation.

Overview of Procedures

Perform the procedures in the following subtopics in the order presented, unless otherwise directed below:

- "Install the core IBM FileNet P8 Platform documentation" on page 23 -- Begin here if you are installing the IBM FileNet P8 Platform documentation for the first time, or re-installing the core IBM FileNet P8 Platform documentation.
- "Install expansion-products, custom and auxiliary documentation" on page 24 -- Begin here if you have already installed the core IBM FileNet P8 Platform documentation and need to add on the documentation for IBM FileNet expansion products (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services, or IBM FileNet Records Manager). Optionally, you could also add web-posted auxiliary or customized documentation.
- "Update documentation search index" on page 25 -- Perform this procedure only after you have added on documentation for all of your expansion products. If you only need to install the core IBM FileNet P8 Platform documentation, which has a pre-configured baseline search index, you can skip this procedure. You should also plan for disaster recover of the installed IBM FileNet P8 Platform documentation.
- "Complete and verify the documentation installation" on page 26 -- Perform this final procedure in all cases, but only after you have installed the core IBM FileNet P8 Platform documentation, added on documentation for your expansion products, and regenerated the search index.

Install the core IBM FileNet P8 Platform documentation

This procedure establishes the core web application for documentation associated with the IBM FileNet P8 Platform. It assumes that your application server is already installed and operational.

To install the core IBM FileNet P8 Platform documentation

1. Access the IBM FileNet P8 Platform Documentation packageand perform the following actions, depending on your application server type:

   WebSphere

   a. Copy the IBM FileNet P8 ecm_help.war file from the IBM FileNet P8 documentation package to an appropriate location on the local hard drive.

   b. From the WebSphere administrative console, deploy the *ecm_help.war* file using **ecm_help** as the Context Root.

WebLogic

a. Unpack the IBM FileNet P8 **ecm_help.war** file or copy the entire **ecm_help** directory structure from the IBM FileNet P8 documentation package to a location on the local hard drive.

b. From the WebLogic administrative console, install a new web application for the IBM FileNet P8 Platform documentation.

JBoss

a. Unpack the IBM FileNet P8 **ecm_help.war** file or copy the entire **ecm_help** directory structure from the IBM FileNet P8 documentation package to a location on the local hard drive.

2. Continue as follows:

- If you need to add documentation for any IBM FileNet P8 expansion product or any web-posted auxiliary documentation, go on to the procedure in the following topic, "Install expansion-products, custom and auxiliary documentation" on page 24.
- If you have no further documentation to install, go on to the procedure in the topic "Complete and verify the documentation installation" on page 26.

Install expansion-products, custom and auxiliary documentation

Use the procedure in this topic to install documentation for expansion products or custom documentation onto an existing IBM FileNet P8 Platform documentation server.

Optionally, you can install auxiliary documentation that is available solely on the IBM (www.ibm.com) web site. For example, the *IBM FileNet P8 Release Notes* and *IBM FileNet P8 Troubleshooting Guide plus many others can be found there*.

If you do not have additional documentation to add on, skip to the procedure "Complete and verify the documentation installation" on page 26.

To install expansion product documentation

1. Determine the expansion products documentation location, either on local media source or network location. For most expansion products, use the *Documentation* package included as part of the particular software product(s).

2. Stop the IBM FileNet P8 Platform ecm_help documentation application (by using the WebSphere or WebLogic administrative console, or by shutting down the JBoss application server) and verify that no processes are accessing the documentation web application.

3. **For all web application servers:** Copy the IBM FileNet P8 Platform expansion product documentation (all directories and files below the ecm_help directory), to the core IBM FileNet P8 Platform documentation location (ecm_help).

   NOTE You can copy more than one expansion product documentation set to the documentation application server (for example, FileNet P8 eForms and FileNet Content Federation Services). When you are finished, you end up with one ecm_help directory

structure containing the core documentation set and one or more sets of expansion product documentation files.

4. Download the latest web-posted updates of planning and preparation, installation and upgrade guide PDF files for the IBM FileNet P8 Platform and various expansion products.

   Note Check the *Documentation* page on the IBM Information Management support page on www.ibm.com for the latest versions of these guides. See also, "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19 for details.

5. Repeat these steps for each of your expansion products, custom applications and auxiliary documentation.

6. Go on to the procedure in the following topic, "Update documentation search index" on page 25.

Update documentation search index

Perform this procedure only if you have refreshed the core IBM FileNet P8 Platform documentation, and installed expansion-products or auxiliary documentation onto your IBM FileNet P8 Platform documentation server. Otherwise, skip to the procedure in the following topic, "Complete and verify the documentation installation" on page 26.

NOTE Every time you update the documentation search index, a backup copy of the files in the existing Index/core subdirectory will be automatically copied to the Index/IndexOld subdirectory. You can reapply these backed-up files to the core subdirectory (after first moving or deleting the files created there) if you should need to return to the previous indexed state.

To update the documentation search index

1. If the IBM FileNet P8 Platform documentation application (ecm_help) is running, stop it. Verify that no processes are accessing the IBM FileNet P8 Platform documentation application.

   NOTE Make sure you have copied all of the IBM FileNet P8 Platform expansion-products and auxiliary documentation that you intend to use to the application server location containing the core IBM FileNet P8 Platform documentation. Otherwise, you will have to repeat this procedure if you install additional documentation later.

2. Open a terminal or command prompt on the application server. From the command line, navigate to the search subdirectory under the application root directory, for example, ecm_help.

3. Using a text editor, open the search-indexing script file that is appropriate for your operating system:

   UNIX

   indexFiles.sh

   Windows

   indexFiles.bat

   NOTE You might need to set the permissions on the file used, as it is set to read-only in the documentation package.

4. If necessary for your environment, set the JAVA_HOME variable in the script file with the path to your Java™ Runtime Environment (JRE) installation location.

   NOTE The Java JRE installation subdirectory can be user-defined, so substitute your actual location, as appropriate.

5. Save your changes and close the text editor.

6. From the ecm_help/search directory, run the updated search-indexing script file that is appropriate for your server operating system:

   UNIX

   indexFiles.sh

   Window

   indexFiles.bat

   NOTE As you run the search-indexing script file, you might notice periodic Parse Abort errors. You can ignore these error conditions, as they are benign and do not affect the overall indexing process.

(Optional) Prepare for disaster recovery or easy redeployment

Now that you have built your complete IBM FileNet P8 Platform documentation location (*ecm_help*), you can regenerate the *ecm_help.war* file using the *ecm_help* directory structure that contains your core IBM FileNet P8 Platform documentation and all added-on IBM FileNet P8 expansion products. This regenerated *ecm_help.war* file could then be placed into safe disaster recovery storage or used to quickly bring additional P8 Platform documentation servers on-line.

Complete and verify the documentation installation

Perform this procedure after you have installedand, if necessary, re-indexed the IBM FileNet P8 documentation on the application server.

To complete and verify the documentation installation

1. Start the IBM FileNet P8 documentation (*ecm_help*) application.

2. From your web browser, access the URL for your web environment, using your documentation server name, and port number, as in these examples:

   WebSphere

   http://yourdocserver:9080/ecm_help/

   WebLogic

   http://yourdocserver:7001/ecm_help/

JBoss

http://yourdocserver:8080/ecm_help/

The help system should open.

NOTE Apply the URL for your application server when it is time to configure the online help location for the various IBM FileNet P8 components, either while running installation programs or later via site preferences settings (for example, in Workplace or Workplace XT).

3. Click the **Search** link on the Help Directory toolbar. The documentation Search page should open.
4. Enter a value for your Search query.
5. Select one of the Search query result links. The associated help page should open.

Install and configure Content Engine

This section contains the following major topics:

Task 1: Install Content Engine

Use a procedure in this topic to install the following Content Engine components interactively or silently:

- Content Engine Server
- Configuration Manager
- .NET Clients (including FileNet Enterprise Manager)
- Content Engine Upgrader
- FileNet Deployment Manager

On UNIX-based or Windows-based application servers, you can install Content Engine Server, Configuration Manager, or both. The remaining components are installable only on Windows-based machines.

Install the .NET Clients software only on machines where you intend to run either the FileNet Enterprise Manager administrative client or a customized client application.

The Content Engine Upgrader is required only to upgrade Content Engine and Content Search Engine data from an earlier version.

The FileNet Deployment Manager is an optional component that you can use to deploy test systems into full production. See the IBM FileNet P8 help topic Application Deployment > Get Started with FileNet Deployment Manager for details about how you might use this tool.

If you are installing in a managed environment, the application server must be the Deployment Manager node (WebSphere) or the Administrator node (WebLogic)

To install Content Engine

1. Log on as *ce_install_user* on the machine where you are going to install Content Engine software.
2. Access the Content Engine software package.
3. (Red Hat Linux only) On the machine where you are going to run the Content Engine installation program, install the shared library libstdc++.so.5:
   a. Access the Red Hat 5 installation software.
   b. Install the package compat-libstdc++-33-3.2.3-61.i386.rpm, which contains the necessary shared library libstdc++.so.5.

4. Start the Content Engine installation using your installation worksheet. For information on the Content Engine parameter values, see “Installation and Upgrade Worksheet” in *Plan and Prepare Your Environment for IBM FileNet P8*.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select CE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - To install interactively:

     i. Run one of the following commands in the software package, depending on your operating system, and then follow the instructions on the wizard screens:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN`
HPUX	`P8CE-4.5.0-HPUX.BIN`
HPUXi	`P8CE-4.5.0-HPUXI.BIN`
Linux	`P8CE-4.5.0-LINUX.BIN`
Solaris	`P8CE-4.5.0-SOL.BIN`
Windows	`P8CE-4.5.0-WIN.EXE`
zLinux	`P8CE-4.5.0-ZLINUX.BIN`

     Remember to select the Content Engine Server option to install the core server software, and the Tools option if, on this machine, you intend to install Configuration Manager, which is the tool you must run to complete the configuration of the Content Engine Server instance. On Windows, the Tools option also installs Content Engine Upgrader and FileNet Deployment Manager.

     At the Install Complete screen of the wizard, selecting the Launch Configuration Manager check box will automatically start the tool for configuring Content Engine.

   - To install silently:

     i. Open the CE_silent_install.txt file in the software package for editing.

     ii. Set the parameter values in the CE_silent_install.txt file for your site.

     iii. Run one of the following commands in the software package, depending on your operating system:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN -f CE_silent_install.txt -i silent`
HPUX	`P8CE-4.5.0-HPUX.BIN -f CE_silent_install.txt -i silent`

Platform	Command
HPUXi	`P8CE-4.5.0-HPUXI.BIN -f CE_silent_install.txt -i silent`
Linux	`P8CE-4.5.0-LINUX.BIN -f CE_silent_install.txt -i silent`
Solaris	`P8CE-4.5.0-SOL.BIN -f CE_silent_install.txt -i silent`
Windows	`P8-CE-4.5.0-WIN.EXE -f CE_silent_install.txt -i silent`
zLinux	`P8CE-4.5.0-ZLINUX.BIN -f CE_silent_install.txt -i silent`

5. Check for errors in the Content Engine error log file ce_install_log_4_5_0.txt.
6. Continue at "Configure Content Engine instances" on page 32.

Task 2: Configure Content Engine instances

Use the Configuration Manager procedures in this topic to configure an instance of Content Engine Server. Once you complete the configuration tasks in this topic and some related update tasks in subsequent topics, you will be instructed to use this same Configuration Manager tool to deploy the Content Engine instance.

Configuration Manager comes in two variations:

- command-line interface
- graphical user interface

You can run either version to configure an instance. After configuring an instance, you will deploy it in "Deploy Content Engine instances" on page 46. You can configure all the instances before deploying any of them, or you can configure and deploy one instance at a time.

NOTES

- (WebSphere only) For best results, configure no more than one Content Engine instance in a profile.
- On a machine that runs Novell SUSE Linux Enterprise 9, you can configure Content Engine only from a command line.

Refer to the values in your installation worksheet to provide the Content Engine configuration property values needed for completing the procedures in this chapter. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

Delete existing data sources as needed

One of the tasks you will run in Configuration Manager is to create data sources for the GCD. You will do this either with the graphical user interface (using the Configure JDBC Data Sources tab) or the command line interface (using the ConfigureJDBC task).

Configuration Manager will not create a new data source with the same name as that of an existing data source. If you want to reuse a data source name, manually delete the existing data source before creating the new data source.

Grant permissions to the Configuration Manager user

Perform the following procedure to grant the file and directory permissions required by *config_mgr_user*, the user who will run Configuration Manager. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.

To grant permissions to the Configuration Manager user

1. If you have not done so already, log on to the machine where you installed Content Engine as *ce_install_user*.

2. Navigate to *ce_install_path*/tools/configure, which contains both the command line and graphical user interfaces of Configuration Manager, where *ce_install_path* is the path in which you installed Content Engine.
3. Grant *config_mgr_user* execute permission (UNIX) or read & execute permission (Windows) on the executable file of the interface of Configuration Manager you intend to use:
   - To enable use of the graphical user interface, grant permission to one of the following files in the /configure/CMUI directory:

     UNIX

     cmui

     Windows

     cmui.exe
   - To enable use of the command line interface, grant permission to one of the following files in the /configure directory:

     UNIX

     configmgr.sh

     Windows

     configmgr.bat
4. Grant write permission to the directory where you want Configuration Manager to place the configuration XML files it will generate.

   If you are not going to specify this directory when you run Configuration Manager, grant write permission on the default directory, *ce_install_path*/tools/configure/ and all its files and subdirectories.
5. Log off the machine hosting the application server and log back on as *config_mgr_user*, the Configuration Manager user.
6. Continue at one of the following subtopics:
   - "Configure instances using the graphical user interface" on page 33
   - "Configure instances using the command line interface" on page 36. If you need an accessible version of Configuration Manager, use the command line interface instead of the GUI.

Configure instances using the graphical user interface

In this subtopic you will configure a Content Engine Server instance on an application server using the graphical user interface version of Configuration Manager. Use the information in your worksheet to specify the values for the parameters required to configure Content Engine. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

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 you must specify for Configuration Manager:

- Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select Configuration Manager.
- Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

Refer to the appendix "Configuration Manager user interface reference" on page 418 for complete information on using the graphical user interface.

To create a configuration profile

1. Navigate to the *ce_install_path*/tools/configure/CMUI directory, which contains the graphical user interface of Configuration Manager.
2. Start Configuration Manager by running one of the following commands, depending on the operating system that runs on the machine where you installed Content Engine:

   UNIX

   ```
   cmui
   ```

   Windows

   ```
   cmui.exe
   ```

3. Start the Create New Configuration Profile wizard, by selecting **File > Create New Configuration Profile** or by clicking the icon in the tool bar, and use the information in your worksheet to specify the values for the parameters in the wizard screens. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.
4. Enter a name for the profile. The name must be a valid directory name for your operating system.
5. Specify the path for the profile. Either type in the full path to the profile directory or click **Browse** to locate the directory. The default path is *ce_install_path*/tools/configure/profiles, where *ce_install_path* is the location where Content Engine is installed.
6. Click **Next**.
7. Select the application server type for this profile.
8. Click **Next**.
9. Enter the property values for the application server. The specific properties displayed depend on the server type that you selected in Step 7.
10. Click **Next**.
11. Select the tasks that you want to include in this profile. For a new Content Engine installation, you need to complete all four configuration tasks: Configure JDBC Data Sources, Configure LDAP, Configure Bootstrap, and Deploy Application.

12. Click **Finish** to create the profile.

    The profile you created will be displayed as an icon in the left-hand pane, along with three icons for the configure tasks, and one icon for the deploy task:

    - Configure JDBC Data Sources
    - Configure LDAP
    - Configure Bootstrap Properties
    - Deploy Application

To configure a Content Engine instance using the Configuration Manager graphical user interface

You will perform the configure tasks in this procedure, and the deploy task in "Deploy Content Engine instances" on page 46. Performing a task means editing it, running it, and then checking its completion status.

NOTE The Configuration Manager allows you to deploy a Content Engine Server instance at the same time you perform the required configuration tasks. However, you must also perform the deployment task any time you install Process Engine Client files on the Content Engine Server machine, which is required if you set up Process Engine in your IBM FileNet P8 environment. Therefore, to avoid having to perform the Content Engine deployment task twice in succession, perform the deployment after you install the Process Engine Client files.

You can perform the configure tasks in any order, and you do not need to complete work on one configure task before starting another.

1. (Optional) By default, Configuration Manager does not save passwords for the tasks you will edit. That is, each time you start Configuration Manager and open a saved profile, you will need to specify any passwords required by the tasks and for the application server properties. For more information, see "Handling passwords in Configuration Manager" on page 417.

2. Edit a configure task by double-clicking one of the configure icons, or by right-clicking the icon and choosing **Edit Selected Task**. Then specify the parameter values in the right pane of the window, using the appropriate information from your installation and upgrade worksheet. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

   After specifying all the required values, choose **File > Save** to save your work, or save your work if you need to exit Configuration Manager before completing your edits. If you do not save your edits, Configuration Manager prompts you to save your changes before running the task.

3. To run a configure task, select the **Task enabled** check box, right-click the configure icon, and choose **Run Task**. (Running the configuration task may take a few minutes.)

4. To check the completion status of a configure task, right-click the configure icon, and choose **Check Task Status**.

5. If any configure tasks remain for this profile, return to Step 2 and choose another configure task.

   NOTE You must run the Configure JDBC task at least 2 times to configure the minimum data sources. First enter the GCD data source values, save, and run the task. Then, enter the

object store data source values, save, and run the task again. You will need to edit and run the Configure JDBC task once for each additional object store.

6. Continue at "Install Content Engine software updates" on page 43.

Configure instances using the command line interface

In this subtopic you will configure a Content Engine Server instance on a given application server using the command-line version of Configuration Manager. Configuring a Content Engine instance involves four major steps, which are documented in detail in the procedural subtopics that follow. You must repeat all these steps to configure each Content Engine Server instance:

- Generate the configuration XML files that contain the properties and values used to perform various tasks required to configure the environment for a Content Engine instance. See "Generate the configuration XML files for a Content Engine instance" on page 36.
- Edit the configuration XML files by inserting your site's properties and values. See "Edit the configuration XML files for a Content Engine instance" on page 38.
- Execute the configuration XML files you edited. See "Execute the configuration XML files for a Content Engine instance" on page 39.
- Check that the configuration XML files that you executed have resulted in a correct configuration of the Content Engine instance. See "Check the completion status of Content Engine configuration tasks" on page 41.

If you are deploying multiple Content Engine instances on the same machine, you will need to generate, edit, and deploy a complete set of configuration files for each instance. Store the configuration files for each instance in a separate directory.

You can navigate through the steps above by generating all the configuration XML files before editing, executing, or verifying any of them; or you can generate, edit, execute, and verify one file at a time.

NOTE The following subtopic refers to Configuration Manager as `configmgr`. When you run the tool, substitute `configmgr.sh` (on UNIX) or `configmgr.bat` (on Windows) in place of `configmgr`, depending on your operating system.

Generate the configuration XML files for a Content Engine instance

The following table lists the configuration XML files that you will generate using Configuration Manager:

File Name	Description
configurejdbc.xml	Settings for configuring JDBC connections to the databases used by Content Engine. You will need to generate, edit, and execute the configure the JDBC task once for the data source for the Global Configuration Data (GCD) and once for the data source for each object store.
configureldap.xml	Settings for connecting to and searching within a directory server

File Name	Description
configurebootstrap.xml	Settings for creating the Global Configuration Data (GCD) and starting Content Engine
deployapplication.xml	Settings for deploying a Content Engine instance
configureapplicationserver.xml	Parent file of the above listed configuration XML files.
applicationserver.xml	Settings for the application server, including the location of the application server software and the name of the server.

You can generate the configuration XML files, in either of two ways:

- Run the tool multiple times, generating one configuration XML file at a time.
- Run the tool once to generate a “parent” file, configureapplicationserver.xml, and automatically generate all the configuration XML files. The parent file points to the individual configuration XML files.

Running the tool also generates the applicationserver.xml file (used in several configuration tasks). Subsequent executions of the tool will not overwrite applicationserver.xml.

To generate configuration XML files

1. Set the current directory to *ce_install_path*/tools/configure.
2. At the command prompt, run Configuration Manager to generate the configuration XML files all at once (Step a) or one at a time (Step b).

   See “Configuration Manager command-line reference” on page 429 for details on Configuration Manager syntax.

   a. To generate all the configuration files at once with a parent configuration file, run the following command. Do not type any line breaks when you type the command.

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureApplicationServer -path mypath
   ```

   where:

   app_server_type is WebSphere, WebLogic, or JBoss

   -path *mypath* is optional and specifies the path to where the generated configuration XML files will be placed. If you are deploying multiple Content Engine instances on the same machine, you will need to specify a separate directory for each instance.

   Continue at “Edit the configuration XML files for a Content Engine instance” on page 38.

   b. To generate a single configuration XML file, run the command in one of the following substeps:

   - To generate the configurejdbc.xml file:

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureJDBC -path mypath
   ```

- To generate the configureldap.xml file:

  ```
  configmgr generateConfig -appserver app_server_type -task ConfigureLDAP
  -path mypath
  ```

  ```
  configmgr generateConfig -appserver app_server_type -db db_type
  -ldap ldap_type -deploy deploy_type -task ConfigureLDAP -path mypath
  ```

- To generate the configurebootstrap.xml file:

  ```
  configmgr generateConfig -appserver app_server_type -task ConfigureBootstrap
  -path mypath
  ```

  ```
  configmgr generateConfig -appserver app_server_type -db db_type
  -ldap ldap_type -deploy deploy_type -task ConfigureBootstrap [-path mypath]
  ```

where *mypath* is optional and specifies the path to where the generated configuration XML files will be placed. If you are deploying multiple Content Engine instances on the same machine, you will need to specify a separate directory for each instance.

Repeat Step b to generate one of the other configuration XML files, or continue at "Edit the configuration XML files for a Content Engine instance" on page 38.

You will eventually need to generate each of the configuration XML files to configure a Content Engine instance.

Edit the configuration XML files for a Content Engine instance

Perform the following procedure for each file you generated in "Generate the configuration XML files for a Content Engine instance" on page 36 to insert your site's properties and values.

To edit values in the configuration XML files

1. Use a text editor or XML editor to open one of the following configuration XML files that you generated in "Generate the configuration XML files for a Content Engine instance" on page 36:

   - configurejdbc.xml
   - configureldap.xml
   - configurebootstrap.xml
   - applicationserver.xml

   NOTE If you generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 36, you will also have generated the deployapplication.xml file. You will open this file for editing in "Deploy Content Engine instances" on page 46.

2. In the configuration XML file, replace each occurrence of `****INSERT VALUE****` with a value appropriate for your site. Refer to the descriptions in the file for more information.

   Set the <TaskEnabled> value to `true` in any configuration XML file you edit if you want to run the configuration task in "Execute the configuration XML files for a Content Engine instance" on page 39.

NOTE If you previously specified values in the configureldap.xml file to add a realm to a federated repository, and want to put an additional realm in the repository, replace the previous values with the values for the additional realm.

3. (Optional, WebSphere only) If you have previously created XA and non-XA data sources that you want to use for the Global Configuration Data (GCD) database, make the following edits:

   a. In the configurejdbc.xml file, set the <TaskEnabled> value to `false` to avoid creating another pair (XA and non-XA) of data sources.

   b. In the configurebootstrap.xml file, set the <JDBCDataSourceXAFileName> and <JDBCDataSourceFileName> values to the XA and non-XA JNDI names, respectively, associated with the GCD database.

4. (Optional) Encrypt any passwords that you need to insert into the file by running the password encryption utility (see "To encrypt a password for Configuration Manager" on page 453), and then copy the encrypted value into the file. It is a best practice to encrypt the passwords for the following accounts:

   - The application server administrator account used in applicationserver.xml
   - The database administrator account used in configurejdbc.xml
   - The LDAP provider service principal account used in configureldap.xml
   - The master key—a word or phrase for encrypting sensitive FileNet P8 Global Configuration Data (GCD) entries—used in configurebootstrap.xml.

   CAUTION Any password you do not encrypt will be stored and sent as clear text.

5. Save your edits.

6. Continue at one of the following procedures:

   - Repeat Step 1 through Step 5 of this procedure for any other configuration XML file that you have not yet edited.
   - Continue at "Execute the configuration XML files for a Content Engine instance" on page 39 to execute a configuration XML file(s) you have edited.
   - Return to "Generate the configuration XML files for a Content Engine instance" on page 36 to generate additional configuration XML files.

Execute the configuration XML files for a Content Engine instance

Perform the procedure in this subtopic to execute the settings in the configuration XML files you have generated and edited for a Content Engine instance.

NOTES

- Any task with the <TaskEnabled> element value set to false will not run (see Step 2 on page 38 in "To edit values in the configuration XML files" on page 38).
- If you are executing tasks for a profile that was created or edited in the Configuration Manager GUI, verify that the XML files contain values for the required passwords. See "Handling passwords in Configuration Manager" on page 417 for more information.

To execute the configuration XML file settings

1. (WebSphere only) Start the application server if it is not already running.
2. Set the current directory to *ce_install_path*/tools/configure.
3. At the command prompt, run Configuration Manager to execute the configuration XML files all at once (Step a) or one at a time (Step b). See "Configuration Manager command-line reference" on page 429 for command syntax details.

   When running the tool, the `-path mypath` parameter is optional and specifies where you placed the configuration XML file. If you do not specify a path, the file must be in *ce_install_path*/tools/configure/profiles.

   a. To execute all the configuration files at once with a parent configuration file, run the following command.

      ```
      configmgr execute -task ConfigureApplicationServer -path mypath
      ```

      NOTES

      - To execute all tasks at once, you must have generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 36.
      - If you did not enable the deployment task in the deployapplication.xml file, the tool will display an informational message indicating that the deployment did not occur. In which case you will perform the deployment in "Deploy Content Engine instances" on page 46.

   b. To execute a single configuration XML file, type and run the command in one of the following substeps:

      - To execute the configurejdbc.xml file:

        ```
        configmgr execute -task ConfigureJDBC -path mypath
        ```

      - To execute the configureldap.xml file:

        ```
        configmgr execute -task ConfigureLDAP -path mypath
        ```

      - To execute the configurebootstrap.xml file:

        ```
        configmgr execute -task ConfigureBootstrap -path mypath
        ```

      Repeat Step b to execute one of the other configuration XML files, or continue at "Check the completion status of Content Engine configuration tasks" on page 41.

      You will eventually need to execute each of the configuration XML files to complete the configuration of a Content Engine instance.

      NOTE You must run the ConfigureJDBC task at least 2 times to configure the minimum data sources. First enter the GCD data source values, save, and run the task. Then, enter the object store data source values, save, and run the task again. You will need to edit and run the ConfigureJDBC task once for each additional object store.

4. Continue at "Check the configuration status of a Content Engine instance" on page 49.

Check the completion status of Content Engine configuration tasks

Use the procedure in this subtopic to verify that one or more Content Engine configuration tasks that you executed in "Execute the configuration XML files for a Content Engine instance" on page 39 have completed. You can check the status of all the tasks you executed or just check individual tasks.

NOTE Checking the completion status does not validate the information in the XML files.

To check the status of a Content Engine configuration

1. At the command prompt, run Configuration Manager to check the status of the configuration tasks all at once (Step a) or one at a time (Step b).

   When running the tool, the `-path mypath` parameter is optional and specifies where you placed the configuration XML file. If you do not specify a path, the file must be in *ce_install_path*/tools/configure/profiles.

   See "Configuration Manager command-line reference" on page 429 for command syntax details.

   a. To check the status of all the configuration tasks at once, run the following command:

      ```
      configmgr checkStatus -task ConfigureApplicationServer -path mypath
      ```

      NOTE To check the status of all tasks at once, you must have generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 36.

      Continue at Step 3.

   b. To check the completion status of a single configuration task, run the tool in one of the following substeps.

      - To check the status of the ConfigureJDBC task using the configuration file in the specified path:

        ```
        configmgr checkStatus -task ConfigureJDBC -path mypath
        ```

      - To check the status of the ConfigureLDAP task using the configuration file in the specified path:

        ```
        configmgr checkStatus -task ConfigureLDAP -path mypath
        ```

      - To check the status of the ConfigureBoostrap task using the configuration file in the specified path:

        ```
        configmgr checkStatus -task ConfigureBootstrap -path mypath
        ```

   If you performed Step b, repeat the step to check the status of any other configuration tasks you have executed; otherwise, continue at Step 2.

2. Continue at one of the following procedures, as needed; otherwise, continue at Step 3.

   - If you have any other configuration tasks to execute, continue at "Execute the configuration XML files for a Content Engine instance" on page 39.
   - If you have any other configuration XML files to generate, continue at "Generate the configuration XML files for a Content Engine instance" on page 36.

- If you want to add a realm to a federated repository, continue at "Edit the configuration XML files for a Content Engine instance" on page 38.

3. Continue at "Install Content Engine software updates" on page 43.

Task 3: Install Content Engine software updates

Perform the procedure in this topic for each Content Engine instance to install software updates fix packs, or interim fixes.

If no Content Engine software updates are available, skip to "Install the latest Process Engine Client files on Content Engine servers" on page 44.

To install the Content Engine software updates

1. For instructions on how to obtain the latest Content Engine software updates, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
2. Open the readmes for the Content Engine software updates and perform the installation procedures in the readmes on each Content Engine instance.

   If you installed an instance into a managed environment, perform the procedure on the Deployment Manager node (WebSphere) or the Administrator node (WebLogic).
3. Continue at "Install the latest Process Engine Client files on Content Engine servers" on page 44.

Task 4: Install the latest Process Engine Client files on Content Engine servers

To install the Process Engine Client files, perform the following steps on all application server machines where Content Engine Server is to be deployed. Except where noted, these steps apply to WebSphere, WebLogic, and JBoss.

To install the Process Engine Client files

1. On the machine where Content Engine is to be deployed, log on as *ce_install_user*.
2. Access the Process Engine Client install software from the Process Engine installation software. The version of the install software must match the version of Process Engine.
3. Expand the (TAR or ZIP) Process Engine Client install software.
4. (WebSphere only) Make a note of the max and min connection settings from the FileNetConnectionFactory.

   You will need to reset these when you deploy Content Engine. For details, see the task "Deploy Content Engine instances" on page 46.
5. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Content Engine will be deployed. Run the program either interactively (using the install wizard) or silently.

   To run the program interactively, run one of the commands in the table below, depending on the operating system and follow the wizard instructions:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN`

   To run the program silently, perform the following steps:

   a. In the expanded install software, open the file PEClient_silent_install.txt and edit it as follows:

      i. Change the `Variable_CheckboxCE` line to the following:

      ```
      -V Variable_CheckboxCE="true"
      ```

      ii. Save your edit.

b. Run one of the commands in the following table to perform the silent install:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN -silent -options "PEClient_silent_install.txt"`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN -silent -options "PEClient_silent_install.txt"`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN -silent -options "PEClient_silent_install.txt"`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN -silent -options "PEClient_silent_install.txt"`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN -silent -options "PEClient_silent_install.txt"`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE -silent -options "PEClient_silent_install.txt"`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN -silent -options "PEClient_silent_install.txt"`

Task 5: Deploy Content Engine instances

To deploy one or more instances of Content Engine, perform the procedures in this topic after you have performed the configuration tasks in "Configure Content Engine instances" on page 32.

To deploy a Content Engine instance, perform the procedures in one of the following subtopics:

NOTE On a machine that runs Novell SUSE Linux Enterprise 9, you can deploy Content Engine only from a command line.

Deploy instances using the Configuration Manager graphical user interface

Use the procedure in this subtopic to deploy a Content Engine Server instance on an application server using the Configuration Manager's graphical user interface.

Refer to the appendix "Configuration Manager user interface reference" on page 418 for complete information on the tool interface itself.

To deploy a Content Engine instance using the Configuration Manager graphical user interface

1. If you have not already done so, log on to the application server machine as *config_mgr_user*, the Configuration Manager user. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Navigate to the *ce_install_path*/tools/configure/CMUI directory, where:

   ce_install_path is the path to the Content Engine Server software.
3. Start Configuration Manager by running one of the following commands, depending on the operating system on the machine where you installed Content Engine Server:

   UNIX

   ```
   cmui
   ```

   Windows

   ```
   cmui.exe
   ```
4. Open the configuration profile for the Content Engine instance that you configured in "Configure Content Engine instances" on page 32.
5. Complete the deployment of the Content Engine Server instance:

   a. Perform the Deploy Application task by double-clicking its icon in the left-hand pane, or by right-clicking the icon and choosing **Deploy Task**.

b. Specify values for the parameters in the right pane of the window. Refer to the values in your installation worksheet. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

c. Select the **Task enabled** check box.

d. Choose **File > Save** to save your values. You can also save your work if you need to exit the tool before completing a task.

e. Right-click the deploy application icon.

f. Choose **Run Task**.

Running the deploy task may take a few minutes.

6. (Optional) Check the completion status of the Deploy Application task:

a. Right-click the deploy icon.

b. Choose **Task Status**.

7. Continue at "Complete post-deployment steps" on page 52.

Deploy instances using the Configuration Manager command line interface

Use the procedure in this subtopic to deploy a Content Engine Server instance on an application server using the Configuration Manager's command-line interface. Deploying a Content Engine Server instance involves four major steps:

1. Generate the deploy-application XML file, deployapplication.xml.
2. Edit the deployapplication.xml file by inserting your site's properties and values.
3. Execute the deployapplication.xml file.
4. Check that the deployapplication.xml file has resulted in a correct deployment of the Content Engine instance.

If you are deploying multiple Content Engine instances on the same machine, you must generate, edit, and execute a deployapplication.xml file for each instance. Store the deployapplication.xml file for each instance in a separate directory.

NOTE This subtopic refers to the Configuration Manager tool as `configmgr`. When you run the tool, substitute `configmgr.sh` (on UNIX) or `configmgr.bat` (on Windows) in place of `configmgr`, depending on your operating system.

Generate the deployapplication.xml file

If you did not create the deployapplication.xml file in "Configure Content Engine instances" on page 32, perform the following procedure for the Content Engine instance you are deploying; otherwise, skip to "To edit the values in the deployapplication.xml file" on page 48.

To generate the deployapplication.xml file

1. If you have not already done so, log on to the application server machine as *config_mgr_user*, the user who will run Configuration Manager. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Set the current directory to *CE_install_path*/tools/configure, where:

   CE_install_path is the path where you installed Content Engine.

   As an example, *CE_install_path* might be /opt/FileNet/ContentEngine.
3. Run the following command without carriage returns to generate the deployapplication.xml file:

   ```
   configmgr generateConfig -appserver app_server_type -deploy deploy_type
   -task DeployApplication -path mypath
   ```

   where:

   app_server_type specifies the type of application server (WebSphere, WebLogic, or JBoss).

   deploy_type indicates one of the following deployment types: `standard`, `cluster`, or `netdeploy` (network deployment). The netdeploy value applies only to Content Engine deployment on WebSphere.

   `-path` *mypath* is optional and specifies the path for the deployapplication.xml file. If you do not specify a path, the file will be placed in the *CE_install_path*/tools/configure/profiles directory.

   If you are deploying multiple Content Engine instances on the same machine, you must specify a separate path for each instance.

   See "Configuration Manager command-line reference" on page 429 for configmgr syntax details.

Edit the deployapplication.xml file for a Content Engine instance

Perform the following procedure to insert your site's properties and values in the deployapplication.xml file.

To edit the values in the deployapplication.xml file

1. Use a text editor or XML editor to open the deployapplication.xml file.
2. Edit the values in the deployapplication.xml file as appropriate for your site. Refer to the descriptions in the file for more information. In particular, perform the following substeps as applicable:

   a. By default, the value inside the <TaskEnabled> element in the deployapplication.xml file is set to false. For deployment to occur, set the value to true.

   b. (WebSphere only) Specify values for the following XML element values that apply to your deployment type:

      - For standard deployment or network deployment, specify values for <applicationservername> and <applicationservernode>

- For cluster deployment, specify values for <applicationservername>, <applicationservernode>, and <applicationserverclustername>

3. Save your edits.
4. Continue at "Execute the deployapplication.xml file for a Content Engine instance" on page 49.

Execute the deployapplication.xml file for a Content Engine instance

Perform the procedure in this subtopic to execute the settings in the deployapplication.xml file.

To execute the deployapplication.xml file settings

1. (WebSphere and WebLogic only) Start the application server if it is not already running.
2. Set the current directory to *CE_install_path*/tools/configure, where:

   CE_install_path is the path where you installed Content Engine server software.
3. At the command prompt, run Configuration Manager to execute the deployapplication.xml file.

   NOTE If you did not enable the deployment task in the deployapplication.xml file, the tool will display an informational message indicating that the deployment did not occur.

   ```
   configmgr execute -task DeployApplication -path mypath
   ```

   The `-path mypath` entry is optional and specifies the path where you placed deployapplication.xml. If you do not specify a path, the file must be in the following directory:

   CE_install_path/tools/configure/profiles

   If you are deploying multiple Content Engine instances on the same machine, you will need to specify a separate path to the deployapplication.xml file for each instance.

   See "Configuration Manager command-line reference" on page 429 for configmgr syntax details.
4. Continue at "Check the configuration status of a Content Engine instance" on page 49.

Check the configuration status of a Content Engine instance

Use the procedure in this subtopic to check all the tasks, or just the one you executed in "Edit the deployapplication.xml file for a Content Engine instance" on page 48.

To check the status of a Content Engine configuration

1. At the command prompt, run the following commands to check the status of all required Content Engine configuration tasks, or just the deploy-application task:

   ```
   configmgr checkStatus -task DeployApplication -path mypath
   ```

   The `-path mypath` entry is optional and specifies the path where you placed the deployapplication.xml file. If you do not specify a path, the file must be in the following directory:

   CE_install_path/tools/configure/profiles directory.

   See "Configuration Manager command-line reference" on page 429 for configmgr syntax details.

2. Continue at one of the following topics:
   - If the Content Engine instance is deployed in a JBoss cluster, continue at "Deploy Content Engine to additional JBoss servers in a cluster" on page 50.
   - If you have more Content Engine instances to deploy, continue at "Generate the deployapplication.xml file" on page 47.
   - If you have no more Content Engine instances to deploy, "Complete post-deployment steps" on page 52.

Deploy Content Engine to additional JBoss servers in a cluster

Use the following procedure to a deploy a Content Engine Server instance on each additional JBoss server in your cluster.

CAUTION Because JBoss clusters communicate state over multicast, it is critical that you configure clusters for Content Engine to segregate multicast traffic from each other and from other JBoss clusters. You can find instructions regarding JBoss cluster configuration at http://www.jboss.com.

To deploy Content Engine Server on additional JBoss servers in a cluster

1. Leave Content Engine deployed on the application server where you installed it.
2. Copy the following files from the initial Content Engine instance to each JBoss node you are setting up:
   - All data source (*JBoss_HOME*/server/all/deploy/*-ds.xml) files created as part of the configuration process.

     NOTE Depending on how you installed JBoss, there might be data source files not created by Configuration Manager, such as hsqldb-ds.xml, in the same location. Do not copy these files.
   - *JBoss_HOME*/server/all/deploy/Engine-ds.xml
   - The JDBC driver in *JBoss_HOME*/server/all/lib
   - *JBoss_HOME*/server/all/conf/login-config.xml
   - *JBoss_HOME*/server/all/deploy/Engine-jbc.ear

     NOTE Clustered JBoss instances are deployed from the "all" directory and not the "default" directory. You can find instructions regarding JBoss cluster configurations at http://www.jboss.com.
3. Start or restart the JBoss application server.
4. Verify that the server.log file located in the *JBOSS_home*/server/all/log directory lists deployment of the WAR or EAR file you used. If it does, then the deployment of Content Engine Server is successful.

5. Continue at one of the following topics:
   - If you have any more Content Engine Server instances to deploy, continue at "Execute the deployapplication.xml file for a Content Engine instance" on page 49
   - If you have no more Content Engine Server instances to deploy, continue at "Complete post-deployment steps" on page 52.

Task 6: Complete post-deployment steps

Before you can put an IBM FileNet P8 system into production, perform one of the following post-deployment procedures in this topic, depending on your application server type:

- "To complete post-deployment steps (WebSphere)" on page 52
- "To complete post-deployment steps (WebLogic)" on page 53
- "To complete post-deployment steps (JBoss)" on page 53

Then perform the procedure "To verify the deployment of Content Engine" on page 53.

To complete post-deployment steps (WebSphere)

1. If you are using federated user repositories, perform the following substeps; otherwise, continue at Step 2.
   a. Using the WebSphere administrative console, navigate to **Security > Secure administration, applications, and infrastructure** and click **Configure**.
   b. If you have not already done so, specify a unique user (short name) for the *Primary administrative user name*.

      NOTE This name must exist in one of the realms and must be unique.
   c. Specify the *Server user identity*. You can specify *Automatically generated server identity*, or specify one that exists in one of the repositories.

      NOTE This name must exist in one of the realms and must be unique.
   d. Save your changes to the master configuration.
2. Enable administrative security (if it is not already enabled) and application security before creating a FileNet P8 domain.

   NOTE If you want to enable WebSphere application or administrative security, you must do so manually. The deployment of Content Engine does not enable or check these settings.

   CAUTION You must disable Java2 security; otherwise, Content Engine will not be able to start or process requests.
3. Restart the WebSphere application server where Content Engine is deployed, as follows:
   - For a standalone server, stop and start the application server.
   - For a WebSphere network deployment, stop and start the application server where Content Engine is deployed, including the Deployment Manager and all managed nodes.
   - For a cluster, stop and start the cluster.
4. Continue at "To verify the deployment of Content Engine" on page 53.

To complete post-deployment steps (WebLogic)

1. Restart the WebLogic application server (to reinitialize the WebLogic cache) where Content Engine is deployed, as follows:
   - For a standalone server, stop and start the application server.
   - For a cluster, stop and start the cluster.
2. Continue at "To verify the deployment of Content Engine" on page 53.

To complete post-deployment steps (JBoss)

1. If your JBoss application server is not running, start it now.
2. Continue at "To verify the deployment of Content Engine" on page 53.

To verify the deployment of Content Engine

1. Verify the state of Content Engine deployment by browsing to the following web page:

   http://*server:port*/FileNet/Engine

   where:

   server is the host name of the machine where Content Engine Server is deployed.

   port is the WSI port used by the application server where Content Engine Server is deployed. Example web page addresses are shown in the following table:

Application Server Type	Web Page Address
WebSphere	http://*server:*9080/FileNet/Engine
WebLogic	http://*server:*7001/FileNet/Engine
JBoss	http://*server:*8080/FileNet/Engine

2. Verify that the web page contains the build and version number of Content Engine and its associated JAR files, and an indication that no errors occurred.

   NOTE It is a best practice to bookmark the web page address in your browser.

3. If you want to deploy another Content Engine instance on the application server, continue at "Configure Content Engine instances" on page 32.

Task 7: Install FileNet Enterprise Manager

If you did not install FileNet Enterprise Manager on the Content Engine server machine in "Install Content Engine" on page 29, you can install it now on the same Windows server or on some other Windows machine.

Microsoft .NET Framework 2.0 and Web Services Enhancements 3.0 must already be installed on the machine where you intend to install FileNet Enterprise Manager.

NOTE You cannot install FileNet Enterprise Manager on a UNIX machine.

To install FileNet Enterprise Manager

1. On the machine where you will install FileNet Enterprise Manager, log on as a member of the Local Administrators group or the Power Users group.
2. Access the Content Engine software package.
3. Start the FileNet Enterprise Manager installation. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select CE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - To install interactively:

     i. Run the following command in the software package:

     P8CE-4.5.0-Win.exe

     ii. When prompted, specify the following components to be installed:

     – .NET Clients (required for FileNet Enterprise Manager)

     – FileNet Enterprise Manager

     iii. Specify the install path, for example C:\Program Files\FileNet\Content Engine.

   - To install silently:

     i. Open the CE_silent_install.txt file in the software package for editing.

     ii. Set the parameter values in the CE_silent_install.txt file for your site. Be sure to set the CHOSEN_INSTALL_FEATURE_LIST parameter value to:

     `DotNetClients,AdminTools`

     iii. Save your edits.

iv. Run the following command in the software package:

```
P8-CE-4.5.0-WIN.EXE -f CE_silent_install.txt -i silent
```

4. Continue at "Establish the FileNet P8 domain and Global Configuration Data (GCD)" on page 56.

Task 8: Establish the FileNet P8 domain and Global Configuration Data (GCD)

With Content Engine installed and deployed, you will use Enterprise Manager to create a FileNet P8 domain.

To create a FileNet P8 domain

NOTE If you run Enterprise Manager as a limited user account, you cannot update the *Base URL for the FileNet P8 Platform help files* field in the General tab of the Enterprise Manager properties dialog box.

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager icon on the desktop, or by choosing **Start > All Programs > FileNet P8 Platform > Enterprise Manager**.
2. In the FileNet P8 Logon dialog box, click **Add**, to create a FileNet P8 domain configuration.

   NOTE For subsequent logons to Enterprise Manager, you can access an existing FileNet P8 domain by clicking **Connect**.
3. In the Add Domain Configuration dialog box, specify the required information and then click **OK**. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   In a non-managed application server environment (such as a JBoss cluster or a standalone WebSphere or WebLogic), set *ce_server_name* and *port* for just one of the application servers. You can configure multiple Enterprise Manager connections to the other non-managed application servers if you wish.
4. Continue at "To configure directory service authentication" on page 56.

To configure directory service authentication

NOTE For multi-realm authorization, run the Create a Directory Configuration wizard once for each realm. Refer to FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms.

1. In the FileNet P8 Logon dialog box, click **Connect**.
2. Specify the required values in the Create a Directory Configuration wizard screens. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   The configuration parameters required by the Create a Directory Configuration wizard are in many cases the same as those you provided to Configuration Manager when you configured LDAP in "Configure Content Engine instances" on page 32.

   Refer also to the topic for your directory service within the IBM FileNet P8 help topic System Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers.
3. In the Configure New Domain Permissions message box, click **OK** to acknowledge that the directory configuration is complete but remains in restricted mode. The Configure New Domain

Permissions wizard automatically starts. Continue at "To configure permissions for a FileNet P8 domain" on page 57.

To configure permissions for a FileNet P8 domain

1. Specify the required values in the Configure New Domain Permissions wizard screens. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   NOTE To edit the list of accounts having administrative access to the FileNet P8 domain, refer to the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to... > Add or remove a GCD administrator.

2. In the Configure New Domain Permissions message box, click **OK**.

To set the Statement Cache Size value for the GCD database

If you are using Microsoft SQL Server 2005 JDBC Driver or Oracle JDBC Driver, you need to set the Statement Cache Size parameter value to 0 for each data source you created to access the GCD database, as shown in the following steps:

1. Access the page containing the Statement Cache Size parameter:
   - (WebSphere) Navigate to the WebSphere administrative console page containing the field Statement Cache Size property. For example, in WebSphere 6.1.x, navigate to Resources > JDBC Providers > JDBC_provider > Data sources > *data_source* > WebSphere Application Server data source properties.
   - (WebLogic) Navigate in the tree view of WebLogic Administration Console to *FNCEDomain* > Services > JDBC > Data Sources > *Data_Source_Name* > Connection Pool.
2. Set the Statement Cache Size to 0 and save your change.
3. Continue at "Create the data sources for an object store" on page 58.

Task 9: Create the data sources for an object store

Perform either of the procedures in this topic to use Configuration Manager to create the distributed (XA) and non-distributed (non-XA) data sources for an object store.

NOTE Configuration Manager will not create a new data source with the same name as an existing data source. To reuse a data source name, manually delete the existing data source before creating the new data source.

To create the data sources using the graphical user interface to Configuration Manager

1. If you have not done so already, log on to the machine where you installed Content Engine as *ce_install_user*. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Start the GUI version of Configuration Manager by running one of the following commands, depending on the operating system on the machine where you installed Content Engine:

   UNIX

   cmui

   Windows

   cmui.exe
3. Start the Create New Configuration Profile wizard by choosing **File > New Configuration Profile** or by clicking the wizard icon in the tool bar.
4. In the New Configuration Profile Wizard screen, perform the following substeps:

   a. Specify a profile name. The name must be a valid directory name for your operating system.

   b. Select (highlight) an application server type on which to base the profile.

   c. Specify the application server properties using your install and upgrade worksheet. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

      HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for creating JDBC data sources:

      – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select "Create JDBC data sources."

      – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select "Installation."

   d. Clear the check boxes for all except the **Configure JDBC Data Sources** task.

   The profile you create will be displayed as an icon in the left-hand pane, along with the Configure JDBC Data Sources icon.
5. Edit the Configure JDBC task by double-clicking its icon, or by right-clicking the icon and choosing **Edit Selected Task**.

6. Specify the parameter values in the right pane of the window, using the appropriate information from your installation and upgrade worksheet, and save your profile. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.
7. Run the Configure JDBC task by double-clicking its icon, or by right-clicking the icon and choosing **Run Selected Task**.
8. Run the Check Status JDBC task by double-clicking its icon, or by right-clicking the icon and choosing **Run Selected Task**.
9. Continue at "Create the initial object store" on page 60.

To create the data sources using the command-line interface to Configuration Manager

1. If you have not done so already, log on to the machine where you installed Content Engine as *ce_install_user*. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Set the current directory to *ce_install_path*/tools/configure, where:

   ce_install_path is the path where you installed Content Engine.
3. Generate the configurejdbc.xml file by running the following command:

   ```
   configmgr generateConfig -appserver app_server_type -db db_type -ldap ldap_
   type -deploy deploy_type -task ConfigureJDBC -path mypath
   ```

   where *mypath* is optional and specifies the path to where configurejdbc.xml will be placed.
4. Use a text editor or XML editor to open the configurejdbc.xml file and edit it as follows:
   a. Replace each occurrence of `****INSERT VALUE****` with a value appropriate for your site. Refer to the descriptions in the file for more information.
   b. Set the <TaskEnabled> value to `true` so you can run the configuration task in Step 5.
   c. (Optional) Encrypt the password for the user of the database associated with the data sources you are creating by running the password encryption utility (see "Encrypt passwords" on page 453), and then copy the encrypted value into the file.

      CAUTION Any password you do not encrypt will be stored and sent as clear text.
   d. Save your edits.
5. Execute the configurejdbc.xml file by running the following command:

   ```
   configmgr execute -task ConfigureJDBC -path mypath
   ```
6. Check the completion status by running the following command:

   ```
   configmgr checkStatus -task ConfigureJDBC -path mypath
   ```
7. Continue at "Create the initial object store" on page 60.

Task 10: Create the initial object store

In this task you will create the initial object store. If you have not already done so, perform the following preparation before starting the Create an Object Store wizard.

- Set up the initial storage area for the object store (see "Prepare storage areas for object stores" on page 53 in *Plan and Prepare your Environment for FileNet P8*).
- Install the latest Content Engine service pack, as indicated in "Install Content Engine software updates" on page 43.
- Create XA and non-XA data sources for the object store, as indicated in "Create the data sources for an object store" on page 58).
- Read the FileNet P8 help topic System Administration > Content Engine Administration > Object stores > Concepts

NOTES

- The Create Object Store wizard will fail if you try to assign a new object store to a database/ tablespace that is not completely empty.
- Once an object store is created, you can refine its definition and add content to it.

To create an object store

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager desktop icon, or by choosing **Start > All Programs > FileNet P8 Platform > Enterprise Manager**.
2. In the FileNet P8 Logon screen, select the FileNet P8 domain in which you will create an object store, and then click **Connect**.
3. In the tree view, right-click the **Object Stores** container and choose **New Object Store** to start the Create an Object Store wizard.
4. Work through the wizard screens. For information on parameter names and values you will specify when running the wizard, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*. For additional details, click **Help** in the wizard screens.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select "Create Object Store wizard."

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header and select "Installation."

Task 11: Verify the Content Engine installation

To verify that the Content Engine installation succeeded, do the following procedure to confirm that, via Enterprise Manager, you can

- Create folders and documents
- Check documents in and out.

To verify the Content Engine installation

1. In Enterprise Manager, create a subfolder as follows:
   a. Expand the **Object Stores container** and expand the node for the object store you just created.
   b. Right-click the **Root Folder icon**, and select **New Sub Folder**.
   c. Enter the Folder Name and click **Create**.
2. Create a document as follows:
   a. Expand the **Root Folder container**.
   b. Right-click the **subfolder** you just created, and select **New Document**.
   c. Enter the Document Title (for example, Coffee Bean.bmp), click **With content**, and click **Next**.
   d. Click **Browse/Add** to select a file (for example, c:\winnt\Coffee Bean.bmp), click **Open**, and then click **Create**.

      NOTE The new containment name should be *Coffee Bean* with a major version of 1.
3. Check out a document as follows:
   a. Right-click the document object (**Coffee Bean.bmp**) you just created, and select **Exclusive Check Out** (Default).
   b. Navigate to the folder where you want the checked-out document to reside and click **Open**.
   c. Click **Yes** to edit the file.
   d. Make some change to the file and save it.
   e. Close the application you used to edit the file.
4. Check in a document, as follows:
   a. Right-click the document object (**Coffee Bean.bmp**) you just edited, and select **Check In**.
   b. In the File Name field, click **Browse/Add**.
   c. Select the file you have checked out (e.g., *Coffee Bean.bmp)* and click **Open**.
   d. Click **Check In**.

e. Right-click the **document object** you just checked in and click the **Versions** tab of the Properties dialog box.

f. You should now see the Major Version field change once the document is checked in.

To enable Enterprise Manager to display file storage area status

The Windows user who logged on to Enterprise Manager as a FileNet P8 user and then created a file storage area will see the status of the file storage area as *online*, and will be able to add content to it. All other Windows users, even if logged on to Enterprise Manager as the same FileNet P8 user who created the file storage area, will see its status as *offline*, and thus will not be able to add content to it.

1. Enable a Windows logon user to see an online status for a file storage area and to add content to it, by doing one of the following:
   - Add the user to the directory security of the file storage area with the following permissions: Modify, Read & Execute, List Folder Contents, Read, and Write.
   - As a member of the Local Administrators group, add the Windows logon user to a group account, as explained in the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Storage area security.

Install and configure Content Search Engine

This section contains the following major topics:

Task 1: Install Content Search Engine

Use this procedure to install and configure IBM FileNet P8 Content Search Engine, an optional component based on the Autonomy K2 product.

Overview

The Autonomy K2 software must be installed on each machine that is part of your Content Search Engine configuration. You must designate one machine in the configuration as the Master Administration Server. The Master Administration Server can be used as a standalone Content Search Engine, or additional K2 Administration Servers can be added. All K2 Administration servers are configured and controlled through the K2 Dashboard of the Master Administration server.

The Autonomy K2 software that underlies IBM FileNet P8 Content Search Engine has many inherent features that you might want to configure that are not discussed in the IBM FileNet documentation. For details, see the Autonomy documentation that is installed with the Autonomy K2 Master Administration Server located at:

http://*Master Administration Server hostname*:9990/verity_docs/

The Autonomy documentation set is not searchable from the IBM FileNet P8 Help but it does have its own internal index and search functionality.

CAUTION Although the K2 Dashboard provides you with documentation for, and direct interfaces to, the K2 collections, IBM FileNet requires that you use Enterprise Manager to manage collections associated with Content Search Engine index areas (for example, to add and remove index areas).

NOTES

- Autonomy K2 was previously known as Verity, and you will see Verity still used in many of the interfaces described in the following procedures.
- The machines in your Content Search Engine configuration must run the same operating system as your Content Engine. For UNIX, it to doesn't have to be the same type.
- Where machine name variables are required, IP addresses will not validate. A valid name must be entered.
- If you unimport the style set, the original files will be deleted from your system. In this scenario, if you wish to re-import the style set, you will need to recover it from your installation package. In order to avoid this situation, you can either enter a unique name for the Style Set Alias during the initial Content Search Engine (Autonomy K2) installation, or make a backup copy of the original style set. If you enter a unique name for the style set during installation, ensure you use that name when you configure Content Engine for Content-Based Retrieval.
- Stop word files can be used to increase performance by about 30%. You can put a file named style.stp into the stylefiles directory to list words you do not want full-text indexed (for example, short words such as a, the, and). However, using a stop word file also prevents searching on these words. See the K2 documentation for more details. To create a stop word file you can typically copy a file named vdk30.stp from either the main K2 install directory or the foreign language locales package over to the main stylefile directory, and then rename it to style.stp. You must do this copy operation before you create collections.

Indexing is case-sensitive, so the style.stp file should include capitalized versions of words in the stop word list in addition to the lower-case version. For example, use “the”, “The” and “THE” if you think all of these would be encountered on a search.

Install Content Search Engine on Windows

Install the Autonomy K2 software on each machine in your Content Search Engine configuration.

You must complete a K2 Master Administration Server software installation before you can add any additional K2 Administration Server installations. The Master Administration Server’s dashboard is the central control point for configuring all additional Administration Servers that are part of the Content Search Engine configuration.

You must designate each installation as either a Master Administration Server software installation or an Administration Server software installation.

CAUTION Only one K2 Master Administration Server can be installed for each Content Engine Domain.

To install Autonomy K2 on Windows

1. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.

   NOTE Ensure *k2_os_user* has administrator privileges on this machine.

2. Set the JAVA_HOME environment variable as follows:

   a. Open the System control panel.

   b. Click the **Advanced** tab.

   c. Click **Environment Variables**.

   d. Click **New** under System Variables.

   e. Set the variable information as follows:

      Variable name: *JAVA_HOME*

      Variable value: *Java (JDK) install path*

   NOTE The installer will not allow you proceed with the installation until the JAVA_HOME environment variable is set.

3. Install the Content Search Engine software, using the appropriate values from your worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

- Interactively
  i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-WIN.EXE file.
  ii. On the Select Autonomy K2 Server Type panel, choose whether to install the Master Administration Server software or the Administration Server software.
- Silently
  i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.
  ii. Launch the Content Search Engine installer by executing the following command:

```
P8CSE-4.5.0-WIN.EXE -f CSE_silent_install.txt -i silent
```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 6.2.1 Administration Server service.
- Verity K2 Administration Web Server service (Tomcat server).

Install Content Search Engine on UNIX

You must install the K2 Master Administration Server software first. The Master Administration Server's dashboard is the central control point for configuring Content Search Engine for single-server or multi-server installations.

You must designate each installation as either a Master Administration Server software installation or an Administration Server software installation.

CAUTION Only one K2 Master Administration Server can be installed for each Content Engine Domain.

To install Content Search Engine on UNIX

1. Configure permissions for the host machine so that permissions of the user who runs the IBM FileNet Content Search Engine installation are root. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic. If *k2_os_user* can not have root privileges on this machine, you can set permissions as follows, to allow *k2_os_user* to run the installation:
   a. Access the host machine and log on as a user with root privileges.
   b. Enter the following commands to set the vspget program's sticky bit such that the service runs as root and is in the same group as *k2_os_user* (default path shown):

```
chown root /opt/verity/k2/_rs6k43/bin/vspget
chmod g+rs /opt/verity/k2/_rs6k43/bin/vspget
```

2. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.
3. Set the following environment variables and place the entries in the .profile file for *k2_os_user*.

All UNIX types

```
JAVA_HOME=java_(JDK)_install_path/jdkversion
export JAVA_HOME
```

HP-UX

```
PATH=$PATH:/verity_install_path/k2/_hpux/bin
export PATH
SHLIB_PATH=$SHLIB_PATH:/verity_install_path/k2/_hpux/bin
export SHLIB_PATH
```

AIX

```
PATH=$PATH:/verity_install_path/k2/_rs6k43/bin
export PATH
LIBPATH=$LIBPATH:/verity_install_path/k2/_rs6k43/bin
export LIBPATH
```

Solaris

```
PATH=$PATH:/verity_install_path/k2/_ssol26/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ssol26/bin
export LD_LIBRARY_PATH
```

Linux

```
PATH=$PATH:/verity_install_path/k2/_ilnx21/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ilnx21/bin
export LD_LIBRARY_PATH
```

4. Install the Content Search Engine software, using the appropriate values from your worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

- Interactively
  i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-*UNIX TYPE*.BIN file.
  ii. On the appropriate panel, choose whether to install the Master Administration Server software or the Administration Server software.
- Silently
  i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.
  ii. Execute the following command:

```
./P8CSE-4.5.0-UNIX_type.BIN -f CSE_silent_install.txt -i silent
```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 6.2.1 Administration Server service.
- Verity K2 Administration Web Server service (Tomcat server).

To start or stop the Autonomy K2 Services on UNIX

To manually start or stop the Autonomy K2 services, use the following commands, according to your environment:

HP-UX

Start Services:

```
nohup /verity_install_path/k2/_hpux/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_hpux/bin/k2adminstop
```

AIX

Start Services:

```
nohup /verity_install_path/k2/_rs6k43/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_rs6k43/bin/k2adminstop
```

Solaris

Start Services:

```
nohup /verity_install_path/k2/_ssol26/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ssol26/bin/k2adminstop
```

Linux

Start Services:

```
nohup /verity_install_path/k2/_ilnx21/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ilnx21/bin/k2adminstop
```

Task 2: Configure Content Search Engine

Use this procedure to configure services required on the K2 Master Administration Server, and on additional Administration Servers you may install for IBM FileNet P8 Content Search Engine. All servers are configured through the Master Administration Server Dashboard.

NOTES

- When naming particular servers you create with this procedure, it is a good idea to indicate the type of server you've created. Otherwise, when you configure Content Engine through the Enterprise Manager, determining which server is which could be confusing. For example, use *servername*_broker to indicate that this is a Broker Server service.
- Ensure you carefully record the server names, ports and settings that you define. Much of the following information will be required later when you configure the IBM FileNet P8 Content Engine for Content-Based Retrieval later in this guide.
- A range of ports is recommended in the Verity K2 Dashboard for each service you create. You do not have to choose a port number from within that range.
- Repeat the related step in the procedure below to add additional services. Some guidelines must be adhered to when adding additional services:
  - Multiple Brokers can be assigned, so that if one goes down the others will be used. However, each broker must have all K2 Servers (search servers) attached that are needed to access collections (index areas). The Content Engine Server will not call multiple brokers and merge the results.
  - If you add additional Index Servers and K2 Servers (search servers), they will not be activated until you enable them through Enterprise Manager. See *Configure Content Engine for Content Based Retrieval* for details.
  - Each K2 Administration Server must contain a Ticket Server for Content Engine.
  - For good stability and performance, Broker Servers must be attached to local Ticket Servers for security on each machine.

To configure Content Search Engine

1. Configure the Autonomy K2 Dashboard to use SSL security. The Autonomy K2 Dashboard web application, by default, uses a non-SSL web site and sends username and password information in plain text. For information on how to modify your Tomcat web applications to use SSL, access the following address:

   http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html

2. Access the K2 Dashboard by launching your browser and entering:

   http://*Master Administration Server hostname*:9990/verity_dashboard/main.jsp

3. Create a K2 Index Server, as follows:
   a. Click **K2 Index Servers** under System View**.**
   b. Click **Add a K2 Index Server** on the K2 Index Server Summary page.

c. Enter the following information on the Configure basic settings for the new K2 Index Server page:
   - Service Alias: *server_name*_index_server
   - Port: 9960 - 9979 (suggested range)

d. Click **Next** to continue with the installation.

e. Enter the following information on the Configure threads for the K2 Index Server page:
   - Synchronous Threads: 25
   - Asynchronous Threads: 3
   - Access Type: **Authorized Administrator**

f. Click **Finish** to continue with the installation.

4. Set the Index Server logging properties:
   a. Click the Index Server, Under System View, that you want to adjust.
   b. Click **Edit Properties** under Actions.
   c. Click the **Logging** tab.
   d. For Status Log Nominal Size, enter the following value:

      9000 kilobytes

   e. Click **Modify**.
5. Create a K2 Broker Server:
   a. Click **K2 Brokers** under System View**.**
   b. Click **Add a K2 Broker** on the K2 Broker Summary page.
   c. Enter the following information on the Configure basic settings for the new K2 Broker page:
      - Service Alias: *servername*_broker
      - Port: 9900 - 9909 (suggested range)
   d. Click **Finish**.
6. Create a K2 Server (search server) and attach the Broker.
   a. Click **K2 Servers** under System View.
   b. Click **Add a K2 Server** under Actions on the K2 Server Summary page.
   c. Enter the following information on the Configure basic settings for the new K2 Server page:
      - Service Alias: *server_name*_search_server
      - Port: 9920 - 9949 (suggested range)
   d. Click **Next**.
   e. Click **Next** on the *Set security options for this service* page.

f. Enter the following information on the Attach to K2 Brokers page:

- Select the K2 Brokers that will provide access to this service: Select the K2 Broker you created in step 3 from the drop-down menu, *servername*_broker.

g. Click **Finish**.

7. Import the IBM FileNet Styleset.

NOTE The K2 Dashboard displays a notification that the StyleSet Editor web application cannot be accessed. This message can be ignored as it relates to a function that is not utilized by IBM FileNet Content Search Engine.

a. Click **Collections** under System View.

b. Click **Manage Style Sets** under Actions on the Collection Summary page.

c. Click **Import** on the Manage Style Sets page (top right).

d. Enter the following information on the Import page:

- Style Set Alias: FileNet_FileSystem_PushAPI
- Gateway Type: --Auto-detect--
- Source Administration Server (if multiple servers are installed). Choose which server.
- Source Path (default listed):

  Windows

  C:\Program Files\filenet\contentengine\verity\data\stylesets\FileNet_FileSystem_PushAPI

  UNIX

  /opt/verity/data/stylesets/FileNet_FileSystem_PushAPI

e. Click **Import**.

8. Create a K2 Ticket Server.

a. Click **K2 Ticket Servers** under System View.

b. Click **Add a K2 Ticket Server** under Actions on the K2 Ticket Server Summary page.

c. Enter the following information on the Configure basic settings for the new K2 Ticket Server page:

- Service Alias: *server_name*_ticket_server
- Port: 9910 - 9919 (recommended range)

d. Click **Next**.

e. Enter the following information on the Configure the login module to use with this K2 Ticket Server page:

- Select which Login Module type to use with this K2 Ticket Server:
  - Windows
  - UNIX

  NOTE LDAP Ticket Servers are not currently supported.
- Default Domain (Windows only): Enter the domain on which this K2 Server is authenticated.

f. Click **Next**.

g. Enter the following information on the Configure the persistent store module to use with this K2 Ticket Server page:

- Select the Persistent Store Module type to use with this K2 Ticket Server: Choose File and Memory.

h. Click **Finish**.

i. (Windows only) Specify local login settings:

i. Click **Edit Properties**.

ii. Click Windows Login Module.

iii. Check **Use Local Credentials:**.

iv. Check Enable Built-in Groups.

v. Click **Modify**.

9. Set Autonomy K2 Administration Security.

a. Click the K2 Ticket Server you created.

b. Click **Manage Administration Security** under Actions.

c. Enter the following information on the Manage Administration Security page:

- Select a K2 Ticket Server to configure for administration security: From the drop-down menu, select the K2 Ticket Server you just created.
- User Name: Enter the *k2_os_user*. For UNIX installs, this is the user you logged in as to run the installation. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.
- Password: Enter the authentication password.
- Default Domain (Windows only): Enter the domain on which this user and K2 Server are authenticated.

d. Click **Modify**.

K2 will authenticate the user based on the information you entered. If the check fails, an error message will indicate what failed and request that you re-enter the information.

If administrator access is successful, Autonomy K2 will close the Dashboard and require that you log on again as the Dashboard Administrator to complete the configuration.

10. Launch the K2 Dashboard and log on.
11. Restart K2 services:
    a. Under *Notifications* on the K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You must perform either a **Quick Restart** or a **Full Restart**.

       Click Home in the top-left corner of the page after each restart to view the remaining notifications. Repeat the process until there are no notifications remaining.
12. Enable additional K2 Admin Users (optional).
    a. From the K2 Dashboard home page, click **Administration Servers**.
    b. Click **Manage K2 Administrative Users**.
    c. Click **Add User** on the Manage K2 Administrative Users page.
    d. Enter the name of an authenticated user on the directory service that you want to make a K2 Administrator and click **Add**.
13. Enable security on the K2 services you have created.
    a. From the K2 Dashboard home page, click **K2 Ticket Servers**.
    b. Click your ticket server ***server_name*_ticket_server**.
    c. Click **Manage K2 Broker/K2 Server Security** in the Services Secured by this K2 Ticket Server section at the bottom of the page.
    d. Click the **K2 Servers** button on the Manage K2 Broker/K2 Server Security page.
    e. Click the service you just created which is listed in the window on the right to enable security.

       CAUTION If you have a multi-server configuration, numerous services, installed on other machines, will be listed also. Select only the service to which you want to attach a broker. Brokers must be attached to local ticket servers for Content Search Engine.
    f. Click the **K2 Brokers** button on the Manage K2 Broker/K2 Server Security page.
    g. Click the Broker in the window on the right that you want to attach to the local K2 ticket server you selected above.
    h. Click **Modify** to save your changes.

14. Restart K2 services, as follows:
    a. Click **Home** in the top left corner of the page.
    b. Under Notifications on the Verity K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You must perform either a **Quick Restart** or a **Full Restart**.

       Click **Home** in the top left corner of the page after each restart to view the remaining notifications. Repeat the process until there are no notifications remaining.

To configure services on Administration Servers

Use this procedure to create and configure services on specific Content Search Engine machines (Administration Servers), other than the Master Administration Server. Services for all machines in your Content Search Engine configuration are configured through the Master Administration Server Dashboard.

1. Click **Administration Servers** under System View**.**
2. Click the Administration Server to which you want to add services.
3. Click **Add a Service** under K2 Services on this Administration Server.
4. Select the service you want to add.
5. Follow the instructions and guidelines for the appropriate service in the *To configure Content Search Engine* procedure above to complete the service addition.

Task 3: Install additional locales (optional)

Complete this procedure only if you require locales other than English.

NOTES

- The panels in the Locales installer indicate that the installer is only compatible with Autonomy K2 version 2.6.0. However, the Locales installer is compatible with version 2.6.1, the version you installed with IBM FileNet Content Search Engine 4.5.
- On one of the installer panels, you will be asked to select your installed version of Autonomy K2 from a dropdown menu. Select K2 6.2.0 as the valid entry.

CAUTION Although the Locales installer offers the option of installing to a location other than the Verity install location, do not choose this option. For Content Search Engine, you must install the locales to the k2 directory of the Autonomy K2 installation path.

Windows

1. Access the Autonomy K2 Master Administration Server machine and log on as *k2_os_user*.
2. Copy the P8CSE450WIN.zip file from the installation package to the machine and extract the contents to a local directory.
3. Stop the Autonomy K2 services.
   a. Access Component Services.
   b. Stop the Verity K2 Administration Server service and the Verity K2 Administration Web Server service.
4. Create a directory in the following specific location (default drive shown):

   C:\Program Files\Common Files\InstallShield\Universal*WinVersion*\x86*host_name*\Gen1_vpddb

   a. Replace *WinVersion* with one of the following valid options:
      - Windows XP
      - Windows 2000
      - Windows 2003
      - Windows 2008

   b. Replace *host_name* with the name of the Master Administration Server machine.

      For example:

      C:\Program Files\Common Files\InstallShield\Universal\Windows 2003\x86\myMachine\Gen1_vpddb
5. Locate the vpd.script file in the K2 install directory and open it in a text editor. Make the following modifications in the file:
   - Replace the instance of *<K2InstallDir>* with the K2 install directory. For example:

     C:\Program Files\FileNet\ContentEngine\verity

- Replace the instance of <*myHostName*> with the Master Administration Server machine name.
- Replace the instance of <*WinVersion*> with the same option you chose in step 4 above.

6. Copy the vpd.script file to the directory location you created in step 4 above.
7. Navigate to the decompressed Locales installer location and execute the setupwin32.exe file.The installer will locate the K2 installation and start, based on the settings you completed above.

   NOTE The following license key is required:

   2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS

8. Start the Verity K2 services after the installation is complete.

UNIX

1. Access the Autonomy K2 Master Administration Server machine and log on as *k2_os_user*.
2. Copy the appropriate compressed file for your platform from the installation package to the machine and extract the contents to a local directory. For example:

   P8CSE450AIX.tar.gz

3. Stop the Autonomy K2 services.
   a. Access opt/verity/appserver/bin
   b. Use the following command, according to your environment:

      HP-UX

      `/verity_install_directory/k2/_hpux/bin/k2adminstop`

      AIX

      `/verity_install_directory/k2/_rs6k43/bin/k2adminstop`

      Solaris

      `/verity_install_directory/k2/_ssol26/bin/k2adminstop`

      Linux

      `/verity_install_directory/k2/_ilnx21/bin/k2adminstop`

4. Prepare your system with the following environment variable:

```
VERITY_CFG=/opt/verity/k2/common/verity.cfg
export VERITY_CFG
```

5. Navigate to the decompressed Locales installer location and execute the setup*UNIX_type*.bin. For example:

   setupaix.bin

   NOTE The following license key is required:

   2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS

6. Start the Verity K2 services after the installation is complete.

Task 4: Create a Content Search Engine collections directory

Use this procedure to configure the storage location and related security for your K2 collections.

Collections directory requirements

CAUTION Contrary to information outlined in the Autonomy-supplied documentation set, remote collections are not supported for use with IBM FileNet Content Search Engine. Collections must be written locally to the Autonomy K2 server. Using a remote-mounted disk that is accessed via the network (NFS, PCNFS, or CIFS) will cause stability problems under load and corrupt your collections. Any existing configurations that utilize non-local collections directories must be re-configured.

On a Windows system, the requirement to write Verity collections to local disks means that you are writing to a path such as D:\collections, and you cannot use a UNC path such as: *servername*\collections.

You can configure a second server that just reads these collections (note that you cannot have a second server that writes the collections). You must map a drive on the machine that reads collections to the file system on the machine that writes collections.

To do this (using the example above), share out the D: drive as some name other than D$, since you can't set permissions on D$. So, for example, set it as “DDrive”, and then map the D: drive on the reader box to *servername*\DDrive. Now D:\collections on the reader box references the same file system as the D: drive on the writer box.

For Windows machines, Autonomy K2 runs as a Windows service by default. Windows services do not mount mapped drives. One solutions is to run the Verity Administration service from a command line instead, so that it doesn’t run as a Windows service. To launch the service manually from a command line on a default installation, execute the following command:

```
C:\Program Files\verity\k2_61\k2\_nti40\bin\k2admin.exe" -cfg "C:\Program
Files\verity\k2_61\k2\common\verity.cfg
```

NOTE Only the Verity K2 6.1.1 Administration Server service needs to be started this way. The Verity Administration Web Server may be left as a Windows service.

Another solution is to use a tool like srvany.exe (supplied as part of the Windows Resource kit) to run a .cmd file that first maps the drives, and then issues the command above to start the k2admin.exe. The command to map drives, using the above example, is:

```
net use D: \\servername\DDrive
```

There are also third-party products available that function the same as the svrany.exe program.

The one drawback to using svrany is that although it will start the service correctly, it will not shut it down. You must use the Verity rcadmin command line tool to stop the service, or else use TaskManager to end all the processes that start with the prefix k2.

Relevant `rcadmin` commands from a command line:

- `rcadmin` to start the program
- `login` *`k2_os_user`* to log on

- `adminsignal` to initiate the shutdown and then one of the following in response to the `Type of Signal` then type one of the following:
  - `2 - Shutdown`
  - `3 - WS Refresh`
  - `4 - Restart all servers`

The service can now be restarted with the Windows Services tool.

To create a collections directory

For performance reasons, create one collections directory for each index area you create in IBM FileNet P8 Content Engine. Each collections directory you create must be set to provide proper security access as outlined below. The path to both the collections directory and collections temp directory must be entered in the index area properties when you create the index area.

Security and communication between Autonomy K2, Content Engine, and the collections directory is handled through the user accounts and permissions provided to those accounts. For more information on the accounts required, see For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.

For detailed information on security, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Security for integrated components and third-party products > Autonomy K2 Server > Security for Autonomy K2 Server.

1. Create a directory in which you will store collections (VerityIndexArea.RootDirectoryPath). This directory must be located on a disk that is local to the Verity server. Set permissions to allow access to *k2_os_user*.

   NOTE This path must be local to the index server that will be assigned to write collections.

2. Create a temp directory (VerityIndexArea.TempDirectoryPath) which will be used by the K2 Index Server and Content Engine Server during operations.

   NOTE This path must be visible to both the Content Engine and the Autonomy K2 servers. If the K2 Master Administration Server and Content Engine are not installed on the same machine, they both must be on a network-mounted file system.

3. Provide read/write access to the collections directory for Autonomy K2 by entering the full path to the location and record the temp collections path.

   Windows

   a. Open the following K2 configuration file in a text editor (default path shown):

      C:\Program Files\filenet\contentengine\verity\k2\common\verity.cfg

   b. Modify the next available alias settings by adding the collections path, where new collections will be written. For example, change alias6, mapping6, and dirmode6 to the following:

      ```
      alias6=path1
      mapping6=C:\Collections_Directory
      dirmode6=wr
      ```

UNIX

a. Open the following K2 configuration file in a text editor (default path shown):

   /opt/verity/k2/common/verity.cfg

b. Modify the next available alias settings by adding the collections path. For example, change alias6, mapping6, and dirmode6 to the following:

```
alias6=path1
mapping6=/Collections_Directory_Path
dirmode6=wr
```

   NOTE The *Collections_Directory_Path* must be a local path and not a mount point.

4. Set file store access. Each file store that will be full text indexed must be accessible by the Autonomy K2 server that will perform the full text indexing. Permissions on the file store must be set the same as the permissions on the collections directories, allowing both the Content Engine Operating System User and the K2 Operating System User to access them. The names of the file store directories must also be the same on each server that access it.

Task 5: Configure Content Engine for content-based retrieval

This task covers how to use Enterprise Manager to configure an index area and enable the content-based retrieval (CBR) feature provided by the IBM FileNet P8 Content Search Engine.

For more detail on content-based retrieval and Content Engine, see the IBM FileNet P8 help topic Configure CBR found at FileNet P8 Administration > Content Engine Administration > content-based retrieval > How to... > Configure CBR.

NOTES

- Before you complete the procedures in this topic, ensure the IBM FileNet P8 Content Search Engine has been installed and configured on at least one server (In effect, this means you have already installed and configured an Autonomy K2 Master Administration Server).
- Numerous K2 security accounts are referenced within this procedure. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.
- Various server names and related ports assigned during the installation and configuration of Autonomy K2 will be required during this procedure. If you do not have a record of the servers created and the ports that have been assigned, log on to the Autonomy K2 dashboard to obtain the information necessary.
- If you unimport the style set, the original files will be deleted from your system. In this scenario, if you wish to re-import the style set, you will need to recover it from your installation disk. In order to avoid this situation, you can either enter a unique name for the Style Set Alias during the initial Content Search Engine (Autonomy K2) installation, or make a backup copy of the original style set. If you entered a unique name for the style set during installation, ensure you use that Style Set Alias name for this procedure.
- Where machine name variables are required, IP addresses will not validate. In these cases, you must enter the host name for the machine.

To configure Content Engine for content-based retrieval

1. Launch Enterprise Manager and log on as the GCD Administrator.
2. Create a Verity Domain Configuration (VDC).
   a. Right-click **Enterprise Manager [domain]** in the Enterprise Manager tree and select Properties.
   b. Click the **Verity Domain Config.** tab.
   c. Enter the following K2 Master Administration Server access information:
      - Host Name - the name of the host of the K2 Master Administration Server.
      - Port - the K2 Master Administration Server port.
      - User Domain - the authentication domain in which your K2 services are installed.
      - User Group - K2 Security Group.

- Verity Username - K2 Security User.
- Password - the K2 Security User password.

d. Click Create Configuration to save your settings and create a Verity Domain Configuration object.

3. Assign a K2 Broker Server:

   a. Click the **Verity Server** tab.

   b. From the Brokers AVAILABLE pane, select the broker and click **Add** to move the server to the Brokers Selected pane.

   c. Click **OK**.

   NOTE You can assign multiple Broker Servers to a K2 Administration Server, primarily for failover. If one Broker Server goes down, then K2 can switch to another. In this configuration, you must ensure that all Search Servers required to access K2 Collections (index areas) are attached to each Broker Server. Be aware that a given Content Engine server will neither call multiple Broker Servers nor merge associated search results. See the appropriate steps of the *Configure Content Search Engine* task for instructions on how to create Broker Servers and assign Search Servers.

4. Enter a CBR Locale.

   NOTE If you click the **Set to Default** button, *uni* will be set as the CBR Locale. The *uni* locale is slower than the language-specific locales. The uni locale handles all languages but does not allow word-stemming. Therefore, a language-specific locale is recommended instead. Refer to the K2 documentation for a listing of K2-specific locales. Naming of K2 locales is specific to the K2 product and not the same as localization locales.

   a. Right-click the Object Store for which you want to enable indexing and select **Properties**.

   b. Click the **Locale** tab.

   c. Enter a valid Autonomy K2 locale for CBR.

5. Create an index area. Launch the Create an Index Area wizard:

   a. In the tree view, expand the **Object Stores** container.

   b. Right-click the object store to which you want add an index area and select **New > Index Area**.

   c. Work through the wizard screens. Refer to the *Installation and Upgrade Worksheet* in *Plan and Prepare your Environment for FileNet P8* for the parameter names and values you will specify when running the wizard. For additional details, click **Help** in the wizard screens.

   For more information on creating indexes, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > content-based retrieval > How to... > Create Verity index area.

6. Enable CBR for class definitions by activating the CBR Enable flag of the class you want available for CBR, as follows:

   a. Right-click the class you want to configure in the Enterprise Manager tree and click **Properties**.
   b. Select CBR Enabled and click **OK**.
7. Enable CBR for the class properties you want available for CBR, as follows:
   a. Right-click the class you want to configure and click **Properties**.
   b. Click the **Property Definitions** tab.
   c. Click the string property you want to enable for CBR indexing and click **Edit**.
   d. Check **CBR Enabled** and click **OK**.

To enable additional K2 Index Servers and Search Servers

If you add additional K2 Index Servers or Search Servers to an existing configuration, you must enable them through Enterprise Manager to utilize them.

1. Log on to Enterprise Manager as the GCD Administrator and expand the Enterprise Manager tree.
2. Open the Index Area folder.
3. Right-click the index area that you want to add the new services to and select **Properties**.
4. Enable the new Search Servers as follows:
   a. Click **Edit Search Servers**.
   b. In the Search Servers Available pane, highlight any servers you want to enable for this index area and click **Add** to add the server to the Search Servers Selected list.
   c. Click **OK** to save the settings and enable the new servers.
5. Enable the new Index Servers as follows:
   a. Click **Edit Index Servers**.
   b. In the Index Servers Available pane, Highlight any servers you want to enable for this index area and click **Add** to add the server to the Index Servers Selected list.
   c. Click **OK** to save the settings and enable the new servers.

Task 6: Verify the Content Search Engine installation

To verify that the Content Search Engine installation and configuration succeeded, perform the following procedure.

To verify the Content Search Engine installation

1. In Enterprise Manager, access the Root Folder and open a sub folder that contains a document.
2. Right-click a document and select **All Tasks > Index for Content Search**.

   A Notice should indicate that a content search index Job was created successfully.
3. Right-click the object store node for the document you set to index for content search and select **All Tasks > Index Jobs Management**.
4. In the Index Jobs Manager, check to see that the Job Status indicates the job has terminated normally.

   NOTE Depending on the size of the job, it may take a few minutes for the indexing job to complete. If the Job Status indicates that the indexing is in progress, wait a few minutes to give the job time to complete.

Install and configure Process Engine

This section contains the following major topics:

Task 1: Install Process Engine Interactively

Interactively install Process Engine software by performing one of the following procedures.

You will find references to logging on as the root and fnsw users within the following procedures. For all UNIX operating systems, the root user must run in the Bourne or Korn shell and the fnsw user must run in the Korn shell.

Before starting the Process Engine installation, verify that you have completed the steps in the *To configure your /etc/hosts file* section in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

To verify the database connection

Verify the database connection between the Process Engine and the database. See "Verify the ability to connect to the database" on page 125 for details.

To install the Process Engine software interactively (UNIX)

1. Log on to the server as the root user.

   NOTE This user does not need to be a database administrator unless you will be executing the SQL scripts from the Process Engine installation program.

2. Access the Process Engine software package.
3. From the console, launch the appropriate P8PE-4.5.0-*platform*.bin installation program.
4. Wait for files to finish unpacking.
5. Complete the Process Engine installation screens using the appropriate information from your configuration worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   a. Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select PE installer.

   b. Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

6. Monitor /fnsw/local/logs/wizard to check the progress of the installation since the installation program will run for several minutes,and its progress, though displayed, might not visibly advance for an extended period of time.

7. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	/fnsw/local/logs/PE (if the install completes successfully) or /fnsw/tmp_installer (if the install has errors)
IS mini-installer logs	/fnsw/local/logs/wizard and other subdirectories under /fnsw/local/logs

8. Log off as the root user and log on as fnsw (or the alias).
9. Proceed to "To reset administrative user passwords" on page 89.

To install the Process Engine software interactively (Windows)

The Process Engine installation occurs in two parts that are separated by a restart of the system. After the restart, the second part of the installation continues automatically.

1. Log on as a member of the local Administrators group or a user with equivalent permissions. If you plan to run the SQL scripts from the Process Engine installation program, the user you log on as must also be a database administrator. See "Specify IBM FileNet P8 accounts" on page 65 for information on requirements for logging on as a Windows domain user for Process Engine installation.

   NOTE This user does not need to be a database administrator unless you will be executing the SQL scripts from the Process Engine installation program.

2. Access the Process Engine software package, and start the P8PE-4.5.0-Win.exe installation program.

   NOTE To run the Process Engine installation from disk, you must copy the installation files to a disk volume where 8.3 name generation is enabled, or if 8.3 name generation is disabled, you must copy the installation to a path the uses only short (8.3) names.

   When running from disk, either interactively or silently, be aware that the Process Engine installation program has a 64-character path limitation when the path is expressed in 8.3 format. This limitation applies to the IMSInst subdirectory where the underlying Image Services (IS) mini-installer setup.exe is located. For example, the original path where the IS mini-installer resides is:

   \\server08\Software\InstallationDisks\FileNet\Release P8 4.5.0\ProcessEngine\Windows\IMSInst

   When expressed in 8.3 format the path might be:

   \\server08\Software\INSTAL~1\FileNet\RELEAS~1.0\PROCES~1\Windows\IMSInst

   This compressed path is 73 characters long, exceeding the 64-character limit.

3. Complete the Process Engine installation screens using the appropriate information from your configuration worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   a. Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select PE installer.

   b. Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

4. The installation program will run a number of steps but to complete the installation the computer must be restarted. When prompted, select **Yes, restart my computer** and click **Finish**.

5. When the system restarts, log on using the same account you used in Step 1. After you log on, Process Engine installation will continue. Click **Next** to continue the installation.

   If installation fails at this point, correct the errors that caused the failure and run the post-boot installation program by navigating to the Process Engine\IMSInst install folder and executing:

   ```
   setup.exe  -postboot
   ```

   NOTE Not all error conditions can be resolved in this way. It might be necessary to uninstall the Process Engine software and re-run the installation program.

6. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE450_setup.log C:\Program Files\FileNet\PE C:\FNSW
IS mini-installer logs	*Windir*\mini_installer.log, Windows Event logs, and log files under \FNSW_LOC\logs
Output files from SQL Script validation (only if the SQL scripts were executed from the Process Engine installation program).	C:\Program Files\FileNet\PE

To reset administrative user passwords

Process Engine installation automatically creates several internally required local users. For Oracle and SQL Server databases, if default users and passwords were configured, Process Engine installation also creates several database users. Because these users are created with default passwords, it is a best practice to reset the passwords for these users. The following table

lists the users created, the level of system access each user has, and the tool used to change the password. See the IBM FileNet Image Services documentation for instructions on using the referenced tools

User Name	User Type	Description	How to modify
f_sw	Database (Oracle. SQL) OS user for DB2	Has DBA privileges. Used as the Process Engine run-time user.	Execute Xdbconnect. See steps below.
f_maint	Database (Oracle, SQL) OS user for DB2	Has DBA privileges. Used for RDBMS maintenance.	Execute Xdbconnect. See steps below.
fnsw or alias	Operating System	Primary Process Engine user. Used to execute Process Engine software and services.	**Windows** After you change the password for the fnsw user, you must also use the Windows Services tool to update the Log On tab for the IMS ControlService and the Process Engine Services Manager because the fnsw user is used to start these services.**NOTE** Process Engine activity will cease if the fnsw user's password expires.
SysAdmin	SEC (internal Process Engine security software)	Primary administrator user for IBM FileNet software tools.	Execute Xapex -> Security Administration. Log on as SysAdmin.
FieldService	SEC	Used internally by Process Engine software.	
Operator	SEC	Used internally by Process Engine software.	

To verify the connection to the Process Engine database

Verify the ability to connect to the Process Engine database using the run-time user name and password.

At a command prompt, execute:

```
vwcemp -l
```

If the connection is successful, you will receive a message indicating that the Content Engine has not been configured yet. Because there are no security settings configured, the above command should return a message similar to "....no object service is configure...."

Unsuccessful connection to the database will return an exception.

If the connection is not successful, there could be a mismatch between the password entered for the f_sw user during Process Engine installation and the password created by execution of the SQL scripts (Oracle) or for the operating system runtime user (DB2).

To determine if a failure to connect to the database is due to a password mismatch, use the steps in "To set the f_maint and f_sw passwords" on page 91.

To set the f_maint and f_sw passwords

For added security, Process Engine stores an encrypted version of the passwords for the f_sw and f_maint users, or their aliases, in a file called rdbup.bin. This is in addition to passwords for these users in the Oracle or SQL Server database, or the DB2 operating system user's password. The encrypted password and the database or operating system user's passwords must match.

To verify that the passwords match, use the following procedure to start the Xdbconnect utility. Xdbconnect works only if the passwords in the encrypted file and the database match.

Use the following procedure to change the passwords for the f_maint and f_sw users after installing the Process Engine software. For Oracle and SQL Server databases, both the encrypted file and the database passwords will be updated. For DB2, only the encrypted file will be updated.

1. Start the Database Server Connect application by executing the following:

   ```
   Xdbconnect -r
   ```

2. Log on as SysAdmin.
3. Change the primary password for the users f_sw and f_maint (or their alias) to match the database password (Oracle and SQL Server) or operating system user's password (DB2).
4. Exit the application.

To re-enable Oracle Password Complexity Verification

If, as directed earlier, you disabled the the Oracle Password Complexity Verification feature prior to installing Process Engine, you can now re-enable it.

(Optional: Oracle databases only) To remove fnsw user from the Oracle database administrators group (UNIX)

Process Engine installation creates a user that is no longer required after installation is complete. Remove the fnsw (or the alias) user from the *Oracle Database Administrators* group.

(Optional; Oracle database only) To remove fnsw and oracle users from the ORA_DBA group (Windows)

Remove the fnsw (or its alias) and oracle users from the ORA_DBA group. Process Engine installation creates these users, which are no longer required after installation is complete.

To restore any custom modifications for root and fnsw users (UNIX)

Process Engine installation creates a new versions of a number of files. If the previous versions of these files contained any custom settings, edit the new files for the fnsw and root users accordingly. Saved files are in *.filename*.old.*nn*, where *nn* is a sequential number. The latest saved version will be in the highest numbered file. The following files are modified by the Process Engine installation program:

.Xdefaults

.Xresources

.dbxinit

.dtprofile

.env

.login

.mwmrc

.xinitrc

.profile

.cshrc

To verify the /etc/services file settings (UNIX)

1. Log on as the root user and check the /etc/services file to verify the following parameters:

   ALL UNIX

   ```
   tms          32768/tcp
   cor          32769/tcp
   nch          32770/udp
   fn_trapd     35225/udp
   ```

   AIX only

   ```
   smux         199/tcp                # snmpd smux port
   ```

   Solaris only

   ```
   fn_snmpd  161/udp
   ```

   HP only

   ```
   snmp        161/udp  snmpd     # Simple Network Management Protocol Agent
   snmp-trap   162/udp  trapd     # Simple Network Management Protocol Traps
   ```

2. If necessary, add the parameters to the file and save the changes.

To verify the /etc/services file settings (Windows)

1. Log on as the Administrator user and check the \Windows\system32\drivers\etc\services file to verify the following parameters:

   tms 32768/tcp
   cor 32769/tcp
   nch 32770/udp
   fn_snmpd 161/udp
   fn_trapd 35225/udp

2. If necessary, add the parameters to the file and save the changes.

To modify the integrity level for executables (Windows 2008)

Modify the mandatory integrity level of the Image Services executables to High. Complete the following procedure:

1. Open a command prompt with Administrative privileges.
2. Change the current directory to *installLocation*\fnsw\bin.
3. Execute the following command:

```
icacls  *.exe  /setintegritylevel  H
```

To redirect log messages to the Image Services error log (Windows)

Enable the redirection of log messages to the Image Services error log. This redirection will log message to the Image Services error log as well as to the default Windows Event Log. By enabling this redirection, you will be able to monitor the progress of the database object upgrade in a command window.

To enable the redirection, change the LogToFiles value from 0 to 1 for the following registry key.

Windows 2003

```
HKEY_LOCAL_MACHINE>SOFTWARE>FileNET>IMS>CurrentVersion
```

Windows 2008

```
HKEY_LOCAL_MACHINE\SOFTWARE\WoW6432Node\FileNet\IMS\CurrentVersion
```

To clean up before starting Process Engine (UNIX)

1. Log on as fnsw.
2. Execute:

   ```
   killfnsw -DAyS
   ```

3. Execute:

   ```
   ipcs -a
   ```

4. Verify there is no entry with 0x464 pattern. If there are any entries with this pattern, use `ipcrm` to remove them.

To edit the /etc/inittab file (UNIX)

By default, the Process Engine software starts automatically when you restart the server and needs its database started beforehand. If the database is not automatically started on server restart, edit the /etc/inittab file on the Process Engine machine to comment out the autostart of Process Engine. The following are examples of changes to make.

AIX

```
rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

```
#rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

HP

```
rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

#rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1

Solaris

```
fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

to

```
#fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

To edit the pe_start file (HP)

If the value for the maxdsiz kernel parameter is > 1GB, edit the pe_start file.

Change:

```
nohup /usr/ccs/lbin/dldd32 2>&1 >/dev/null
```

to

```
nohup /usr/ccs/lbin/dldd32 +a 0x70000000 2>&1 >/dev/null
```

Proceed to “Install Process Engine software updates” on page 102.

Task 2: Install Process Engine Silently

Silently install Process Engine software by performing one of the following procedures.

Before starting the Process Engine installation, verify that you have completed the steps in the *To configure your /etc/hosts file* section in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

To verify the database connection

Verify the database connection between the Process Engine and the database. See "Verify the ability to connect to the database" on page 125 for details.

To install the Process Engine software silently (UNIX)

Complete the following procedures to silently install Process Engine.

1. Access the Process Engine software package, and copy the contents to a local temporary directory on the local disk.
2. Edit the PE_silent_install.txt file to reflect the appropriate responses for your installation. See "Encrypt passwords" on page 453 for information on use of the password encryption tool.
3. Save the edited response file to your temporary directory.
4. Log on as the root user in the Korn shell.
5. Navigate to the temporary directory on the local disk.
6. Open a command prompt and execute:

   `P8PE-4.5.0-`*platform*`.bin -silent -options PE_silent_install.txt`
7. Monitor /fnsw/local/logs/wizard to check the progress of the installation since the installation program will run for several minutes,and its progress, though displayed, might not visibly advance for an extended period of time.
8. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	/fnsw/local/logs/PE (if the install completes successfully) or /fnsw/tmp_installer (if the install has errors)
IS mini-installer logs	/fnsw/local/logs/wizard and other subdirectories under /fnsw/local/logs

9. Log off as the root user and log on as fnsw (or the alias).

10. Proceed to "To reset administrative user passwords" on page 97.

To install the Process Engine software silently (Windows)

1. Access the Process Engine software package, and copy its contents to a temporary directory on the local disk.

   The installation program will run a number of steps but to complete the installation the computer must be restarted.

2. Edit the PE_silent_install.txt file to reflect the appropriate responses for your installation. See "Encrypt passwords" on page 453 for information on use of the password encryption tool.

3. Save the edited response file to your temporary directory.

4. 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. See "Specify IBM FileNet P8 accounts" on page 65 for information on requirements for logging on as a Windows domain user for Process Engine installation.

5. Open a command prompt and navigate to the temporary directory. Execute:

   ```
   P8PE-4.5.0-Win.exe -silent -options PE_silent_install.txt
   ```

   The Process Engine installation program will reboot the server, after which it will continue with the post-boot operations.

6. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE450_setup.log C:\Program Files\FileNet\PE C:\FNSW
IS mini-installer logs	*Windir*\mini_installer.log, Windows Event logs, and log files under \FNSW_LOC\logs
Output files from SQL Script validation (only if the SQL scripts were executed during Process Engine installation).	C:\Program Files\FileNet\PE

7. Proceed to "To reset administrative user passwords" on page 97.

To reset administrative user passwords

Process Engine installation automatically creates several internally required local users. For Oracle and SQL Server databases, if default users and passwords were configured, Process Engine installation also creates several database users. Because these users are created with default passwords, it is a best practice to reset the passwords for these users. The following table lists the users created, the level of system access each user has, and the tool used to change the password. See the IBM FileNet Image Services documentation for instructions on using the referenced tools

User Name	User Type	Description	How to modify
f_sw	Database (Oracle. SQL) OS user for DB2	Has DBA privileges. Used as the Process Engine run-time user.	Execute Xdbconnect. See steps below.
f_maint	Database (Oracle, SQL) OS user for DB2	Has DBA privileges. Used for RDBMS maintenance.	Execute Xdbconnect. See steps below.
fnsw or alias	Operating System	Primary Process Engine user. Used to execute Process Engine software and services.	**Windows** After you change the password for the fnsw user, you must also use the Windows Services tool to update the Log On tab for the IMS ControlService and the Process Engine Services Manager because the fnsw user is used to start these services.**NOTE** Process Engine activity will cease if the fnsw user's password expires.
SysAdmin	SEC (internal Process Engine security software)	Primary administrator user for IBM FileNet software tools.	Execute Xapex -> Security Administration. Log on as SysAdmin.
FieldService	SEC	Used internally by Process Engine software.	
Operator	SEC	Used internally by Process Engine software.	

To verify the connection to the Process Engine database

Verify the ability to connect to the Process Engine database using the run-time user name and password.

At a command prompt, execute:

vwcemp -l

If the connection is successful, you will receive a message indicating that the Content Engine has not been configured yet. Because there are no security settings configured, the above command should return a message similar to "….no object service is configure…."

Unsuccessful connection to the database will return an exception.

If the connection is not successful, there could be a mismatch between the password entered for the f_sw user during Process Engine installation and the password created by execution of the SQL scripts (Oracle) or for the operating system runtime user (DB2).

To determine if a failure to connect to the database is due to a password mismatch, use the steps in "To set the f_maint and f_sw passwords" on page 98.

To set the f_maint and f_sw passwords

For added security, Process Engine stores an encrypted version of the passwords for the f_sw and f_maint users, or their aliases, in a file called rdbup.bin. This is in addition to passwords for these users in the Oracle or SQL Server database, or the DB2 operating system user's password. The encrypted password and the database or operating system user's passwords must match.

To verify that the passwords match, use the following procedure to start the Xdbconnect utility. Xdbconnect works only if the passwords in the encrypted file and the database match.

Use the following procedure to change the passwords for the f_maint and f_sw users after installing the Process Engine software. For Oracle and SQL Server databases, both the encrypted file and the database passwords will be updated. For DB2, only the encrypted file will be updated.

1. Start the Database Server Connect application by executing the following:

   ```
   Xdbconnect -r
   ```

2. Log on as SysAdmin.
3. Change the primary password for the users f_sw and f_maint (or their alias) to match the database password (Oracle and SQL Server) or operating system user's password (DB2).
4. Exit the application.

To re-enable Oracle Password Complexity Verification

If, as directed earlier, you disabled the the Oracle Password Complexity Verification feature prior to installing Process Engine, you can now re-enable it.

(Optional: Oracle databases only) To remove fnsw user from the Oracle database administrators group (UNIX)

Process Engine installation creates a user that is no longer required after installation is complete. Remove the fnsw (or the alias) user from the *Oracle Database Administrators* group.

(Optional; Oracle database only) To remove fnsw and oracle users from the ORA_DBA group (Windows)

Remove the fnsw (or its alias) and oracle users from the ORA_DBA group. Process Engine installation creates these users, which are no longer required after installation is complete.

To restore any custom modifications for root and fnsw users (UNIX)

Process Engine installation creates a new versions of a number of files. If the previous versions of these files contained any custom settings, edit the new files for the fnsw and root users accordingly. Saved files are in *.filename*.old.*nn*, where *nn* is a sequential number. The latest saved version will be in the highest numbered file. The following files are modified by the Process Engine installation program:

.Xdefaults

.Xresources

.dbxinit

.dtprofile

.env

.login

.mwmrc

.xinitrc

.profile

.cshrc

To verify the /etc/services file settings (UNIX)

1. Log on as the root user and check the /etc/services file to verify the following parameters:

   ALL UNIX

   ```
   tms          32768/tcp
   cor          32769/tcp
   nch          32770/udp
   fn_trapd     35225/udp
   ```

   AIX only

   ```
   smux         199/tcp                # snmpd smux port
   ```

Solaris only

```
fn_snmpd  161/udp
```

HP only

```
snmp        161/udp  snmpd     # Simple Network Management Protocol Agent
snmp-trap   162/udp  trapd     # Simple Network Management Protocol Traps
```

2. If necessary, add the parameters to the file and save the changes.

To verify the /etc/services file settings (Windows)

1. Log on as the Administrator user and check the \Windows\system32\drivers\etc\services file to verify the following parameters:

```
tms          32768/tcp
cor          32769/tcp
nch          32770/udp
fn_snmpd     161/udp
fn_trapd     35225/udp
```

2. If necessary, add the parameters to the file and save the changes.

To modify the integrity level for executables (Windows 2008)

Modify the mandatory integrity level of the Image Services executables to High. Complete the following procedure:

1. Open a command prompt with Administrative privileges.
2. Change the current directory to *installLocation*\fnsw\bin.
3. Execute the following command:

```
icacls  *.exe  /setintegritylevel  H
```

To redirect log messages to the Image Services error log (Windows)

Enable the redirection of log messages to the Image Services error log. This redirection will log message to the Image Services error log as well as to the default Windows Event Log. By enabling this redirection, you will be able to monitor the progress of the database object upgrade in a command window.

To enable the redirection, change the LogToFiles value from 0 to 1 for the following registry key.

Windows 2003

```
HKEY_LOCAL_MACHINE>SOFTWARE>FileNET>IMS>CurrentVersion
```

Windows 2008

```
HKEY_LOCAL_MACHINE\SOFTWARE\WoW6432Node\FileNet\IMS\CurrentVersion
```

To clean up before starting Process Engine (UNIX)

1. Log on as fnsw.
2. Execute:

   ```
   killfnsw -DAyS
   ```

3. Execute:

   ```
   ipcs -a
   ```

4. Verify there is no entry with 0x464 pattern. If there are any entries with this pattern, use `ipcrm` to remove them.

To edit the /etc/inittab file (UNIX)

By default, the Process Engine software starts automatically when you restart the server and needs its database started beforehand. If the database is not automatically started on server restart, edit the /etc/inittab file on the Process Engine machine to comment out the autostart of Process Engine. The following are examples of changes to make.

AIX

```
rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

```
#rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

HP

```
rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

#rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1

Solaris

```
fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

to

```
#fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

To edit the pe_start file (HP)

If the value for the maxdsiz kernel parameter is > 1GB, edit the pe_start file.

Change:

```
nohup /usr/ccs/lbin/dldd32 2>&1 >/dev/null
```

to

```
nohup /usr/ccs/lbin/dldd32 +a 0x70000000 2>&1 >/dev/null
```

Proceed to “Install Process Engine software updates” on page 102.

Task 3: Install Process Engine software updates

Perform the procedure in this topic for the Process Engine to install software updates (mod release, fix pack, or interim fix).

If there are no Process Engine software updates to be installed, skip to "Install the latest Content Engine Client files on Process Engine servers" on page 103.

To install the Process Engine software updates

1. To obtain the latest Process Engine software updates, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the Process Engine software updates and perform the installation procedures in the readmes on the Process Engine.
3. Proceed to "Install the latest Content Engine Client files on Process Engine servers" on page 103.

Task 4: Install the latest Content Engine Client files on Process Engine servers

To install the 4.5 release or fix pack version Content Engine Client files, perform the following procedures on all Process Engine servers. Except where noted, these steps apply to WebSphere, WebLogic, and JBoss.

To install the Content Engine Client files

Refer to the appropriate information from your installation worksheet while performing the following steps.

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. On the machine where Process Engine is installed, log on as fnsw, with these permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Process Engine is installed
   - Execute permission on the Content Engine Client install software
3. Verify that there is a current backup of Process Engine.
4. Copy the Content Engine Client install software from the Content Engine installation software to the temporary directory, and unzip the software package. The version of the installation software must match the version of Content Engine.
5. Start the Content Engine client installation process.
   - To install the Content Engine client interactively:
     i. Access the IBM FileNet Content Engine client update software.
     ii. Run one of the commands in the table below, *CE_version* is the version of Content Engine you intend to run, for example, 4.5.0..

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN`

     iii. Complete the installation program wizard.

- To install the Content Engine client files silently:

  i. Make a back up copy of the input file listed below for your operating system:

     Windows

     `CEClient_silent_install_win.txt`

     UNIX

     `CEClient_silent_install_unix.txt`

  ii. Open the silent input file in a text editor. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your update.

  iii. Navigate to the path containing the Content Engine Client installation program, and run one of the commands in the following table to perform the silent installation, where:

     CE_version is the version of Content Engine you intend to run, for example, 4.5.0.

     path is the path that contains the installation program.

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN -f `*`path`*`/CECLIENT.AIX/ CEClient_silent_install.txt -i silent`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN -f `*`path`*`/CEClient.HPUX/ CEClient_silent_install.txt -i silent`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.EXE -f `*`path`*`/CEClient.HPUXI/ CEClient_silent_install.txt -i silent`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN -f `*`path`*`/CEClient.Linux/ CEClient_silent_install.txt -i silent`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN -f `*`path`*`/CEClient.Solaris/ CEClient_silent_install.txt -i silent`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE -f `*`path`*`\CEClient.Windows\ CEClient_silent_install.txt -i silent`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN -f `*`path`*`/CEClient.zLinux/ CEClient_silent_install.txt -i silent`

Task 5: Configure Process Task Manager

Complete this task to start the Process Task Manager and set initial configuration parameters.

To start the Process Task manager and software on the Process Engine

1. Log onto Process Engine as a member of fnadmin or the fnsw alias you assigned during Process Engine installation.

   On a Windows machine, Process Engine installation either automatically creates the fnadmin group and the fnsw user, and adds the fnsw user to the fnadmin group or aliases for the user and group were designated during Process Engine installation.

   On a UNIX machine, this user and group were manually created before running Process Engine installation.

2. Start the Process Task Manager, as follows:

   Windows Select Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager.

   UNIX Enter `vwtaskman` at the command prompt. The terminal should support X Windows and the DISPLAY environment variable should be set.

3. Right-click your Process Engine server in the feature pane.

4. If the Process Engine software is not already running, start it by choosing Start from the Action menu.

5. Select the Process Engine in the feature pane and click the **Security** tab to configure the Security settings.

   Provide the service username, Content Engine URI, Administrator Group, and Configuration Group identified on your installation worksheet. See the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Configure Process Engine > Security for details on the user and groups.

   NOTE The service username should be entered as a short name.

6. Click the **General** tab and make sure the language settings as well as the rest are all correct. Click **Apply** if any changes had been made.

   After all parameters have been entered, click **Apply** and restart the Process Service when prompted. If errors are returned, additional information is available in:

   Windows \fnsw_loc\logs\TM_daemon\PEDirectoryServerConnectionDebug.txt

   UNIX /fnsw/local/logs/TM_daemon/PEDirectoryServerConnectionDebug.txt

7. Right-click on the **Regions** folder, select **New** to create a new region.

8. Select the **General** subtab and fill in the region number.

9. Indicate whether region recovery will be enabled for this region. If you will be using region recovery:
   a. Enable region backups.
   b. Identify the tablespaces or filegroups to be used by the region.
10. Click the **Security** subtab to set a region password.

    NOTE The password must match the password that is entered when creating a Process Engine Region in "Create a Process Engine Isolated Region" on page 153.
11. After all parameters have been entered, click **OK** and restart the Process Service when prompted.

To verify connection to the Process Engine database

Check the database connection from Process Engine by executing the following command. This will also verify that security settings were saved to the database. Execute:

```
vwcemp -l
```

If there are no security settings configured, the above command should return a message similar to "….no object service is configure…."

If the security settings were saved, the above command will return the information you entered in the Process Engine Task Manager.

Proceed to "Complete Post-Install Process Engine Configuration (Windows Only)" on page 107.

Task 6: Complete Post-Install Process Engine Configuration (Windows Only)

Use the procedure in this task to enable Process Engine to use the largest available contiguous free memory area for shared memory allocations. If you fail to perform this procedure, the system will not allocate shared memory at some point during normal execution and will cease to function correctly.

In the following steps. you will use vwtool to get the address for the largest free memory block and use that address to create a new registry key. You will then verify that address by running ipc_tool.

To configure contiguous free memory for Process Engine (Windows only)

1. Start vwtool at a command prompt. Log on using the Service Username you provided when completing the steps in "Configure Process Task Manager" on page 105.
2. Use the processmap command to find the largest contiguous free memory area, as in:

   ```
   <vwtool:1>processmap
   ```

   vwtool returns the following:

   ```
   Process Id (CR=this vwtool process):
   ```

   Press **Return** (CR) to get the process map for this process, as in the following example, where the process ID is 2592:

   ```
   C:\FNSW\BIN\vwtool.exe (ID:2592)

   Address    Attrib         Size    Owner
   =======    ======           ====     =====

   00000000   Free          65536
   00010000   Private       12288
   00013000   Free          53248
   00020000    Private         4096
   ```

   …….(pages of memory addresses omitted here)

   ```
   7FFDE000   Private        4096
   7FFDF000   Private        4096
   7FFE0000    Private        65536

   C:\FNSW\BIN\vwtool.exe (ID:2592)

   Largest FREE block found : 453873664 bytes at address 0x4B577000
   Rounded up to a 64K boundary, free block address 0x4B580000
   ```

   In this example, `0x4B580000` is the address you will need in the next step. In some cases you might see only the line referencing the largest free block because the value is already at a 64K boundary.

3. Start regedit from the Windows > Start > Run command and perform the following steps to create a DWORD value for IS StartShmAddress, using the address noted in step 2.
   a. Navigate to the following regedit key:

      `HKEY_Local_Machine\Software\FileNet\IMS\CurrentVersion\`

   b. Create a new DWORD value named:.

      `StartShmAddress`

   c. Enter or verify the following in the Edit DWORD Value Screen:

      Value name = `StartShmAddress`

      Value data = *<address of largest free memory block>*

      From the example above the value will be 4B580000.

      Base is hexadecimal.

   d. Click **OK.**
   e. Exit from regedit.
4. Restart the Process Engine software.
5. Verify the setting you just applied for the shared memory address by executing the following at a command prompt:

   ```
   ipc_tool -A
   ```

   The following is an example of the information that is returned.

   ```
   Image Services software shared memory segment limit: 129 segments
   Current configured segment size: 0x01000000 bytes (16 MB)
   Before allocating shared memory for Image Services, the SysV library
   performs a test to determine the system shared memory limit.  This test
   can be used as a reference for performance tuning.  The test results vary
   depending on the amount of memory in use by other processes.  The actual
   amount of shared memory available during operation may be less.  The test
   results are:

      Successfully attached to 27 segments
      Successfully obtained 432 MB of shared memory

   The following table displays the number of shared memory segments currently
   in use by Image Services.  Segment #0 (called the address manager) is small.
   The other segment(s) contain the actual Image Services data.  Note that
   running ipc_tool will force the creation of segments #0 and #1 even when no
   other Image Services process is up.
   Shared Memory Address Manager Information

       Address     Shm id          Creator
   ```

```
Enter <space> to continue, 'q' to quit:
  0 0x4b580000  FNSHM_464d0000  Shared address manager
  1 0x4c580000  FNSHM_464a0000  FileNet server software
Total Image Services shared memory allocated: 16 MB
(This does not include segment #0)
```

NOTE The First shared memory address above 0x4B580000 is the value you would check for this example.

6. Exit ipc_tool. If the shared memory address is correct, proceed to the next installation task. If the value is not correct, verify Step 1 - Step 4 above before proceeding.

Install and configure Application Engine

To set up Application Engine, do the following:

1. Install Application Engine. Do Task 1 on page 111.
2. Install Application Engine Software Updates. Do Task 2 on page 116.
3. Install Content Engine Client file updates. Do Task 3 on page 117.
4. Install Process Engine Client file updates. Do Task 4 on page 120.
5. Configure Application Engine. Do one of the following:
   - Task 5a on page 122 (WebSphere)
   - Task 5b on page 132 (WebLogic)
   - Task 5c on page 138 (JBoss)
6. Deploy Application Engine. Do one of the following:
   - Task 6a on page 141 (WebSphere)
   - Task 6b on page 143 (WebLogic)
   - Task 6c on page 145 (JBoss)

Task 1: Install Application Engine

This topic includes Application Engine installation instructions for all supported application servers, for UNIX and Windows platforms.

NOTES

- If you plan to install and use the IBM FileNet Workplace XT product, installing Application Engine is not required.
- Before installing Application Engine, check the latest version of the IBM *FileNet P8 4.5 Release Notes* for known issues that might impact this software installation. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- Make sure your installation location meets the requirements specific for Application Engine outlined 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.
- If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch.

  Otherwise, if you deploy Application Engine initially as a WAR file and later decide to redeploy as an EAR file, you will have to uninstall Application Engine and then reinstall the application, selecting EAR file deployment, to add the required files to your setup.
- (Highly Available installations) To install Application Engine in a web farm or clustered environment, follow the instructions in the*IBM FileNet P8 Platform High Availability Technical Notice.* The document outlines the required HA install procedure and references this guide for detailed installation and deployment instructions. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- To ensure proper functionality and performance only install one instance of Application Engine per application server (or virtual machine or WebSphere LPAR). You can, however, deploy multiple instances of a single Application Engine version per application server. For details, see "Deploy multiple Application Engine instances" on page 181.
- Before logging on to Workplace for the first time, at least one object store must exist on the Content Engine to hold the site preferences. See "Create the initial object store" on page 60 for more information.
- If you run the installer to upgrade Application Engine, the installer verifies that the currently installed version of Application Engine can be upgraded. See "Prepare for Application Engine upgrade" on page 332 for more information.

To install the Application Engine software

1. Log on to the application server, as appropriate for your operating system:

   UNIX

   Log on as a user with write access to the /bin directory and read, write, and execute access to the directory where you plan to install Application Engine.

   NOTE Although the installing user must have write access to the /bin directory, the Application Engine installer does not write to that directory.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

2. Start the installation process.

   - To install Application Engine interactively:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software.

     ii. Launch the appropriate Setup program (P8AE-4.0.2.0-*operating_system*.bin/.exe) and continue with Step 3 on page 113 below.

   - To install Application Engine silently:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software package, and copy the appropriate AE_silent_input.txt file to a local directory.

     ii. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your installation.

     CAUTION If you are modifying the silent input file to perform an upgrade from AE 3.5.x to AE 4.0.2 you must modify all instances of *AE_install_path* in the script as follows:

     UNIX

     Change `../FileNet/AE` to `../FileNet`

     Windows

     Change `..\FileNet\AE` to `..\FileNet`

     Change `..\\Filenet\\AE` to `..\\FileNet`

     iii. From a command prompt, navigate to, and execute the installer:

     UNIX

     ```
     ./P8AE-4.0.2.0-operating system.bin -options path_to_edited_input_file/
     AE_silent_input.txt -silent
     ```

     Windows

     ```
     P8AE-4.0.2.0-Win.exe -options path_to_edited_input_file\AE_silent_input.txt -
     silent
     ```

NOTE When installing Application Engine remotely on UNIX, run the installer with the standard input stream redirected from `/dev/console`, for example:

```
rsh remote-machine -n P8AE-4.0.2.0-AIX.bin -options AESilentScriptUNIX.txt -
silent < /dev/console
```

If you do not add the redirect, the silent intaller will fail with a "process not attached to terminal" message and the usage message for the "who" command.

iv. Continue with Step 4.

3. Complete the Application Engine Installer screens. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

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 the Application Engine installation program:

- Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select AE installer.
- Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

In addition to the prompts for system-specific information, you will need to consider the following decisions:

- If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the /FileNet/AE part of the path. This affects the following additional paths:
  - Configuration path: Retain the /FileNet/Config/AE part of the path.

    The configuration files for an EAR file deployment, a web farm, or a clustered environment must be located in a shared folder that is accessible by all copies of the Workplace application. For more information, see the *IBM FileNet P8 Platform High Availability Technical Notice*. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
  - Log files: Retain the /FileNet/Logs part of the path.
  - User tokens: Retain the /FileNet/Authentication part of the path.
- You can choose to deploy Application as a WAR file or an EAR file. If you decide to deploy Application Engine as a WAR file and later decide to redeploy as an EAR file you must uninstall Application Engine and then reinstall the application, selecting EAR file deployment. For information, see "If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch." on page 111.
- You can choose whether to use container-managed or application-managed authentication. Choose container-managed if you plan to use SSO.

When you select Container-Managed Authentication, Setup 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 Setup is finished.

- Additional installation notes:
  - To change the Content Engine name later, or to connect to a different Content Engine, edit the WcmApiConfig.properties file. For information, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs.
  - For information on how to reconfigure the Documentation URL after installation is completed, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > Bootstrap properties.
  - The UTCryptoKeyFile.properties file contains the user token cryptography key used by IBM FileNet P8 applications to launch into each other without the need for additional login.

    CAUTION For multiple applications to pass user tokens to each other, each participating application **must** use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all servers that are hosting a token-sharing application.

    For information, see the IBM FileNet P8 Developer Help topic Developer Help > Workplace Integration and Customization Introduction > User Tokens > Configuring Applications to Use Tokens.

4. View the app_engine_install_log_4_0_2.txt file located, located in *AE_install_path*/Logs.

   Verify that no errors or failures were logged. Correct any errors before you proceed.

5. (Solaris only) Set the anon ports.

   To use the IBM FileNet ports listed below for Component Manager on Solaris-based systems, you must first enable the ports by setting the smallest anon port to 32778. When you do this, the ports used by Solaris communication deamons will be 32778 or greater, leaving port 32777 available for IBM FileNet use.

   When Solaris first starts up, it takes the first several ports, called anonports, for its communication daemons. By default, the maximum tcp_smallest_anon_port is 32768. IBM FileNet uses several ports higher than 32768. See the appendix “IBM FileNet P8 ports” in the *Plan and Prepare Your Environment for IBM FileNet P8* for details on which ports IBM FileNet uses.

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

   a. Determine the current tcp_smallest_anon_port setting.

      From a command prompt, enter the following:

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

   b. Enable port 32777.

      If the port returned in the step above is less than 32778, you must enable port 32777.

Solaris 9

Edit the /etc/rc2.d/S69inet file.

Add the following line before the exit 0 entry at the bottom of the file:

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

Solaris 10

Edit the /lib/svc/method/net-init file.

Add the following line before the exit 0 entry at the bottom of the file:

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

c. Reboot the Application Engine server.

You must reboot the Application Engine server to force the release of ports required by the Application Engine that may be in use by the operating system.

CAUTION Failure to reboot after these changes are made can result in the port being unavailable, generating OpenSocket errors.

6. (Optional, UNIX only) Set the the P8TASKMAN_HOME system environment variable.

The Process Engine Client installer will set this variable for you. If you want to specify a different location for the client files, you can do so by setting the P8TASKMAN_HOME system variable. The default value set by the PE Client installer is:

P8TASKMAN_HOME=*AE_install_path*/CommonFiles

7. Continue with Task 2 “Install Application Engine software updates” on page 116.

Task 2: Install Application Engine software updates

Install any fix packs and interim fixes required for Application Engine.

To install the Application Engine software updates

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your support representative.
2. Open the readmes for any subsequent fix packs or interim fixes (typically optional) and perform the installation procedures provided.
3. Install the latest updates for the Content Engine Client and Process Engine Client files using the subsequent tasks.

Task 3: Install the latest Content Engine Client files on Application Engine servers

To install the 4.5 release or fix pack version Content Engine Client files, perform the following procedures on all Application Engine servers. Except where noted, these steps apply to WebSphere, WebLogic, and JBoss.

To install the Content Engine Client files

Refer to the appropriate information from your installation worksheet while performing the following steps.

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. On the machine where Application Engine is installed, log on as any user who has the following permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Application Engine is installed
   - Execute permission on the Content Engine Client install software
3. Verify that there is a current backup of the Application Engine.
4. Copy the Content Engine Client install software from the Content Engine installation software to the temporary directory. The version of the install software must match the version of Content Engine.
5. Expand the (TAR or ZIP) Content Engine Client install software within the temporary directory.
6. Close all instances of Workplace and any other Application Engine client applications. From the application server administrative console, stop and un-deploy Application Engine.

   WebSphere

   Stop and un-deploy the FileNet Application Engine application.

   WebLogic

   Stop and un-deploy the FileNet Application Engine application.

   JBoss

   Execute the shutdown command.

7. (WebLogic only) Manually delete the following application server cache directories, substituting your domain name in place of *mydomain*:

   WebLogic UNIX

   /opt/bea/user_projects/domains/*mydomain*/servers/AdminServer/tmp/_WL_user/app_engine

   /opt/ bea/user_projects/domains/*mydomain*/servers/.wlnotdelete

   /opt/bea/user_projects/domains/*mydomain*/servers/AdminServer/cache/EJBCompilerCache

   WebLogic Windows

   C:\bea\user_projects\domains*mydomain*\servers\AdminServer\tmp_WL_user\app_engine

   C:\bea\user_projects\domains*mydomain*\servers\.wlnotdelete

   C:\bea\user_projects\domains*mydomain*\servers\AdminServer\cache\EJBCompilerCache

8. Start the Content Engine client install process.

   - To install the Content Engine client interactively:

     i. Access the IBM FileNet Content Engine client update software.

     ii. Run one of the commands in the table below, *CE_version* is the version of Content Engine you intend to run, for example, 4.5.0..

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN`

     iii. Complete the installation program wizard.

   - To install the Content Engine client files silently:

     i. Make a back up copy of the CEClient_silent_install.txt input file.

     ii. Open the silent input file in a text editor. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your update.

iii. Navigate to the path containing the Content Engine Client installation program, and run one of the commands in the following table to perform the silent install, where:

CE_version is the version of Content Engine you intend to run, for example, 4.5.0.

path is the path that contains the installation program.

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN -f `*`path`*`/CECLIENT.AIX/CEClient_silent_install.txt -i silent`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN -f `*`path`*`/CEClient.HPUX/CEClient_silent_install.txt -i silent`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN -f `*`path`*`/CEClient.HPUXI/CEClient_silent_install.txt -i silent`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN -f `*`path`*`/CEClient.Linux/CEClient_silent_install.txt -i silent`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN -f `*`path`*`/CEClient.Solaris/CEClient_silent_install.txt -i silent`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE -f `*`path`*`\CEClient.Windows\CEClient_silent_install.txt -i silent`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN -f `*`path`*`/CEClient.zLinux/CEClient_silent_install.txt -i silent`

9. Install Process Engine client updates on the Application Engine. Continue at "Install the Latest Process Engine Client files on Application Engine servers" on page 120.

Task 4: Install the Latest Process Engine Client files on Application Engine servers

To install the Process Engine Client files, perform the following steps on all Application Engine servers.

To install the Process Engine client files

1. On the machine where Application Engine is to be deployed, log on as any user who has the following permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Application Engine is installed
   - Execute permission on the Process Engine Client install software
2. Copy the Process Engine Client install software from the Process Engine installation software to the temporary directory. The version of the install software must match the version of Process Engine.
3. Expand the (TAR or ZIP) Process Engine Client install software within the temporary directory.
4. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Application Engine is deployed. Run the program either interactively (using the install wizard) or silently.

   To run the program interactively, run one of the commands in the table below, follow the wizard instructions, and then continue at step 5:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN`

To run the program silently, perform the following steps:

a. In the expanded install software, open the file PEClient_silent_install.txt and edit it as follows:

   i. Change the `Variable_CheckboxAE` line to the following:

   ```
   -V Variable_CheckboxAE="true"
   ```

   ii. Run one of the commands in the following table to perform the silent install:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN -silent -options "PEClient_silent_install.txt"`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN -silent -options "PEClient_silent_install.txt"`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN -silent -options "PEClient_silent_install.txt"`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN -silent -options "PEClient_silent_install.txt"`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN -silent -options "PEClient_silent_install.txt"`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE -silent -options "PEClient_silent_install.txt"`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN -silent -options "PEClient_silent_install.txt"`

5. Configure Application Engine.

   Follow the instructions for your application server to configure Application Engine:

   - “Configure Application Engine (WebSphere)” on page 122
   - “Configure Application Engine (WebLogic)” on page 132
   - “Configure Application Engine (JBoss)” on page 138

Task 5a: Configure Application Engine (WebSphere)

This topic covers the configuration of your Application Engine web application (Workplace) on WebSphere. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using WebSphere with container-managed authentication or SSO, edit web.xml. See "To edit web.xml for container-managed authentication or SSO" on page 122.
2. If you are using SSO, edit web.xml. See "(SSO Only) To edit web.xml for SSO (optional)" on page 125.
3. Configure the Application Engine. See "To configure Application Engine (WebSphere)" on page 126.
4. Configure the server ports. See "To configure the server ports" on page 130.

To edit web.xml for container-managed authentication or SSO

NOTE Perform this procedure only if your site uses WebSphere with container-managed authentication or Single Sign-On (SSO). If you are using SSO, you must perform additional configuration steps as directed at the end of this procedure.

1. Make a back-up copy of web.xml.

   AE_install_path/Workplace/WEB-INF/web.xml

2. Edit web.xml.

   NOTE Text in **bold** in the examples below indicates changes made to the original web.xml file.

   a. Search for the parameter challengeProxyEnabled and set it to false.

   ```
   <param-name>challengeProxyEnabled</param-name>
   <param-value>false</param-value>
   ```

   b. Search for the first instance of <web-resource-collection>, and uncomment the url-pattern as noted in the file comments below.

   ```
   <web-resource-collection>
      <web-resource-name>action</web-resource-name>
      <description>Define the container secured resource</description>
      <url-pattern>/containerSecured/*</url-pattern>

               <!--

                   Uncomment this section if all resources that require credentials
         must be secured in order to obtain a secured Thread.  If using WebSphere,
         this section must be
                    uncommented.
         --> Move this commenting tag here from just before the </web-resource-
         collection> closing tag below.

         <url-pattern>/containerSecured/*</url-pattern>
         <url-pattern>/</url-pattern>
         <url-pattern>/author/*</url-pattern>
         <url-pattern>/Browse.jsp</url-pattern>
         <url-pattern>/eprocess/*</url-pattern>
         <url-pattern>/Favorites.jsp</url-pattern>
         <url-pattern>/GetPortalSitePreferences.jsp</url-pattern>
         <url-pattern>/GetTokenSignIn.jsp</url-pattern>
   ```

```
<url-pattern>/GetUserInformation.jsp</url-pattern>
<url-pattern>/GetUserToken.jsp</url-pattern>
<url-pattern>/HomePage.jsp</url-pattern>
<url-pattern>/IntegrationWebBasedHelp.jsp</url-pattern>
<url-pattern>/is/*</url-pattern>
<url-pattern>/operations/*</url-pattern>
<url-pattern>/portlets/Author/edit.jsp</url-pattern>
<url-pattern>/portlets/Author/portlet.jsp</url-pattern>
<url-pattern>/portlets/Browse/edit.jsp</url-pattern>
<url-pattern>/portlets/Browse/portlet.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/edit.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/portlet.jsp</url-pattern>
<url-pattern>/portlets/GroupPageDesign.jsp</url-pattern>
<url-pattern>/portlets/GroupPageSettings.jsp</url-pattern>
<url-pattern>/portlets/Inbox/edit.jsp</url-pattern>
<url-pattern>/portlets/Inbox/portlet.jsp</url-pattern>
<url-pattern>/portlets/MultiPagesDesign.jsp</url-pattern>
<url-pattern>/portlets/OrganizePages.jsp</url-pattern>
<url-pattern>/portlets/PortalPageDesign.jsp</url-pattern>
<url-pattern>/portlets/PortalPageInfo.jsp</url-pattern>
<url-pattern>/portlets/PortletAlias.jsp</url-pattern>
<url-pattern>/portlets/PortletSettings.jsp</url-pattern>
<url-pattern>/portlets/PreviewAndSetup.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/edit.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/portlet.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/edit.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/portlet.jsp</url-pattern>
<url-pattern>/portlets/Workflows/edit.jsp</url-pattern>
<url-pattern>/portlets/Workflows/portlet.jsp</url-pattern>
<url-pattern>/properties/*</url-pattern>
<url-pattern>/redirect/*</url-pattern>
<url-pattern>/regions/*</url-pattern>
<url-pattern>/Search.jsp</url-pattern>
<url-pattern>/select/*</url-pattern>
<url-pattern>/SelectReturn.jsp</url-pattern>
<url-pattern>/Tasks.jsp</url-pattern>
<url-pattern>/UI-INF/*</url-pattern>
<url-pattern>/utils/*</url-pattern>
<url-pattern>/WcmAdmin.jsp</url-pattern>
<url-pattern>/WcmAuthor.jsp</url-pattern>
<url-pattern>/WcmBootstrap.jsp</url-pattern>
<url-pattern>/WcmCloseWindow.jsp</url-pattern>
<url-pattern>/WcmDefault.jsp</url-pattern>
<url-pattern>/WcmError.jsp</url-pattern>
<url-pattern>/WcmJavaViewer.jsp</url-pattern>
<url-pattern>/WcmObjectBookmark.jsp</url-pattern>
<url-pattern>/WcmPortletHelp.jsp</url-pattern>
<url-pattern>/WcmPortletSearch.jsp</url-pattern>
<url-pattern>/WcmQueueBookmark.jsp</url-pattern>
<url-pattern>/WcmSignIn.jsp</url-pattern>
<url-pattern>/WcmSitePreferences.jsp</url-pattern>
<url-pattern>/WcmUserPreferences.jsp</url-pattern>
<url-pattern>/WcmWorkflowsBookmark.jsp</url-pattern>
<url-pattern>/wizards/*</url-pattern>
<url-pattern>/Author/*</url-pattern>
<url-pattern>/axis/*.jws</url-pattern>
<url-pattern>/Browse/*</url-pattern>
<url-pattern>/ceTunnel</url-pattern>
<url-pattern>/CheckoutList/*</url-pattern>
<url-pattern>/downloadMultiTransferElement/*</url-pattern>
<url-pattern>/ExternalUrl/*</url-pattern>
<url-pattern>/findRecordTarget</url-pattern>
<url-pattern>/formCallback/*</url-pattern>
<url-pattern>/getAnnotSecurity/*</url-pattern>
<url-pattern>/getCEAnnotations/*</url-pattern>
<url-pattern>/getContent/*</url-pattern>
```

```
        <url-pattern>/getForm/*</url-pattern>
        <url-pattern>/getISAnnotations/*</url-pattern>
        <url-pattern>/getISAnnotSecurity/*</url-pattern>
        <url-pattern>/getISContent/*</url-pattern>
        <url-pattern>/getMultiContent/*</url-pattern>
        <url-pattern>/getPreview</url-pattern>
        <url-pattern>/getProcessor/*</url-pattern>
        <url-pattern>/getRealms/*</url-pattern>
        <url-pattern>/getUsersGroups/*</url-pattern>
        <url-pattern>/Inbox/*</url-pattern>
        <url-pattern>/integrationCommandProxy</url-pattern>
        <url-pattern>/integrationResponse</url-pattern>
        <url-pattern>/integrationResponseProxy</url-pattern>
        <url-pattern>/integrationWebBasedCommand</url-pattern>
        <url-pattern>/keepAlive</url-pattern>
        <url-pattern>/launch/*</url-pattern>
        <url-pattern>/PublicQueue/*</url-pattern>
        <url-pattern>/putContent/*</url-pattern>
        <url-pattern>/QuickSearch/*</url-pattern>
        <url-pattern>/signingServlet/*</url-pattern>
        <url-pattern>/transport/*</url-pattern>
        <url-pattern>/upload/*</url-pattern>
        <url-pattern>/vwsimsoapservlet</url-pattern>
        <url-pattern>/vwsoaprouter</url-pattern>
        <url-pattern>/Workflows/*</url-pattern>  Move the closing comment tag from
        here to the location indicated at the beginning of this example.
    </web-resource-collection>
```

c. Locate the section <auth-constraint>, comment the wild-card (*) role-name as noted in the file comments below.

```
<auth-constraint>
    <!-- <role-name>*</role-name> -->
    <!-- For WebSphere 6, use the role-name line below instead of the wildcard role
    above.-->

    <role-name>All Authenticated</role-name>

    <!-- For WebSphere 6, add this security-role element below the login-config
    element (below).
       <security-role>
          <description>All Authenticated</description>
          <role-name>All Authenticated</role-name>
       </security-role>
             -->
</auth-constraint>
```

d. Locate the end of the </login-config> element, and add the All Authenticated users role-element after the closing tag.

```
<security-role>
    <description>All Authenticated</description>
    <role-name>All Authenticated</role-name>
</security-role>
```

e. Search for the first instance of a <security-constraint> tag, and add the following <security-constraint> tag before that tag.

CAUTION Enter the information below as single lines without line breaks.

```
<security-constraint>
    <web-resource-collection>
        <web-resource-name>action</web-resource-name>
        <description>Define the non-secured resource</description>
```

```
            <url-pattern>/P8BPMWSBroker/*</url-pattern>
        </web-resource-collection>
    </security-constraint>
```

3. Save your changes to web.xml and close the file.
4. If your site uses SSO, Continue on with "(SSO Only) To edit web.xml for SSO (optional)" on page 125, otherwise continue with "To configure Application Engine (WebSphere)" on page 126.

(SSO Only) To edit web.xml for SSO (optional)

NOTE Perform this procedure only if your site uses SSO with a proxy server. You must use this procedure to modify web.xml to enable SSO.

1. Open web.xml for editing.
2. At the end of web.xml, comment out the <login-config> element, as follows:

```
<!--
<login-config>
    <auth-method>FORM</auth-method>
    <realm-name>AE Workplace</realm-name>
    <form-login-config>
        <form-login-page>/ContainerLogin.jsp</form-login-page>
        <form-error-page>/ContainerError.jsp</form-error-page>
    </form-login-config>
</login-config>
-->
```

3. As needed, set the ssoProxyContextPath, ssoProxyHost, ssoProxyPort, and ssoProxySSLPort.

   These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, you must set the corresponding sso*<proxy host element>* parameter for that element in the URL to the value for the SSO proxy host.

   The default settings are (in **bold** below):

```
<init-param>
    <param-name>ssoProxyContextPath</param-name>
    <param-value></param-value>
</init-param>
<init-param>
    <param-name>ssoProxyHost</param-name>
    <param-value></param-value>
</init-param>
<init-param>
    <param-name>ssoProxyPort</param-name>
    <param-value></param-value>
</init-param>
<init-param>
    <param-name>ssoProxySSLPort</param-name>
    <param-value></param-value>
</init-param>
```

   In general, the init parameters above must be configured as follows:

   - **ssoProxyContextPath**: Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

For example, if the Workplace deploy host URL is `http://deploy_server:2809/Workplace` and the SSO proxy host URL is `http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

```
<param-name>ssoProxyContextPath</param-name>
<param-value>/fn/Workplace</param-value>
```

- **ssoProxyHost**: Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

  For example, if the host URL where Workplace is deployed is `http://deploy_server/Workplace` and the corresponding SSO proxy host URL is `http://sso_proxy_server/Workplace`, then use the following:

```
<param-name>ssoProxyHost</param-name>
<param-value>sso_proxy_server</param-value>
```

- **ssoProxyPort**: Set the value to the http port on the SSO proxy host.

  For example:

```
<param-name>ssoProxyPort</param-name>
<param-value>80</param-value>
```

- **ssoProxySSLPort**: Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

  For example:

```
<param-name>ssoProxySSLPort</param-name>
<param-value>443</param-value>
```

4. Save your changes to web.xml and close the file.

To configure Application Engine (WebSphere)

1. Open the WebSphere administrative console.
2. Set JVM settings for JAAS login configuration and memory settings.
   a. In the Java Virtual Machine settings, set the JAAS login entry in the Generic JVM argument field to one of the following (do not enter the linebreaks):

   `-Djava.security.auth.login.config=`*`ae_install_path`*`\CE_API\config\jaas.conf.WebSphere`

   NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

   CAUTION Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the field.

   CAUTION (Windows only) On WebSphere for Windows the path cannot contain a space. You must use 8.3-notation for the install path information.

   If *AE_install_path* is:

   `C:\Program Files\FileNet\AE`

use:

```
C:\Progra~1\FileNet\AE
```

b. Set the Initial and Maximum Heap Size.

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

c. Save your changes.

3. (For installations with Content Engine and Application Engine collocated on the WebSphere server, but in different WebSphere profiles) Create an additional JVM property for different WebSphere profiles.

Perform the following steps on *both* the Content Engine profile and the Application Engine profile:

a. In the Java Virtual Machine settings, create a new Custom Property:

com.ibm.websphere.orb.uniqueServerName

b. Set the Value to true.

c. Save your changes.

d. Restart WebSphere.

4. Verify that Application Integration installation is enabled.

a. In the Environment area, in Virtual Hosts, navigate to **default_host** (or the host your application is deployed under).

b. Verify that MIME type is set to application/octet-stream, or set it if it is not. Include .exe in the Extension(s) field.

c. Save your changes.

5. Configure Light weight Third Party Authentication (LTPA).

NOTE Skip this step if your Application Engine and Content Engine are located on the same WebSphere application server profile.

a. On the Content Engine server, do the following:

NOTE If you are already using LTPA with your Content Engine application server, you only need to export the existing keys per Step iii thru Step iv below.

i. Log in to the WebSphere adminstrative console.

ii. If LTPA is not already configured for Content Engine, create an LTPA password, specify the location for the key file on the Content Engine server, and generate the keys.

NOTE For password restrictions, see the WebSphere documentation. If you have already configured Content Engine for LTPA, use the existing password in the Application Engine configuration below.

iii. Export the keys using the WebSphere administrative console.

iv. Copy the key file from CE server location you specified above to a directory on the Application Engine server. On Windows, for example, C:\LTPA*ltpa_key_name*.

v. Save your changes, and stop and restart WebSphere.

b. On the Application Engine server, do the following:

i. Log in to the WebSphere adminstrative console.

ii. Create an LTPA password, using the same password as the Content Engine.

iii. Set the timeout value for LTPA authentication to a value that is larger than the default.

For example, if the timeout value is set to 2 hours, the LTPA key expires and end users will not be able to log in to Workplace after having been logged in for 2 hours.

iv. Specify the path for the key file that you copied to the Application Engine server. For example, C:\LTPA*ltpa_key_name*.

v. Import the keys from the Key file.

vi. Set the following Security settings:

– Turn on (check) **Enable Administrative Security flag**.

– Turn on (check) **Enable application security flag**.

– Turn off (uncheck) **Java 2 Security**.

NOTE The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.

– Set the Active Authentication Mechanism to LTPA (Light weight Third Party Authentication).

6. Configure Lightweight Directory Access Protocol (LDAP).

a. Set "Standalone LDAP registry" to current for the available realm definition.

b. Configure the LDAP provider to exactly match the corresponding settings on the Content Engine application server.

– Primary administrative user name

– Select “Automatically generated server identity."

– Type

– Host

– Port

– Base distinguished name (DN)

– Bind distinguished name (DN)

– Bind password

c. Configure the Advanced Lightweight Directory Access Protocol (LDAP) user registry settings to exactly match the corresponding settings from the Content Engine application server.

   – User filter
   – Group Filter
   – User ID map
   – Group member ID map
   – Certificate map mode
   – Certificate filter

d. Save these settings.

NOTE If you are using Tivoli as your directory server, perform Step e. For all other directory servers, continue with Step f, below.

e. (Tivoli only) Configure WebSphere security for Tivoli.

   i. Shut down Websphere.

   ii. On the Application Engine WebSphere server, back up the security.xml file. The file is in the following location:

   install_path/WebSphere/AppServer/profiles/*appserver_name*/config/cells/*ServerName*Node01Cell/

   iii. On the Content Engine WebSphere server, open the security.xml file, located in a path similar to the one in Step iii. Copy the “userRegistries”, “searchFilter”, and “host” entries, as shown in the example below (in bold). Do not use the example settings, as your settings will be different.

```
  <userRegistries xmi:type="security:LocalOSUserRegistry"
xmi:id="LocalOSUserRegistry" serverId="" serverPassword="{xor}" realm=""
useRegistryServerId="true" primaryAdminId=""/>

  <userRegistries xmi:type="security:CustomUserRegistry"
xmi:id="CustomUserRegistry_1" useRegistryServerId="true" primaryAdminId=""
customRegistryClassName="com.ibm.websphere.security.FileRegistrySample"/>

  <userRegistries xmi:type="security:LDAPUserRegistry"
xmi:id="LDAPUserRegistry_1" serverId="CEAdmin"
serverPassword="{xor}Dz4sLCgwLTtu" realm="vauxhall:389" ignoreCase="true"
useRegistryServerId="false" primaryAdminId="CEAdmin"
type="IBM_DIRECTORY_SERVER" sslEnabled="false" sslConfig=""
baseDN="dc=gardens"
bindDN="cn=CEAdmin,ou=Shared,ou=Engineering,ou=FileNet,dc=gardens"
bindPassword="{xor}Dz4sLCgwLTtu" searchTimeout="120" reuseConnection="true">

    <searchFilter xmi:id="LDAPSearchFilter_1"
userFilter="(&amp;(cn=%v)(objectclass=person))"
groupFilter="(&amp;(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUn
iqueNames)))" userIdMap="cn" groupIdMap="*:cn"
groupMemberIdMap="memberof:member" certificateMapMode="EXACT_DN"
certificateFilter=""/>

    <hosts xmi:id="EndPoint_1" host="vauxhall" port="389"/>

  </userRegistries>
```

```
  <userRegistries xmi:type="security:WIMUserRegistry"
xmi:id="WIMUserRegistry_1" serverId="" serverPassword="{xor}"
realm="defaultWIMFileBasedRealm" ignoreCase="true"
useRegistryServerId="false" primaryAdminId="administrator"
registryClassName="com.ibm.ws.wim.registry.WIMUserRegistry"/>

  <authConfig xmi:id="AuthorizationConfig_1" useJACCProvider="false">

    <authorizationProviders xmi:id="AuthorizationProvider_1"
j2eePolicyImplClassName="com.tivoli.pd.as.jacc.TAMPolicy" name="Tivoli Access
Manager"
policyConfigurationFactoryImplClassName="com.tivoli.pd.as.jacc.TAMPolicyConf
igurationFactory"
roleConfigurationFactoryImplClassName="com.tivoli.pd.as.jacc.TAMRoleConfigur
ationFactory"
initializeJACCProviderClassName="com.tivoli.pd.as.jacc.cfg.TAMConfigInitiali
ze" requiresEJBArgumentsPolicyContextHandler="false"
supportsDynamicModuleUpdates="true"/>

  </authConfig>
```

iv. On the Application Engine WebSphere server, open the security.xml file for editing, and paste in the “userRegistries”, “searchFilter”, and “host” entries you copied from the Content Engine file, in the location shown above.

v. Save and close the file.

vi. Restart WebSphere, and log in to the administrative console.

vii. Navigate to Security > Secure administration, applications and infrastructure.

NOTE If you make any changes to your LDAP server settings in the WebSphere Administrative Console after updating the security.xml file, these settings will be overwritten and you will have to add the entries back to the file.

f. Turn on (check) the Enable administrative security flag.

g. Turn off (uncheck) Use Java 2 security to restrict application access to local resources.

h. Save your changes to the master configuration.

i. Test the connection on the Standalone LDAP registry page. If the test fails, correct the error before proceeding. If it passes, click **OK** to return to the previous page.

j. Save your changes.

7. Continue with “To configure the server ports” on page 130.

To configure the server ports

This configuration is not required but is recommended.

1. Stop the WebSphere server.

2. Make a backup copy of serverindex.xml located in:

   WAS_HOME\profiles\default\config\cells*machine_name*Node01Cell\nodes*machine_name*Node01\

3. Edit serverindex.xml.

   Locate the <specialEndpoints> section, and change the port numbers for the three SSL listener addresses to “0” as shown below:

```
<specialEndpoints xmi:id="NamedEndPoint_1155689929072"
endPointName="SAS_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929072" host="host_name" port="0"/>
</specialEndpoints>
<specialEndpoints xmi:id="NamedEndPoint_1155689929073"
endPointName="CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929073" host="host_name" port="0"/>
</specialEndpoints>
<specialEndpoints xmi:id="NamedEndPoint_1155689929074"
endPointName="CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929074" host="host_name" port="0"/>
</specialEndpoints>
```

4. Save and close the serverindex.xml file.

5. Restart WebSphere.

6. Deploy the Workplace application. See “Deploy Application Engine (WebSphere)” on page 141.

Task 5b: Configure Application Engine (WebLogic)

This topic covers the configuration of your Application Engine application (Workplace) on WebLogic. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using SSO, edit web.xml. See "To edit web.xml for SSO (optional)" on page 132.
2. Modify the application server startup script. See "To modify the application server startup script" on page 133.
3. Configure Application Engine. See "To configure Application Engine (WebLogic)" on page 135.
4. Modify config.xml to support passing user credentials to clients such as Application Integration and WebDAV. See "To enable passing user credentials to client applications (WebLogic)" on page 136.

To edit web.xml for SSO (optional)

NOTE Perform this procedure only if your site uses SSO with a proxy server. You must modify web.xml to enable SSO.

1. Make a backup copy of web.xml.

   AE_install_path/Workplace/WEB-INF/web.xml

2. Edit web.xml.

   a. Set the parameter perimeterChallengeMode to true, as in:

   ```
   <init-param>
      <param-name>perimeterChallengeMode</param-name>
      <param-value>true</param-value>
   </init-param>
   ```

   b. As needed, set the ssoProxyContextPath, ssoProxyHost, ssoProxyPort, and ssoProxySSLPort.

   These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, then you must set the corresponding sso* parameter for that element in the URL to the value for the SSO proxy host.

   The default settings are (in **bold**):

   ```
   <init-param>
      <param-name>ssoProxyContextPath</param-name>
      <param-value></param-value>
   </init-param>
   <init-param>
      <param-name>ssoProxyHost</param-name>
      <param-value></param-value>
   </init-param>
   <init-param>
      <param-name>ssoProxyPort</param-name>
      <param-value></param-value>
   </init-param>
   ```

```
<init-param>
    <param-name>ssoProxySSLPort</param-name>
    <param-value></param-value>
</init-param>
```

In general, the init parameters above should be configured as follows:

- **ssoProxyContextPath**: Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

  For example, if the Workplace deploy host URL is `http://deploy_server:2809/Workplace` and the SSO proxy host URL is `http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

  ```
  <param-name>ssoProxyContextPath</param-name>
  <param-value>/fn/Workplace</param-value>
  ```

- **ssoProxyHost**: Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

  For example, if the host URL where Workplace is deployed is `http://deploy_server/Workplace` and the corresponding SSO proxy host URL is `http://sso_proxy_server/Workplace`, then use the following:

  ```
  <param-name>ssoProxyHost</param-name>
  <param-value>sso_proxy_server</param-value>
  ```

- **ssoProxyPort**: Set the value to the http port on the SSO proxy host.

  For example:

  ```
  <param-name>ssoProxyPort</param-name>
  <param-value>80</param-value>
  ```

- **ssoProxySSLPort**: Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

  For example:

  ```
  <param-name>ssoProxySSLPort</param-name>
  <param-value>443</param-value>
  ```

3. Save your changes to web.xml and close the file.

To modify the application server startup script

1. Stop the WebLogic application server if running.
2. Make a backup copy of the application server startup script.

   Back up startWebLogic.cmd for Windows or startWebLogic.sh for UNIX.

   NOTE If you are not using a WebLogic domain, backup startWLS.cmd for Windows or startWLS.sh for UNIX.

   a. Edit the MEM_ARGS variable.

      Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

NOTE If the MEM_ARGS variable doesn't exist, add it to the startup script.

- For all systems except those using JRockit JAVA.

  Append the following to the MEM_ARGS variable:

  ```
  -XX:MaxPermSize=sizem
  ```

  where *size* is the value, in MB, of the MaxPermSize.

  Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

- For systems using JRockit JAVA.

  Append the following to the MEM_ARGS variable:

  ```
  -Xgc:gencon
  ```

b. Configure JAAS login.

Add one of the following right after the classpath entry for WebLogic.

CAUTION Enter the jaas_login entry (bold below) as a single line without line breaks. Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the script.

NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

Windows

```
@REM Jaas Login configuration setting
set JAAS_LOGIN=%JAAS_LOGIN% -Djava.security.auth.login.config=
"AE_install_path\CE_API\config\jaas.conf.WebLogic"
```

UNIX

```
# Jaas Login configuration setting
JAAS_LOGIN="${JAAS_LOGIN}" -Djava.security.auth.login.config=
"AE_install_path/CE_API/config/jaas.conf.WebLogic"
```

AIX

```
# Jaas Login configuration setting
JAAS_LOGIN="${JAAS_LOGIN}" -Djava.security.auth.login.config=
"AE_install_path/CE_API/config/jaas.conf.WebLogic"
JAAS_LOGIN="${JAAS_LOGIN}" -
Dlogin.configuration.provider=com.ibm.security.auth.login.ConfigFile
```

c. Add %JAAS_LOGIN% section as indicated in the examples below in **bold**.

Windows - in the WLS_REDIRECT_LOG settings

```
If "%WLS_REDIRECT_LOG%"==""(
    echo Starting WLS with line:
   echo %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN%
-Dweblogic.Name=%SERVER_NAME% -
```

```
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS%
    %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN% -
Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
Dweblogic.management.password=%WLS_PW% -
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS%
) else (
echo Redirecting output from WLS window to %WLS_REDIRECT_LOG%
%JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN% -
Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
Dweblogic.management.password=%WLS_PW% -
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS% >”%WLS_REDIRECT_LOG%” 2>&1
)
```

UNIX - in the WLS_REDIRECT_LOG settings

```
${JAVA_HOME}/bin/java ${JAVA_VM} -version

if [ "${WLS_REDIRECT_LOG}" = "" ] ; then
        echo "Starting WLS with line:"
        echo "${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS} -
Dweblogic.Name=${SERVER_
NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_CLASS
}"
        ${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
${JAAS_LOGIN} -Dweblogic.Name=${S
ERVER_NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_
CLASS}
else
        echo "Redirecting output from WLS window to ${WLS_REDIRECT_LOG}"
        ${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
${JAAS_LOGIN} -Dweblogic.Name=${S
ERVER_NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_
CLASS}  >"${WLS_REDIRECT_LOG}" 2>&1
fi
```

3. Save and close the server startup script.

To configure Application Engine (WebLogic)

1. (If you selected Container-Managed Authentication during the installation) Enable trust between WebLogic domains for the Content Engine domain and the Application Engine domain.

   Do the following on both the Content Engine application server and the Application Engine application server.

   a. In the WebLogic Administration Console, in the security settings, enter a password for the domain. You must enter the same password for both the Content Engine domain and Application Engine domain.

   b. Save your changes.

   c. Restart the server if needed.

   d. Repeat this procedure in each domain for which you want to enable trust.

2. (If you selected Container-Managed Authentication during the installation) Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
   a. Refer to your Content Engine installation worksheet items and the WebLogic Administration Console settings for Compatibility Security > Realms for Authentication Provider, users, and groups on Content Engine.

      Configure the LDAP provider to exactly match the settings from the Content Engine server.

      - Group Base DN:
      - User Name Attribute:
      - Port:
      - User Base DN:
      - Principal:
      - Credential:
      - Confirm Credential:
      - Host:
      - User From Name Filter:
      - Group From Name Filter:

   b. Restart WebLogic.
3. Set permissions for the user running the application server.

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed AE (*AE_install_path*).
4. Continue with "To enable passing user credentials to client applications (WebLogic)" on page 136.

To enable passing user credentials to client applications (WebLogic)

Perform this procedure to enable passing user credentials between Application Engine and its client applications such as WebDAV and Application Integration.

CAUTION If you do not make this change to config.xml, then end users will be prompted to enter their user name and password to complete any client operations, such as adding a document.

1. Stop the WebLogic server.
2. Make a backup copy of config.xml located in deployment directory.

   For example:

   BEA_home/bea/user_projects/domains/*domain_name*/config/config.xml
3. Edit config.xml.

   CAUTION The enforce-valid-basic-auth-credentials entry should be entered as a single line without line breaks.

a. Locate the <security-configuration> section and add the following line to the end of the section, just before the </security-configuration> tag:

```
<enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-
credentials>
```

b. Save your changes to config.xml and close the file.

4. Restart WebLogic.

5. Deploy the Workplace application. See "Deploy Application Engine (WebLogic)" on page 143.

Task 5c: Configure Application Engine (JBoss)

This topic covers the configuration of your Application Engine application (Workplace) on JBoss. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

To modify the application server startup script

1. Stop JBoss if it is running.
2. Make a backup copy of the application server startup script.

   UNIX

   run.sh

   Windows

   run.bat
3. Edit the application server startup script Java settings.

   a. Add a line to specify the path to the JDK provided by JBoss, as shown in the following example (Windows):

   ```
   set JAVA_HOME=C:\Program Files\Java\jdk1.5.0_06
   ```

   NOTE If your JDK is different from version 1.5.0, substitute your version for the one listed above.

   b. Update the JAVA_OPTS memory settings.

   Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

   In the `JAVA_OPTS` line, change the -Xms and -Xmx values (**bold**) for your configuration.

   Example (Windows):

   ```
   set JAVA_OPTS=%JAVA_OPTS% -Xms128m -Xmx512m
   ```

   Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

   c. Save your edits.
4. Configure JAAS login.

   Add one of the following right after the "$JAVA" $JAVA_OPTS (UNIX) or "%JAVA%" %JAVA_OPTS (Windows) entry in the startup script.

   CAUTION Enter the jaas_login entry (bold below) as a single line without line breaks. Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the script.

NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

UNIX

```
"$JAVA" $JAVA_OPTS -Djava.security.auth.login.config=”/opt/FileNet/AE/CE_API/
config/jaas.conf.JBoss” "-Djava.endorsed.dirs=$JBOSS_ENDORSED_DIRS" -classpath
"$JBOSS_CLASSPATH" org.jboss.Main $@
```

Windows

```
"%JAVA%" %JAVA_OPTS% "-Djava.security.auth.login.config=C:\Program
Files\FileNet\AE\CE_API\config\jaas.conf.JBoss" "-
Djava.endorsed.dirs=%JBOSS_ENDORSED_DIRS%" -classpath "%JBOSS_CLASSPATH%"
org.jboss.Main %*
```

5. Save and close the server startup script.
6. Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
   a. On the Application Engine server, open login-config.xml, located in *JBoss_home*/server/*server_name*/conf, for editing.
   b. Set the <application-policy name="FileNet"> entry identical to the corresponding entry in the login-config.xml file on the Content Engine server.
   c. Restart JBoss.
7. Set permissions for the user running the application server.

   NOTE On Windows, the following is only required for NTFS formatted partitions

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed Application Engine (*AE_install_path*).
8. (Optional) Disable JBoss logging.

   In development mode, JBoss creates a large number of HTTP Access, “INFO”, “DEBUG” and “TRACE” log messages. This can cause unexpected behavior in the deployed IBM FileNet software. Using the following procedure, you can limit this type of excessive JBoss logging.

   NOTE When logging is disabled, error messages will still be displayed in the JBoss console.

   a. Edit the log4j.xml file (*JBOSS_home*/server/*server_name*/conf/log4j.xml).
      i. Change all threshold values and priority values from "INFO", "DEBUG", or "TRACE" to "ERROR".
      ii. Delete or comment out the "Preserve messages in a local file" to turn off the server log.
   b. To turn off HTTP access logging, open jboss-service.xml with a text editor and delete or comment out the "Access logger" section.

      Location of jboss-service.xml:

      JBoss_Home/server/*server_name*/deploy/jbossweb-tomcat55.sar/META-INF

c. Open web.xml and change the logVerbosityLevel to "FATAL".

   Location of web.xml:

   JBoss_Home/server/*server_name*/deploy/jbossweb-tomcat55.sar/conf

d. Restart the JBoss server.

9. Deploy the Workplace application. See "Deploy Application Engine (JBoss)" on page 145.

Task 6a: Deploy Application Engine (WebSphere)

This topic covers the deployment and start of your Application Engine application (Workplace) on WebSphere.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.exe (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.exe (Windows) from the following location:

     AE_install_path/deploy/

To deploy Application Engine

1. Log on to the WebSphere administrative console.
2. Navigate to the dialog for installing a new application.
3. Select the file to deploy.
   - (If the administrative console is running *locally*) Select **Local Path** and enter or browse to the location of the app_engine.war or app_engine.ear file created by the installation program (see below for the default path). Do not enter the machine name.
   - (If the administrative console is *remote*) Select **Server path** and enter the fully-qualified pathname to the app_engine.war or app_engine.ear file. Do not enter the machine name.

     AE_install_path/deploy
4. If you are deploying a WAR file, enter the context root.

   Enter `Workplace` and click **Next** to proceed to deploying a new application.

   NOTE The context root is the name of the application you log in to using the web interface, such as:

   `http://`*`ApplicationEngineServerName`*`:`*`port`*`#/`*`Context_Root`*.

5. Complete the dialogs for installing a new application, using the following settings:

   Application name: `Workplace`, or the name you chose to call the application.

   WebServer: The server you are planning to use. Verify that your application name is selected and associated with the correct WebServer.

   virtual host: Choose the default_host.

6. Save your configuration.

7. Configure the Classloader settings and change the polling interval to a number appropriate for your environment.

   Change Classloader order to have the classes loaded with parent classloader last.

   NOTE Do this only for the specific web application. Do not change the similar settings for the entire application server.

8. In the Manage Modules area, configure the Web Module Classloader setting.

   Change Classloader order to have the classes loaded with parent classloader last.

   NOTE Do this only for the specific web application. Do not change the similar settings for the entire application server.

9. (WebSphere 6.1) If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Map security roles to users/groups, and verify that the **All Authenticated** column is checked for the "All Authenticated" role.

10. (WebSphere 7.0) If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Security roles to user/group mapping. Select the “All authenticated” role and map it to “All Authenticated in Applications realm”.

11. Set permissions for the user running the application server.

    The user that will be running the application server must have read/write permissions on the following (default) folders:

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

    AE_install_path

12. Save all your changes.

13. Stop and restart WebSphere.

14. Start Workplace (or whatever you named your application) from the administrative console.

Task 6b: Deploy Application Engine (WebLogic)

This topic covers the deployment of your Application Engine application (Workplace) on WebLogic.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.bat (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.bat (Windows) from the following location:

     AE_install_path/deploy/

To deploy as “Workplace” or custom name using a WAR file

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users enter to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the app_engine.war file to reflect the name you want to use using the format *Application Name*.war.

Example:

The default app_engine.war will generate the following context root:

```
http://server_name:port#/app_engine
```

Renaming the WAR file Workplace.war will generate the following context root:

```
http://server_name:port#/Workplace
```

CAUTION You must rename the WAR file every time you regenerate it. The create_app_engine_war.sh/.bat script will by default create a file with the name app_engine.war.

To deploy Application Engine

1. From the WebLogic Administration Console, navigate to the domain you initially created for the Application Engine.
2. Prepare the WebLogic Administration Console to deploy the application.
3. Choose whether to deploy from an exploded folder (*AE_install_path*) or from the WAR or EAR file (default: app_engine.war or app_engine.ear in *AE_install_path*/deploy).
4. Accept the defaults for the deployment, except for the name for the deployment. Use “Workplace” instead of “appengine”.
5. Finish the deployment, and save and activate your changes.

   NOTE To verify that the deployment was successful, expand **Web Applications**. The web application Workplace will be listed.

6. After deployment is complete, start Workplace (or your custom application name) in the WebLogic Administration Console.

Task 6c: Deploy Application Engine (JBoss)

This topic covers the deployment and start of your Application Engine application (Workplace) on JBoss.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.exe (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.exe (Windows) from the following location:

     AE_install_path/deploy/

To deploy as “Workplace” or custom name using a WAR file

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users type to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the app_engine.war file to reflect the name you want to use using the format *Application Name*.war.

Example:

The default app_engine.war will generate the following context root:

```
http://server_name:port#/app_engine
```

Renaming the WAR file Workplace.war will generate the following context root:

```
http://server_name:port#/Workplace
```

CAUTION You must rename the WAR file every time you regenerate it. The create_app_engine_war.sh/.exe script will by default create a file with the name app_engine.war.

To deploy and start Application Engine

1. Deploy the Workplace application:

   To deploy from exploded directory

   a. On the JBoss server, copy the /Workplace folder from:

      AE_install_path

      to:

      JBOSS_home/server/default/deploy/

   b. Append the extension .war to the Workplace folder:

      JBOSS_home/server/default/deploy/Workplace.war

   To deploy from a WAR file

   On the JBoss server, copy the app_engine.war file from:

   AE_install_path/deploy

   to:

   JBOSS_home/server/default/deploy/

   To deploy from an EAR file

   On the JBoss server, copy the app_engine.ear file from:

   AE_install_path/deploy

   to:

   JBOSS_home/server/default/deploy/

2. Set permissions for the user running the application server.

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the following folders:

   NOTE For Windows this is only required for NTFS formatted partitions:

   JBOSS_home/server/default/deploy/app_engine.war/.ear

   AE_install_path

3. Start or restart the JBoss application server.

4. Verify that the application deployed successfully.

   Verify that the server.log file located in *JBOSS_home*/server/default/log lists deployment of the WAR or EAR file you used.

Configuration and startup tasks

To configure the IBM FileNet P8 Platform components

NOTE You can perform the following configuration and startup tasks listed below in any order.

- "Set Application Engine bootstrap preferences" on page 148.
- "Create a Process Engine Isolated Region" on page 153.
- "Create a Process Engine Connection Point" on page 154.
- "Configure the Process Engine Connection Point for Application Engine" on page 155.
- "Set up Content Engine and client transport SSL security" on page 157.
- "Set up Application Engine SSL security" on page 160.
- "Perform additional configuration tasks" on page 163.
- Familiarize yourself with IBM FileNet P8 system startup and shutdown procedures. See the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup.

Set Application Engine bootstrap preferences

Bootstrap preferences are a category of site preferences. The first time you sign into Workplace after Application Engine installation, the Bootstrap Preferences page opens.

Bootstrap Preferences

The following six bootstrap preference groups are available the first time you sign in:

- Security Info (required for SSL only)
- User Token Settings
- Preference Settings (required)
- Banner Image
- Application Integration
- Administrator Access Role

For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site Preferences > Bootstrap preferences.

Enhanced Timezone Detection

In addition to these settings you can also set the *useEnhancedTimeZoneDetection* parameter to accurately detect a client browser's time zone. This setting cannot be modified through the Site Preferences page. To enable this feature you must manually modify the bootstrap.properties file. For more information, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > bootstrap.properties.

NOTES

- By successfully signing in to Workplace and saving the bootstrap preferences, you are verifying the Application Engine's basic functionality such as user authentication as well as communication and storing of data in Content Engine.
- In addition to the preferences covered in this topic, more preferences can be set for the Workplace application using Workplace Site Preferences. For more information, see the the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences.
- After the initial bootstrap configuration, users with the Application Engine Administrators role can change any of these preferences by signing into Workplace and navigating to Admin > Site Preferences > Bootstrap.
- When you access the bootstrap preference page via the Site Preferences application, an additional preference, **Guest info** (to allow guest sign ins), is also available.
- In a web farm/clustered environment, all Application Engines share site preferences by using the same bootstrap.properties file. For more information, see the *IBM FileNet P8 Platform High Availability Technical Notice*. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

- Bootstrap properties are stored in a separate file from other site preference. The default location for this file, bootstrap.properties, is:

  AE_install_path/FileNet/Config/AE

- (New installations only) To allow users to create workflows subscriptions, you *must* configure the PWDesigner access role. For more information, see "(New installations only) To enable user access to the Workflow Subscription Wizard" on page 152.

To set the bootstrap properties on first login

1. Sign in to Workplace:

   a. On any computer, open a browser and type:

   `http://ApplicationEngineServerName:port#/Workplace`

   NOTE *ApplicationEngineServerName* cannot be 'localhost' or an IP address.

   b. Enter a user name and password, and then click **Sign in**. The Bootstrap Preferences page opens.

   NOTE The user who initially logs in and sets the bootstrap preferences is automatically added to the Application Engine Administrators role. For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences.

2. Enter security info (required for SSL only).

   a. Enter the SSL Host and Port information for the SSL server.

   b. Enter the Java Server HTTP port.

   Use the Security info preference to redirect sign-ins through a Secure Socket Layer (SSL) server and to identify a remote Java server. This encrypts the user IDs and passwords when they travel over the network. See "Set up Application Engine SSL security" on page 160 for instructions on setting up SSL security for one or more Application Engines.

   CAUTION Once you've configured SSL, the Site Preferences application also runs under SSL to protect the guest account's user ID and password. This means that when you run Site Preferences on an unsecured server that redirects sign-ins to an SSL server, you will be editing the Bootstrap preferences of the SSL server (stored in the bootstrap.properties file). This does not affect the General, Object Store, and Shortcut preferences, which are retrieved from the preferences file saved in the object store.

3. Configure user token settings.

   User Tokens are used by IBM FileNet P8 applications to launch into each other without the need for an additional login.

   a. Select whether or not to create user tokens for your Application Engine (Default: **Yes**).

   b. Select whether or not the application will pick up generated tokens from other applications (Default: **Yes**).

   c. Specify a Token timeout interval (1 - 15 minutes).

4. (Required) Specify preference settings.

   Preference settings specify the name of the site preference file, its storage location (object store), and the documentation server URL(if installed). The site preferences file is checked into the object store the first time you log on to Workplace. The site preferences are stored as XML data in this file, <Site Preferences for *Preferences name*>.xml. Use Enterprise Manager to access this file, and navigate to Object Stores > *Object Store location* > Root Folder > Preferences.

   NOTE The bootstrap preferences are saved in the bootstrap.properties file, and not in the site preferences file.

   a. Select an object store from the **Object store location** choice list. The preferences file will be saved in this object store. Workplace users must have access to this object store.

   b. Enter a preference file name in the **Preferences name** field.

   c. Enter the documentation server URL in the **Documentation server** field.

      The format of the URL is:

      `http://DocServerName:port#/ecm_help/`

      where *DocServerName* is the name of your Java application server where the documentation is installed,
      port# is the port number,
      and *ecm_help* is the root directory of the documentation web site.

      NOTE If no documentation URL is specified, the Workplace Help link will default to `http://localhost`.

   d. Enter the ISRA Interface Servlet URL.

      For more information, see "Enable Application Engine to use ISRA" on page 184.

5. Set Banner Image.

   The banner image is the graphic that appears at the upper left -hand side of the Workplace application pages. If you have a banner image that you would like to use in place of the default, follow this procedure.

   a. Copy the new graphic file to the location of your choice on Application Engine in the /FileNet/AE/Workplace folder.

   b. In the Path field, type the path (relative to the /Workplace folder) to the new banner graphic file.

   c. In the Image Width field, type the width of the image (in pixels).

   d. In the Image Height field, type the height of the image (in pixels).

6. Configure Application Integration.

   Select **No** (default), if you do not want users to be prompted to add an email to an object store each time the email is sent.

   Select **Yes**, if you want users prompted to add an email to an object store when the email is sent.

   This preference setting only affects Outlook integration.

7. Add Application Engine Administrators.

   Add the users and groups that will perform Application Engine administration tasks to the *Application Engine Administrators* role.

   NOTES

   - The user who initially signs in and sets the bootstrap preferences is automatically added to the Application Engine Administrators role. For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences.
   - To modify the access roles after the initial bootstrap configuration, users with the Application Engine Administrators role can use the access roles section of the Workplace Site Preferences. Launch Workplace and navigate to Admin > Site Preferences > Access Roles.

8. Click **Apply** to save your bootstrap settings.

To verify that a single index has been added for Application Name on the site preferences object store

To properly index access roles and improve login performance on Application Engine, an index is created for Application Name on the object store that contains the Workplace site preferences. Verify this index setting after you have successfully configured the bootstrap preferences.

1. On Content Engine, launch the Enterprise Manager.
2. In the left pane, expand the **Object Stores** folder.
3. Expand the object store node where your preferences are stored. See "(Required) Specify preference settings." on page 150 above.
4. Expand **Other Classes** and then **Custom Object**.
5. Right click **Access Role** and select **Properties**.
6. Select the **Property Definitions** tab.
7. Select **Application Name** and click **Edit**.
8. On the General tab of the Application Name Properties, verify that the **Indexed** field shows **Single Indexed.**
   - If the **Indexed** field shows **Single Indexed.**, continue at Step 13.
   - If the **Indexed** field shows 'not indexed', continue at Step 9.
9. Click **Set/Remove**.

10. Select **Set** and then **Single Indexed**.
11. Click **OK** to set the index.
12. Click **OK** to apply the change and close the Application Name Properties window.
13. Click **OK** to close the Access Role Class Properties window.

(New installations only) To enable user access to the Workflow Subscription Wizard

To allow users to create workflows subscriptions, you must configure the PWDesigner access role using the Workplace Site Preferences, and give the users appropriate access rights to the workflow subscriptions classes. You can perform these steps in any order, and you must perform both steps any time you need to add or remove users.

1. Assign users as members of the PWDesigner access role. See the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences.
2. Run the security script wizard,and load the workplacescript.xml file to add accounts to the Workflow Designer role.

   For more information about how to use the Security Script wizard to assign the Workflow Designer role to user or group accounts, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content Engine Wizard Help > Security Script.

   For more information about the workplacescript.xml file and how roles are defined in the Enterprise Manager, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Managing Security > Security Script Wizard.

Create a Process Engine Isolated Region

Process Engine communicates to its database using a connection point. Each connection point is associated with an isolated region. In this task you will create an isolated region. In "Create a Process Engine Connection Point" on page 154 you will define a connection point to this isolated region.

To create a Process Engine isolated region

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn on the desktop, or by navigating to Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn. Log on as a GCD administrator.
2. Connect to the FileNet P8 domain you created in "Establish the FileNet P8 domain and Global Configuration Data (GCD)" on page 56.
3. Right-click **PE Region ids > New PE Region ids**.
4. Click **Next** on the Specify a Site screen to select a site named *initial site*.
5. Enter the DNS name for the Process Engine server.
6. Enter the region ID.
7. Modify the communication port if needed. This communication port must match that entered in the Process Task Manager for the communication port in "Configure Process Task Manager" on page 105.
8. Click **Next** when done.
9. Enter the password for the isolated region. This password must match that entered in the Process Task Manager for the isolated region in "Configure Process Task Manager" on page 105.
10. Click **OK** on the Confirmation Window.
11. Click **Finish** to finish create a new region for Process Engine.

Create a Process Engine Connection Point

A connection point identifies a specific isolated region of the workflow database, and gives it a name that workflow-related applications use to access the region. Follow these procedures to create a connection point.

To create a Process Engine connection point

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn on the desktop, or by navigating to Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn. Log on as a GCD administrator.
2. Connect to the FileNet P8 domain you created in "Establish the FileNet P8 domain and Global Configuration Data (GCD)" on page 56.
3. Right-click **PE Connection Points** > **New PE Connection Points**.
4. Enter a Process Engine Connection Point name and click **Next.**
5. Choose the region which is created in "Create a Process Engine Isolated Region" on page 153, and click **Next.**
6. Click **Finish** to finish creating the Connection Point.
7. Click **OK.**

Configure the Process Engine Connection Point for Application Engine

Before users can access tasks and work items from Workplace, you must configure the Connection Point on the Application Engine. Make sure that you have already completed these Tasks:

- "Create a Process Engine Isolated Region" on page 153
- "Create a Process Engine Connection Point" on page 154

To configure the connection point

1. Log in to Workplace as an Application Engine Administrator:
   a. On any computer, open a browser and navigate to:

      `http://ApplicationEngineServerName:port#/Workplace`

   b. Log in using the same account that you used to set the bootstrap preferences.
2. Click **Admin**.
3. Click **Site Preferences**.
4. Under General Settings > Tasks, select a **Process Engine Conncetion Point** from the drop down list.
5. Click **Apply**.
6. Click **Exit**.

CAUTION Performing Step 7 in an existing environment will destroy all data in the existing region.

7. (For new installations only) Initialize the isolated region.
   a. Click **Admin**.
   b. Click **Process Configuration Console**.

      NOTE If your computer does not have the appropriate Java Runtime Environment (JRE) installed, you will be prompted to download the JRE at this point; follow the prompts to complete the download. During the installation process, click the **Browser** tab and enter the following settings:

      - De-select (clear) the Internet Explorer option.
      - If you will be using Netscape 6.0, select the Netscape option.

   c. Right-click the **icon** or **name** of the isolated region you want to initialize, and select **Connect** from the context menu.
   d. Click **Action**.

e. Click **Initialize Isolated Region**.

f. Click **Yes** at the prompt asking if you want to continue.

g. Close the Process Configuration Console.

8. In Workplace, click **Tasks** to confirm that Application Engine is communicating with Process Engine.

9. Sign out of Workplace.

Set up Content Engine and client transport SSL security

Configuring SSL enables secure communications between the Content Engine and the directory service, as well as between Content Engine clients and the Content Engine server. In addition, setting up Content Engine SSL provides secure authentication for Process Engine.

CAUTION IBM strongly recommends enabling SSL for the Content Engine and Process Engine web services. Authentication over these two web services is usually performed by providing username and password credentials. If these web services are not configured to run over an SSL connection, clear text passwords will be sent across the network. (However, this is not true when Kerberos-based authentication is used. In the P8 4.5 release, Kerberos authentication is available only for the Content Engine web service.) The option not to use SSL over these two web services is provided primarily for development systems or other non-production systems where the security provided by SSL may not be required.

For access to the Content Engine through the EJB transport (IIOP or T3 protocol), an SSL connection is necessary to provide privacy for data sent across the network, but user passwords would not be compromised if SSL were not used. While it is preferable to use SSL with the EJB transport (IIOP or T3 protocol), it is not a requirement.

NOTES:

- The Content Engine web service is used:
  - By all clients of the Content Engine 4.5 .NET API
  - By all clients of the Content Engine 4.5 COM Compatibility API (CCL)
  - By the Enterprise Manager tool
  - By the Content Engine 3.5.2 to 4.5.0 Upgrade tool
  - By the Process Engine, when making calls to the Content Engine to retrieve user and group information
  - By the Component Manager, running on the Application Engine, which is an integral component for BPM Process Orchestration framework
  - By customer and 3rd party tools written against the CE 3.5 web service API, including Altien Document Manager and the Sharepoint integration done by Vorsite.
- Certain Java applications (written against the Content Engine 3.5.x Java API or the Content Engine 4.5 Java API) may use the Content Engine web service transport, but typically they would use EJB transport (IIOP or T3 protocol).
- The IBM FileNet Application Engine server will use only the EJB transport to communicate with the Content Engine in the P8 4.5 release.
- The Process Engine web service is used by customer and third-party applications to write runtime applications (typically step processor applications) against the Process Engine. The Process Engine Java API does not make use of the Process Engine web service.

To enable SSL for Content Engine

NOTE In the steps below, a server certificate certificate will be added to the Directory Services server (for authentication). In addition, the CA certificate will be added in two different locations on the Content Engine server (the JDK path location is for authorization). Follow the steps closely to ensure that the proper certificate is added to each of the three locations.

1. Obtain and install a server certificate and a CA certificate on the directory service. These certificates are available from third-party certificate authorities, such as VeriSign, or you can generate your own certificates if you have the necessary certificate management software installed.
2. Enable SSL on the directory service and set the SSL port number. The default SSL port number is 636; however, if you have more than one directory service that is using SSL on the server, you may need to use a non-default port number. See your directory server documentation for instructions.
3. On the Content Engine server, add the CA certificate to the application server keystore, if it does not already contain it.
4. On the Content Engine server, add the CA certificate to the JDK (Java) keystore, if it does not already contain it. You can use the default key store, in Step a, or create a custom location, in Step b.
   a. To use the JDK default java key store, do the following:
      i. Determine the java version your application server uses and the JAVA_HOME location.
      ii. Use the keytool to import the CA certificate to the Java keystore at %JAVA_HOME%\jre\lib\security\cacerts.
      iii. To improve security, change the default password.
   b. To use your own key store (rather than the JDK default key store), do the following:
      i. Add the following system parameters to the Java command line in your application server's startup script:

         ```
         -Djavax.net.ssl.trustStore=path_to_your_keystore_file
         -Djavax.net.ssl.trustStorePassword=password_of_your_keystore
         ```

      ii. Use the Java keytool to import the CA certificate to your own keystore.
5. Use Enterprise Manager to enable SSL for Content Engine and set the port number to match the SSL port on the directory server, as described in "To enable SSL between Enterprise Manager and the directory service" on page 159.
6. Obtain another server and CA certificate for the Content Engine.
7. Create a custom identity keystore on the Content Engine server, and add the server certificate to the custom keystore.

8. Using the application server administration tool, enable SSL and point to the custom identity keystore. Directions vary by application server type; see your application server documentation for detailed procedures.

   NOTE (WebLogic only) The name in your certificate must match the host name specified in your WebLogic application server. If the name in the certificate is fully qualified (for example, Host1.filenet.com), the same fully qualified host name must appear in the Host field (WebLogic > Authentication Provider > Active Directory tab > Host field).

9. Configure clients to use a particular URL for connecting to Content Engine based on the application server type and the client transport (protocol) type. The following table provides details:

Protocol	SSL	Port	App Server	Sample URL
HTTP	no	7001	WebLogic	http://mycorp.com:7001/wsi/FNCEWS40DIME/
HTTPS	yes	7002	WebLogic	https://mycorp.com:7002/wsi/FNCEWS40DIME/
T3 (IIOP)	no	7001	WebLogic	t3://mycorp.com:7001/FileNet/Engine
T3S (IIOP)	yes	7002	WebLogic	t3s://mycorp.com:7002/FileNet/Engine
HTTP	no	9080	WebSphere	http://mycorp.com:9080/wsi/FNCEWS40DIME/
HTTPS	yes	9403	WebSphere	https://mycorp.com:9403/wsi/FNCEWS40DIME/
IIOP	no	2809	WebSphere	iiop://mycorp.com:2809/FileNetEngine
IIOP	yes	9443	WebSphere	iiop://mycorp.com:9443/FileNetEngine

   NOTE The port values in the table above are default values. If you change the port that your application server listens on, you must also change the port number used by the Content Engine client.

To enable SSL between Enterprise Manager and the directory service

1. Launch Enterprise Manager and log on as a GCD administrator.
2. In the tree view, right-click the root node and choose Properties.
3. In the Enterprise Manager Properties dialog box, click the Directory Config. tab, select a directory service, and click **Modify**.
4. In the General tab of the Modify Directory Configuration dialog box, set the Is SSL Enabled parameter to True and modify the port number appropriately.
5. Click **OK** in each open dialog box.

Set up Application Engine SSL security

This topic describes how to configure an Application Engine to direct sign-ins through a Secure Socket Layer (SSL) https connection. It assumes that Application Engine(s) have already been installed.

IBM FileNet Application Engine supports the following methods of configuring an SSL environment:

- Full SSL support - A single Application Engine server, where all of the software is running under SSL.
- One server SSL redirect - One Application Engine server set up to redirect logon attempts on the non-SSL port to the SSL port.
- Two server SSL redirect - Two Application Engine servers, where one is SSL-enabled, and the other redirects users to the SSL-enabled Application Engine server to log on.

To set up full SSL support on a single Application Server

1. Enable SSL on the application server that runs Application Engine (see your SSL documentation).
2. Test the SSL connection by signing into Workplace using one of the following URLs:

   `https://Application_Engine_server_name:SSL port/Workplace`

   The entire sign-in process will be handled by the SSL-enabled host.

   For more information about SSL port numbers, see “IBM FileNet P8 Port Numbers” in the *Plan and Prepare Your Environment for IBM FileNet P8 Platform* guide.

To set up SSL redirect on a single Application Engine server

1. Enable SSL on the application server that runs Application Engine (see your SSL documentation).
2. Sign in to Workplace:
   a. On any computer, open a browser and type the following URL address:

      `http://Application_Engine_server_name:port#/Workplace`

   b. Sign in as a user with Application Engine Administrator access role privileges. For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences.
3. Set bootstrap preferences:
   a. Navigate to Admin Site Preferences > Bootstrap.
   b. Set the Security info Site Preference SSL Host:Port to identify the alias host name and port number.

      Use the IP address of the Application Engine server for the SSL Host entry.

      For more information, see “Enter security info (required for SSL only).” on page 149.

c. Click **Apply** to save your bootstrap settings.

4. Update the base URL:

   a. Navigate to Admin > Site Preferences > Refresh.

   b. Enter the Workplace Base URL value in the provided field. The URL must contain a valid host name, and not contain “localhost" or an IP number. For example, http://myserver:7001/Workplace

      For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Refresh preferences.

   c. Click **Refresh** to update the base URL.

   d. Click **Exit** to close Site Preferences.

5. Sign out of Workplace, and close your browser.

6. Test the SSL connection by signing into Workplace using the following URL:

   `http://Application_Engine_server_name:non-SSL port/Workplace`

   NOTE You will be redirected to the SSL-enabled port for sign in, then back to the non-SSL enabled port after sign-in is complete. Before sign-in, you should receive a warning that you are accessing pages over a secure connection (unless you turned this dialog box off), and then Workplace will open.

To set up SSL redirect on two Application Engine servers

1. Install Application Engine on both computers so that both Application Engines use the same bootstrap.properties file and site preferences file (the Setup program will prompt you for a shared location).

   During setup of the first Application Engine, create a share on the folder where the bootstrap.properties file is installed (the \WEB-INF folder). Then during setup of the second Application Engine, specify the shared location from the first installation. The bootstrap.properties file must already exist when specifying a shared location. See “Setup WebLogic clusters” or “Setup WebSphere clones” in the *IBM FileNet P8 Platform High Availability Technical Notice* for specific instructions. To download this guide from the IBM support page, see “Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 19.

   CAUTION The system clocks on the two Application Engine servers must be synchronized to within the Token time-out interval. For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Bootstrap Preferences > User token settings.

2. Copy the UTCryptokeyFile.properties file.

   For SSL redirect to work, each Application Engine must use the same User Token cryptographic key file.

   After installing the second Application Engine, copy the UTCryptoKeyFile.properties file from the first Application Engine server to the same location on the second Application Engine server.

   NOTE IBM recommends copying the file over a secure link.

3. Enable SSL on the application server that you are using for the SSL-enabled Application Engine (see your SSL documentation).
4. Sign in to Workplace on the non-SSL enabled Application Engine.
   a. On any computer, open a browser and type:

      `http://ApplicationEngineServerName:port#/Workplace`

   b. Sign in as a user with Application Engine Administrator access role privileges. For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences.
5. Set bootstrap preferences:
   a. Navigate to Admin > Site Preferences > Bootstrap.
   b. Set the Security info Site Preference SSL Host:Port to identify the alias host name and port number.

      For more information, see "Enter security info (required for SSL only)." on page 149.
   c. Click **Apply** to save your bootstrap settings.
6. Update the base URL:
   a. Navigate to Admin > Site Preferences > Refresh.
   b. Enter the Workplace Base URL value in the provided field. The URL must contain a valid host name, and not contain localhost or an IP number. For example, http://myserver:7001/Workplace

      For more information, see the IBM FileNet P8 help topic User Help > Actions, preferences, and tools > Site preferences > Refresh preferences.
   c. Click **Refresh** to update the base URL.
   d. Click **Exit** to close Site Preferences.
7. Sign out of Workplace, and close your browser.
8. Test the SSL connection by signing into Workplace using the following URL:

   `http://Application_Engine_server_name:non-SSL port#/Workplace`

   NOTE You will be redirected to the SSL-enabled server for sign in, then back to the non-SSL enabled server after sign-in is complete. Before sign-in, you should receive a warning that you are accessing pages over a secure connection (unless you turned this dialog box off), and then Workplace will open.

Using Java Applets in an SSL Environment

If you are using a Java applet in an SSL environment, you may experience an SSLHandshakeException because the appropriate certificate does not exist on your computer. Follow the instructions in the the IBM FileNet P8 help topic User Help > Using Workplace > Basics > Use Java applets to resolve this issue.

Perform additional configuration tasks

Once you have completed the Installation Tasks, your core IBM FileNet P8 system will be up and running. Below is a list of additional configuration tasks you should complete (or at least review) to prepare the system for general use. Except where noted, the links go to the IBM FileNet P8 Help, and start from the Contents panel in:

<Documentation URL, in the form http://*webserver:port*#/ecm_help>/_start_here.htm

- Configure Content Federation Services for Image Services Guidelines. Refer to 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.
- Configure Application Engine to set the file types you want to open in a browser window rather than using the Image Viewer. Refer to FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > content_redir.properties file.
- Set site preferences for the Workplace application. Refer to User Help > Actions, preferences, and tools > Site preferences.
- Design searches and/or search templates for Workplace users. Refer to User Help > Actions, preferences, and tools > Tools > Search Designer > About Search Designer.
- Design publishing templates for Workplace users. Refer to User Help > Actions, preferences, and tools > Tools > Publishing Designer > About Publishing Designer.
- Configure security for publishing. Refer to User Help > Actions, preferences, and tools > Tools > Publishing Designer > Security > Specify publication document security.
- Configure automatic workflow launch. Refer to FileNet P8 Administration > Content Engine Administration > Events and subscriptions > Concepts: workflow subscriptions.
- Create and configure the object stores that will contain business objects, folders, documents, workflow definitions, searches, and other objects. Refer to FileNet P8 Administration > Content Engine Administration > Object stores > How to... > Create object store.
- Define document classes and folders and set security for each class. Refer to FileNet P8 Administration > Content Engine Administration > Classes > Concepts.
- Review and, if necessary, edit the security of the network shared folders containing any file stores created for the object store. Refer to FileNet P8 Administration > Content Engine Administration > Content storage > File storage areas.
- Configure Process Engine for automatic startup. Refer to FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Process Service > Start and stop Process Service > Configure the Process Service for automatic startup (Windows).
- Configure email notification. Refer to FileNet P8 Administration > Process Engine Administration > Workflow administration tasks > Coordinating workflow design > Email notification.

  NOTE Process Engine supports localized email notification. For details on configuring this, see the *IBM FileNet P8 Non-English Support Guide*. To download this document from the IBM FileNet support Web site, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

- Set Process Engine runtime options. Refer to User Help > Integrating workflow > Process Configuration Console > VWServices > View or modify VWService properties > Set runtime options.
- Set the default date/time mask for the Process Service. Refer to Process Engine > Process Task Manager > Process Service > Configuring Process Service > General properties.
- Create content cache area. Refer to FileNet P8 Administration > Content Engine Administration > Content storage > Content cache areas > How to... > Create content cache.
- Create additional authentication realms. Refer to FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms.
- Define additional isolated regions. Refer to User Help > Integrating workflow > Process Configuration Console > Isolated regions.
- For each isolated region:
  - Define workflows. Refer to User Help > Integrating workflow > Process Designer.
  - Configure event logging options. Refer to User Help > Integrating workflow > Process Configuration Console > Isolated regions > View or modify isolated region properties > Configure event logging options.
  - Configure step processors. Refer to User Help > Integrating workflow > Process Configuration Console > Isolated regions > View or modify isolated region properties >Configure custom step processors.
  - Define and configure work queues. Refer to User Help > Integrating workflow > Process Configuration Console > Queues > Configuring work queues.
  - Define and configure component queues. Refer to User Help > Integrating workflow > Process Configuration Console > Queues > Configuring component queues.
  - Define and configure workflow rosters. Refer to User Help > Integrating workflow > Process Configuration Console > Queues > Rosters.

Optional installation tasks

To install optional IBM FileNet P8 components

NOTE You can install the additional or optional IBM FileNet P8 components listed below in any order.

- "Install and Configure IBM FileNet Publishing" on page 166.
- "Enable the Process Engine Component Integrator" on page 167.
- "Install an Additional Instance of Enterprise Manager" on page 170.
- "Create Additional File Storage Areas" on page 171.
- "Install IBM FileNet Deployment Manager" on page 172.
- "Install Application Integration" on page 174.
- "Install File Tracker" on page 178.
- "Deploy multiple Application Engine instances" on page 181.
- "Enable Application Engine to use ISRA" on page 184.
- "Install and Configure IBM FileNet System Manager" on page 189.
- "Modify an Autonomy K2 server configuration" on page 190.
- "Install the COM compatibility layer (CCL)" on page 196.

Install and Configure IBM FileNet Publishing

Install the IBM FileNet Rendition Engine to establish publishing capabilities. For instructions, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide* at FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade.

Enable the Process Engine Component Integrator

Via the Component Integrator functionality included in the IBM FileNet P8 Platform, a step in a workflow can access properties of documents, folders, and other objects in an object store. Using this functionality requires configuration on both Application Engine and Process Engine servers, as described in this task.

As a post-installation task, you will also have to define workflows that incorporate Content Engine (CE) operations in order to use the out-of-the-box Component Integrator functionality. For further details on defining such workflows, see the IBM FileNet P8 help topic Steps > Component Steps > General Properties > Using Content Engine (CE) operations in a workflow.

Once the software is installed, users can extend the out-of-the-box Component Integrator functionality so that a workflow step can interact with an external entity such as a Java object or Java™ Messaging Service (JMS) messaging system. For further information, see the IBM FileNet P8 help topic Developing Process Applications > Developing Work Performers / Component Integrator Operations > Developing Component Integrator-Based Workflow Applications.

To specify the user name and password for the Java adaptor (on an Application Engine server)

1. Sign in to Workplace.

   If you defined the Process Engine Configuration Group on the Security tab of Process Task Manager (when completing "Configure Process Task Manager" on page 105), you must sign in as a member of either that group or the Process Engine Administrators Group, which was also defined on the Security tab, in order to complete the following steps.

2. In Workplace, click **Admin** and then click **Process Configuration Console**.

   NOTE If your computer does not have the appropriate Java Runtime Environment (JRE) installed, you will be prompted to download the JRE at this point; follow the prompts to complete the download. During the installation process, click the **Browser** tab and clear the Internet Explorer option.

   For further information about the JRE download, click **Help** in Process Configuration Console, click **Process Reference** on the help page toolbar, and see the IBM FileNet P8 help topic Concepts > Java Runtime Environment (JRE).

3. Select the Isolated Region icon that corresponds to the isolated region you initialized in "Create a Process Engine Isolated Region" on page 153.

4. Right-click the **CE_Operations** component queue and select **Properties**.

5. On the Adaptor tab of the displayed dialog box, enter a user name and password that will be used for identification and permissions for both Process Engine (PE) and potentially any external systems that will be accessed. By default, the user name and password are set to *Administrator* and *<no password>*, respectively. If you choose to use another user name and password, they must already exist in the directory service.

   For additional information about the fields on the Adaptor tab, click the **Help** button. To use the out-of-the-box functionality, it is necessary to modify only the user name and password fields.

6. Click **OK** and commit the changes.

To configure and start the Component Manager (on an Application Engine server)

Execute Step 1 below if Application Engine is configured to use maximum strength encryption. The JRE used to run the Process Task Manager that contains the Component Manager (which uses JDK 1.4.x) must be updated with Unlimited Strength Jurisdiction Policy Files. Otherwise proceed to Step 2 below.

1. Install unlimited strength JAR files.

   NOTE Perform this step only if you are using JDK 1.4 or higher and have selected the **Create unlimited strength keys** option in the *Application Engine User Security* and/or *User Token Security* steps of the Application Engine Setup program. Failure to perform the step will cause EncryptionException messages or other errors indicating that a Java Security API provider for Blowfish is not available. The EncryptionException is caused by the wrong versions of (or absence of) required JAR files that provide unlimited strength security policy files in a Sun JDK 1.4 or higher environment.

   For more information, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Application Engine Security.

   a. Obtain the JDK version-specific unlimited strength JAR files, as follows:

      - For the IBM JDK, obtain the IBM unlimited jurisdiction policy files from the IBM web site (http://www.ibm.com/developerworks/java/jdk/security).
      - For the Sun JDK, obtain the Sun unlimited strength policy files from the Sun product web site (http://java.sun.com/j2se/).

      CAUTION Make sure you install JAR files specific to the JDK version you are using.

   b. Install the files into the JRE's /jre/lib/security folder by replacing files with the same names.

   c. Restart the application server.

2. Start Process Task Manager on the Application Engine server.

   Launch the Process Task Manager by running one of the following command files from the *AE_install_path*/FileNet/AE/Router directory, depending on your operating system:

   UNIX

   routercmd.sh

   Windows

   routercmd.bat

   NOTE If the port number assigned to Component Manager conflicts with the port number required by another application or service running on the Application Engine server, then Process Task Manager will not start up as expected. See "IBM FileNet P8 ports" on page 265 for details on how to resolve this condition.

3. Select **Component Manager** in the left pane (also referred to as the feature pane).

4. Right-click and select **New** to define a new connection point. You will be prompted to enter the Content URI, Service Username, and Service Password to authenticate to the Content Engine server.

5. Enter or modify the component properties as appropriate. For details, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Application Engine > Component Manager > Configure the Component Manager -> General.

   NOTE In an environment configured for single sign-on (SSO), do not use the SSO server name in the URL, even if Process Task Manager displays it by default.

6. Click **Start** on the toolbar.

To specify connection between Process Engine and Component Manager (on a Process Engine server)

1. On the Process Engine server, start Process Task Manager as follows, depending on your operating system:

   Windows

   Select Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager.

   UNIX

   Enter the following command on the command line:

   ```
   vwtaskman
   ```

2. Select Process Engine in the left pane (also referred to as the feature pane).
3. In the Component Manager connection section, select the Server Connections tab.
4. In the Host field, enter the host name of the Application Engine server where Component Manager is running.
5. In the Event Port field, enter the port that the Component Manager listens to for incoming events. The default is 32773. The port number you enter must match the number you entered in Step 3 of "To configure and start the Component Manager (on an Application Engine server)" on page 168.

Install an Additional Instance of Enterprise Manager

Do this task only if you want to install an instance of Enterprise Manager in addition to the one you installed in "Install FileNet Enterprise Manager" on page 54.

CAUTION Do not install Enterprise Manager 4.5 or the COM Compatibility Clients API on any machine running the 3.5.x version, at least until the Content Engine 4.5 upgrade is complete. Otherwise, you will no longer be able to run Enterprise Manager 3.5.x against any remaining 3.5.x object stores. You may need to use Enterprise Manager to prepare for upgrading the 3.5.x object stores.

NOTE You can install Enterprise Manager only on a Windows machine, and only using the Windows version of the Content Engine installation media.

To install an additional instance of Enterprise Manager,

1. If you have not already done so, install Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0. Enterprise Manager on the Windows machine where you are going to install Enterprise Manager.
2. Do a silent or interactive installation of Enterprise Manager as shown in "Install FileNet Enterprise Manager" on page 54.
3. Install any service packs, fix packs and/or interim fixes required. To determine whether such additional software updates are needed, contact your service representative.

Create Additional File Storage Areas

Perform this task to create additional file storage areas for existing object stores. To create additional fixed storage areas, navigate instead to the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content storage > Fixed storage areas.

Complete the following procedures for each file storage area you want to create.

To create a file storage area

1. Prepare a location for the file storage area, as shown in "Prepare storage areas for object stores" on page 53 in *Plan and Prepare Your Environment for IBM FileNet P8*, and then continue at Step 2.
2. Start Enterprise Manager.
3. Select a FileNet P8 domain and log on as an administrator of the object store in which you will create a file storage area.
4. Right-click the **Storage Areas node** and then choose **New Storage Area**.
5. When the Create a Storage Area wizard opens, click **Next** and complete the wizard screens as shown in the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content Engine Wizard Help > Create a Storage Area.

Before storing content in the file storage area, do the following procedure to verify that it was properly created.

To verify the file storage area

1. Log on to the machine where Content Engine Server is installed.
2. List the contents of the *fsa1* directory you created or designated on the file server in one of the following procedures:
   - **UNIX** "To configure a UNIX-based file server" on page 56 in *Plan and Prepare Your Environment for IBM FileNet P8*
   - **Windows** "To configure Windows-based Content Engine Server to talk to a Windows file server via CIFS" on page 58 in *Plan and Prepare Your Environment for IBM FileNet P8*
   - **Windows** "To configure a Windows-based file server for a UNIX client using NFS" on page 57 in *Plan and Prepare Your Environment for IBM FileNet P8*
3. Verify that *fsa1* contains an XML file, named fn_stakefile, and two subdirectories, content and inbound.

   NOTE On UNIX machines, the content and inbound subdirectories must remain in the same file system.
4. Verify that *fsa1* has the ownership and access permissions you specified.

Install IBM FileNet Deployment Manager

IBM FileNet Deployment Manager works with FileNet Enterprise Manager to deploy test systems into full production. Use the procedure in this topic to install the IBM FileNet Deployment Manager interactively or silently on a Windows machine.

NOTE FileNet Deployment Manager runs only on Windows.

To install IBM FileNet Deployment Manager

1. On the machine where you will install FileNet Deployment Manager, log on as a member of the Local Administrators group or the Power Users group.
2. If you have not already done so, install Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0 on the machine.
3. Access the Content Engine installation software.
4. Start the FileNet Enterprise Manager installation. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select CE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - To install interactively:

     i. Run the following command in the software package:

     ```
     P8CE-4.5.0-WIN.EXE
     ```

     ii. When prompted, select Tools to be installed. If you want to also install FileNet Enterprise Manager on this machine, select .NET Clients as well.

     iii. Specify the install path, for example **C:\Program Files\FileNet\Content Engine**.

   - To install silently:

     i. Open the **CE_silent_install.txt** file in the software package for editing.

     ii. Set the parameter values in the **CE_silent_install.txt** file for your site. Be sure to set the CHOSEN_INSTALL_FEATURE_LIST parameter value to:

     ```
     Tools
     ```

     If you want to also install FileNet Enterprise Manager on this machine, select .NET Clients as well:

     ```
     DotNetClients,Tools
     ```

     iii. Set the LAUNCH_CM value to 0.

iv. Save your edits.

v. Run the following command in the software package:

```
P8-CE-4.5.0-WIN.EXE -f CE_silent_install.txt -i silent
```

NOTE To start FileNet Deployment Manager, choose **Start > All Program > FileNet P8 Platform > P8 Deployment Manager**.

Install Application Integration

Install Application Integration if you want to integrate IBM FileNet Workplace with your Microsoft Office applications and Outlook. Complete the following procedure on each machine that will use Workplace Application Integration.

NOTE You cannot collocate Workplace Application Integration with clients running IDM Desktop Application Integration.

Verify that the client machine meets the platform requirements documented 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 install the Application Integration software interactively

1. Log onto the client machine using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download Application Integration for Microsoft Office,** and then do one of the following:
   - Click **Open** to run the program from its current location.
   - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the ApplicationIntegration.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the installer.

   The Welcome Wizard dialog box for Application Integration appears. Another Welcome dialog box appears.
5. Click **Next**.
6. Read the license agreement, and then select **I accept the terms to the license agreement**, and then click **Next**. If you do not accept the license agreement, you cannot continue with the install.
7. Do the following:
   - Select the applications you want to integrate, and then click **Next**.

     NOTE The Application Integration Toolkit Components option is required to use Application Integration.
   - Under **Install to**, the default installation path is displayed. Click **Change** to specify a different location on the Change Current Destination Folder dialog box, and then click **OK**. Click **Next**.

     NOTE You may see two default installation paths - one for Microsoft Office and Outlook, and another for the Toolkit Components. The Toolkit Components path only appears when the system on which you are installing Application Integration has the Toolkit Components currently installed. You cannot modify the Toolkit Components installation path.

8. Enter the server name, port number and application name that defines the Workplace address. The *server name* is the name of the web server running Workplace, *port number* is the web server's assigned port, *application* is the directory where you installed the IBM FileNet Workplace application files.

   Check **Server uses secure connection (SSL)** if you are running full SSL to encrypt all communication with Workplace.

   NOTE You can also leave these fields blank and enter the information when you log on to Workplace Application Integration.

9. Click **Next**.
10. Click **Install**.
11. After the install is complete, click **Finish** to complete the setup process.

To install the Application Integration software silently

1. Log onto the client machine using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download Application Integration for Microsoft Office,** and then click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the ApplicationIntegration.exe file locally, and then click **Save**.
5. Open a DOS command window and change the current directory to the one where ApplicationIntegration.exe resides.
6. Type the following at the command line:

   ```
   ApplicationIntegration.exe /s/v"/qn <additional msi arguments included in string>
   LICENSE_ACCEPTED=true"
   ```

   Use the /s switch to launch the execution silently and include the /qn switch in the msi string to make msi run silently.

   Refer to the following optional command line values you can also use. Append the values within the string containing the msi arguments. For example:

   ```
   ApplicationIntegration.exe /s/v"/qn /L*v C:\temp\AppIntSetup.txt
   LICENSE_ACCEPTED=true"
   ```

Command Line Values	Installs
ADDLOCAL=ALL	All Features
ADDLOCAL=ALL REMOVE=OutlookIntegrationFeature	Office Only
ADDLOCAL=ALL REMOVE=OfficeIntegrationFeature	Outlook Only
ADDLOCAL=ALL REMOVE=OutlookIntegrationFeature, OfficeIntegrationFeature	Core Only

Command Line Values	Settings
HOST=<host name>	Enter the name of the web server running Workplace.
PORT=<port number>	Enter the web server's assigned port number.
APPLICATION=<application name>	Enter the directory in which you installed the Workplace application files.
SERVER_CONNECTION=1	Set Application Integration to use an https connection
SERVER_CONNECTION=0	Set Application Integration to use http connection. This is the default if this parameter is not passed.
/L*v C:\temp\AppIntSetup.txt	Verbose installation log and specify log location.

To verify your Workplace Application Integration installation

1. Start Microsoft Word.
2. From the **File** menu, click **FileNet P8**, click **Open Document**, and then click **Select Item**. The Logon dialog box opens.
3. Log on using any valid domain account. The available object stores in your environment are displayed.

   NOTE If you did not enter the Workplace Address information in Step 8 in the procedure "To install the Application Integration software interactively" on page 174, enter the server name, port number and application name that defines the Workplace address. The *server name* is the name of the web server running Workplace, *port number* is the web server's assigned port, *application* is the directory where you installed the IBM FileNet Workplace application files.

Check **Server uses secure connection (SSL)** if you use a full SSL to encrypt all communication with Workplace. Do not select this option if you use a SSL redirect during logon.

4. Close all dialog boxes and close Microsoft Word.

To uninstall or modify Workplace Application Integration

1. From the **Start** menu, click **Settings**, and then click **Control Panel**.
2. Click **Add/Remove Programs**, and then click **FileNet Workplace Application Integration 4.0**.
3. Do one of the following:
   - Click **Remove**, and then click **Yes** to confirm you want to uninstall Workplace Application Integration.
   - Click **Change** to access maintenance tasks, and then click **Next**. You can modify, repair, or remove Application Integration using the maintenance tasks.

     Do one of the following:

     – Select **Modify** to add or remove integration with Microsoft applications from your previous install. For example, if you have both Microsoft Office and Outlook installed, you can remove one of the applications using this option. The Custom Setup dialog box appears, where you highlight the option you want to add or remove. Click **Next**, and then click **Install**. Click **Finish** to complete the process.

     – Select **Repair** to re-install Workplace Application Integration to repair installation errors, and then click **Next**. Click **Install** to start the repair process. Click **Finish** to complete the process.

     – Select **Remove** to remove Workplace Application Integration from your system, and then click **Next**. Click **Remove**. Once the application is removed from your system, click **Finish** to complete the process.

To silently uninstall Workplace Application Integration

1. Open a command prompt.
2. Enter the following command to uninstall Workplace Application Integration:

```
msiexec.exe /X{35907B7D-02E2-490C-8F3B-54C4E3729D90} /qn
```

Install File Tracker

Install File Tracker if you want to use the Workplace file tracking feature without installing Application Integration. Complete the following procedure on each machine that will use Workplace File Tracker.

NOTES

- If you have already installed or upgraded to Application Integration 3.5.1-002 or higher, then the File Tracker feature has already been installed. Do not perform this procedure if you already installed a version of Application Integration that includes File Tracker, including Application Integration 4.0.
- If you have already installed Application Integration 3.5.1-001 or earlier, upgrade it to 4.0 plus the latest fix pack, and this will include the File Tracker installation. For details about the upgrade process, see "Upgrade Application Integration and File Tracker" on page 382.

Verify that the client machine meets the platform requirements documented 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 install the File Tracker software interactively

1. Log on to the client machine using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download File Tracker** and do one of the following:
   - Click **Open** to run the program from its current location.
   - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the FileTracker.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the installer.

   The Welcome Wizard dialog box for File Tracker appears. Another Welcome dialog box appears.
5. Click **Next**.
6. Read the license agreement, and then select **I accept the terms to the license agreement**, and then click **Next**. If you do not accept the license agreement, you cannot continue with the install.
7. Do one of the following:
   - Click **Change** if you want to install File Tracker to a different location. Specify the location to which you want to install File Tracker, and then click **OK**. Click **Next**.
   - Click **Next** to accept the default location.
8. Click **Install**.
9. After the install is complete, click **Finish** to complete the setup process.

10. (On Vista using Internet Explorer 7.0 only) Add the Workplace URL to the browser security tab.
    a. In the Internet Explorer 7.0 browser, click **Tools** > **Internet Options**. Click the **Security** tab.
    b. Select Trusted Sites, and click the **Sites** button.
    c. Add the Workplace URL and click **Add**.
    d. Click **OK**, and **OK** again to save changes.

To install the File Tracker software silently

1. Log on to the client machine using an account that has Administrator privileges and sign in to Workplace.
2. Click **Author**, and then click **General Tools**.
3. Scroll down and click **Download File Tracker** and click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the FileTracker.exe file locally, and then click **Save**.
4. Open a DOS command window and change the current directory to the one where FileTracker.exe resides.
5. Type the following at the command line:

   ```
   FileTracker.exe /s /v"/qn <additional msi arguments included in string>
   LICENSE_ACCEPTED=true"
   ```

   NOTE Use the /s switch to launch the execution silently and include the /qn switch in the msi string to make msi run silently. In addition, be aware of spaces specified between switches in the above example. Using the correct spacing ensures a successful silent install.

   Refer to the following optional command line values you can also use. Append the values within the string containing the msi arguments.

   For example:

   ```
   FileTracker.exe /s /v"/qn /L*v C:\temp\FileTrackerSetup.txt LICENSE_ACCEPTED=true"
   ```

Command Line Values	Settings
/L*v C:\temp\FileTrackerSetup.txt	Verbose installation log and specific log location. **NOTE** If you intend to specify a log location, create the directory before running the silent install. If the directory is not created ahead of time, the install will fail.

6. (On Vista using Internet Explorer 7.0 only) Add the Workplace URL to the browser security tab as follows:
   a. In the Internet Explorer 7.0 browser, click **Tools** > **Internet Options**. Click the **Security** tab.
   b. Select Trusted Sites, and click the **Sites** button.
   c. Add the Workplace URL and click **Add**.

d. Click **OK**, and **OK** again to save changes.

To uninstall Workplace File Tracker

1. From the **Start** menu, click **Settings**, and then click **Control Panel**.
2. Click **Add/Remove Programs**, and then click **FileNet Workplace File Tracker**.
3. Click **Remove**, and then click **Yes** to confirm you want to uninstall Workplace File Tracker.

To silently uninstall Workplace File Tracker

1. Open a command prompt.
2. Enter the following command to uninstall Workplace File Tracker:

   msiexec.exe /X{4291FBBC-C585-43ED-9416-5F22D8C6FEE9} /qn

Deploy multiple Application Engine instances

This topic covers deployment of multiple instances of Workplace on a single application server. Each deployment of Workplace must use the same Content Engine, Process Engine, and connection point. Each deployment of Workplace may use different Site Preference settings and may provide access to different object stores.

NOTES

- The following procedure assumes that you have already installed Application Engine and performed the following configuration tasks according to your application server type:
  - "Configure Application Engine (WebSphere)" on page 122
  - "Configure Application Engine (WebLogic)" on page 132
  - "Configure Application Engine (JBoss)" on page 138
- When deploying multiple instances of Workplace, make copies of all the Workplace configuration and working files. Each instance of Workplace will use separate configuration, deploy, download, upload, and Workplace files. Leave the default installed files unmodified.
- For more information on how to deploy and manage multiple identical applications, see your application server documentation.

To deploy a second instance of the Workplace application

1. Make a copy of the /FileNet/Config/AE directory, including all of its contents, for each instance you plan to deploy.

   For example, if you are deploying two instances, you would create:

   install_path/FileNet/Config/AE1

   install_path/FileNet/Config/AE2

2. Make copies of the upload and download directories in the <*install_path*>/FileNet/AE directory.

   For example, you would create:

   install_path/FileNet/AE/Download1

   install_path/FileNet/AE/Upload1

   install_path/FileNet/AE/Download2

   install_path/FileNet/AE/Upload2

3. Make a copy of the deploy directory and all of its contents for each Workplace instance.

   For example, you would create:

   install_path/FileNet/AE/deploy1

   install_path/FileNet/AE/deploy2

4. Make a copy the Workplace directory and all of its contents for each Workplace instance.

   For example, you would create:

   install_path/FileNet/AE/Workplace1

   install_path/FileNet/AE/Workplace2

5. Navigate to each custom copied Workplace web.xml instance and update the path for the configuration directory, upload directory, and download directory locations.

   For example, in the *install_path*/FileNet/AE/Workplace1/WEB-INF/web.xml, you would make the following changes (in **bold**):

```
<context-param>
   <param-name>configurationDirectory</param-name>
   <param-value>/opt/FileNet/Config/AE1</param-value>
 </context-param>

<context-param>
   <param-name>uploadDir</param-name>
   <param-value>/opt/FileNet/AE/Upload1</param-value>
 </context-param>
 <context-param>
   <param-name>downloadDir</param-name>
   <param-value>/opt/FileNet/AE/Download1</param-value>
 </context-param>
```

To deploy each additional Workplace instance as an EAR file

NOTE Perform the following steps for each custom Workplace instance you plan to deploy.

1. Modify the application.xml file located in the copied deploy directories, as follows:

   a. Open each instance of the application.xml file, for example, <*install_path*>/FileNet/AE/deploy1/META-INF/application.xml.

   b. Change the <display-name> and the <context-root> elements to your custom name, for example, Workplace1 (shown in **bold**, below).

```
<display-name>Workplace1</display-name>
<description>FileNet Application Engine</description>


<module>
   <web>
      <web-uri>app_engine.war</web-uri>
      <context-root>Workplace1</context-root>
   </web>
</module>
```

2. In the create_app_engine_war file, change the *install_home* path and the *deploy* directory to match your custom names.

   For example, you would make the following changes, shown in bold:

```
install_home="/opt/FileNet/AE/Workplace1"

"${install_home}/../_AEjvm/bin/jar" -cf "${install_home}/../deploy1/app_engine.war"*
```

3. In the create_app_engine_ear file, set the install home, deploy directory, and EAR file to match your custom names.

   For example, you would make the following changes, shown in **bold**:

   ```
   install_home="/opt/FileNet/AE/Workplace1"
   cd "${install_home}/../deploy1"
   "${install_home}/../_AEjvm/bin/jar" -cvf "${install_home}/../deploy1/
   app_engine1.ear" META-INF *.war
   ```

4. Delete the existing app_engine.war and app_engine.ear files.
5. Create your custom WAR and EAR files by running the create_app_engine_war and then the create_app_engine_ear files.
6. Deploy the EAR file for each custom Workplace instance according to the procedures for your application server type, as presented in the following topics:
   - “Deploy Application Engine (WebSphere)” on page 141
   - “Deploy Application Engine (WebLogic)” on page 143
   - “Deploy Application Engine (JBoss)” on page 145
7. Install any service packs, fix packs and/or interim fixes required. To determine whether such additional software updates are needed, contact your service representative.

Enable Application Engine to use ISRA

Image Services Resource Adapter (ISRA) is a J2EE connector to the FileNet Image Services (IS) libraries. Using ISRA, Workplace users can view IS documents and their associated annotations in the FileNet P8 Image Viewer and, if they have the appropriate permissions, update the annotations.

To enable Workplace users to view documents using ISRA, the following steps must be completed:

- Install Application Engine.
- Install FileNet ISRA.

  For information on installing, configuring and deploying FileNet ISRA, refer to the Image Services Resource Adapter documentation in the FileNet ISRA installation package.

  TIP Use the Sample Application shipped with FileNet ISRA to confirm that the ISRA installation was successful.

  WARNING In an ISRA upgrade situation, take care to use the same library name (JNDI connection factory name) that has been previously set in the ISRA install. Changing this variable can cause conflicts when accessing documents.
- Install the Application Engine ISRA Servlet and take the following into account:
  - Install and deploy ISRA before installing and deploying the ISRA Servlet.
  - Deploy the ISRA Servlet on the same application server as FileNet ISRA.
  - You need not necessarily install the ISRA Servlet on the Application Engine server. See "ISRA SSL Support" on page 184 for details that might affect your collocation plans.
- Configure Workplace Site Preferences.

ISRA SSL Support

The following table details supported SSL configurations for ISRA.

SSL Configuration	SSL Support
ISRA Servlet and Application Engine Collocated. Application Engine configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet and Application Engine Collocated. Application Engine Configured for SSL logon redirect to a local host.	Supported
ISRA Servlet and Application Engine Collocated. Application Engine and ISRA Servlet running under SSL.	Not Supported

SSL Configuration	SSL Support
ISRA Servlet remote from Application Engine. Application Engine configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet remote from Application Engine. Application Engine configured for SSL logon redirect to a local host.	Supported
ISRA Servlet remote from Application Engine. Application Engine running under SSL, ISRA Servlet not running under SSL.	Supported
ISRA Servlet remote from Application Engine. Application Engine and ISRA Servlet running under SSL.	Not Supported

Install and Deploy the Application Engine ISRA Servlet

Use the procedure in this topic to install and deploy the ISRA Servlet on the operating systems supported by Application Engine by running the associated ISRA setup program found in the Application Engine software package.

To install and deploy the Application Engine ISRA Servlet

The FileNet P8 Application Engine installation software contains the ISRA servlet installation programs for the supported P8 AE operating systems.

1. Log on to the application server using the following account, depending on your operating system:

   UNIX

   User account with write access to the /bin directory and read, write, and execute access to the directory where you plan to install ISRA Servlet.

   Windows

   A member of the local Administrators group or as a user account with equivalent permissions.

2. Stop the application server if it is running.

3. Access the ISRA installation package, and start the following Application Engine ISRA Servlet setup program, depending on your operating system:

   UNIX

   AE-ISRA-Servlet-4.0.2.0-*operating_system*.bin

   Windows

   AE-ISRA-Servlet-4.0.2.0-WIN.exe

4. Complete the installation program screens. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   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 the ISRA servlet installation program:

   – Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select ISRA installer.

   – Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

5. Check the file AE_ISRA_Servlet_install_log-4_0_2_0.txt, located in the *AE_israservlet_install path*\FileNet directory, to see if any errors occurred during the installation.

6. Install unlimited strength JAR files.

   Perform this step only if the following are true:

   - You selected the option to create Strong keys in the Application Engine User Token Security step of the Application Engine installation.
   - Application Engine ISRA Servlet is deployed on a different application server than Application Engine.

   CAUTION If these conditions are true, failure to perform this step causes an EncryptionException when you log on to the IS server.

7. (WebSphere and WebLogic only) Start the application server.

8. Deploy *AE_israservlet_install_path*/FileNet/ApplicationEngineISRAServlet/ae_isra.war in the same way you deployed the app_engine.war file for Workplace.

9. (JBoss only) Start the application server.

10. Verify the Application Engine ISRA Servlet is installed and deployed correctly, as follows. This step launches a diagnostic tool that does the verification.

    a. Open your web browser.

    b. Enter the URL for the Application Engine ISRA Servlet, for example:

    ```
    http://ApplicationEngineISRAServlet_servername:port/
    ApplicationEngineISRAServlet/ISRA
    ```

    NOTE ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

    If the ISRA Servlet is set up correctly, a congratulations message displays, for example:

    ```
    Congratulations! ISRA Interface Servlet is configured at this URL.

    WcmApiConfigFile = D:\ISRAInterface\jsp\WEB-INF\WcmApiConfig.properties

    WcmApiConfig file exists
    ```

```
CryptoKeyFile/UserToken = C:\Program
Files\FileNet\Authentication\UTCryptoKeyFile.properties

CryptoKeyFile/UserToken exists


FileNet ISRA classes are in the classpath

com.filenet.is.ra.cci.FN_IS_CciConnectionSpec
```

Configure Workplace Site Preferences

Use the following procedure to enable the Image Services external service and set the ISRA Interface Servlet URL in Workplace Site Preferences. The Application Engine setup installs the pre-configured Image Services external service, which includes the parameterized values necessary to access FileNet IS libraries from Workplace.

To configure Workplace Site Preferences for ISRA Servlet support

1. Sign in to Workplace as a user having the Application Engine Administrators access role.
2. Launch Site Preferences as follows:
   a. Select **Admin**.
   b. Select **Site Preferences**.
3. Enable the pre-configured Image Services external service, as follows:
   a. Select **External Services** from the left options list.
   b. Select **Modify** for the Image Service, located under External Reference Services.

      The External Reference Service Settings site preference page displays.
   c. Under General Information, locate *Show on Select File page* and change the value to **Show**.
   d. Accept the setting.
4. Set the ISRA Interface Servlet URL as follows:
   a. Select **Bootstrap**.
   b. Under Preferences Settings, set the value of ISRA Interface Servlet URL. For example:

      `http://servername:port/ApplicationEngineISRAServlet/ISRA`

      NOTE ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.
   c. Accept the setting.
   d. Exit Site Preferences.

Log On to Image Services Using an LDAP Account

To log on to the Image Services library using your LDAP account, configure ISRA and Image Services for LDAP authentication.

NOTE If the LDAP account with which you accessed Workplace is not valid for the Image Services library, or if LDAP authentication is not configured, you will be prompted to log on to the Image Services library.

For information on configuring LDAP authentication for ISRA, refer to the *ISRA Installation and Deployment Guide*. For information on configuring LDAP authentication for Image Services, refer to the *Image Services System Tools Reference Manual*.

Access IS Library Documents

For information about accessing IS library documents, see User Help > Actions, preferences and tools > Actions > Documents > Add a document (Workplace).

Install and Configure IBM FileNet System Manager

Content Engine, Application Engine, and Process Engine install, by default, the necessary software required for the System Manager performance component. To use System Manager, enable associated components and install IBM FileNet Dashboard to perform related configuration procedures to enable System Manager. Installing Dashboard is not necessary if you currently have IBM FileNet System Monitor installed.

Refer to the following IBM FileNet P8 help topic FileNet P8 Documentation > FileNet P8 Administration > Enterprise-wide Administration > System Manager for instructions on how to enable the associated System Manager components.

Refer to the documentation provided with IBM FileNet Dashboard for instructions on how to use Dashboard.

Modify an Autonomy K2 server configuration

Add additional K2 Administration Servers

You will likely need to install additional Autonomy K2 Administration Servers to handle the indexing load. The Master Administration Server is the main hub for configuring all the servers you install. The additional servers are managed through the K2 Dashboard that is installed with the Master Administration Server.

NOTE A K2 Master Administration Server must be installed before you add additional K2 Administration Servers.

To install additional Autonomy K2 Administration Servers on Windows

1. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.

   NOTE Ensure *k2_os_user* has administrator privileges on this machine.

2. Set the JAVA_HOME environment variable as follows:

   a. Open the System control panel.

   b. Click the **Advanced** tab.

   c. Click **Environment Variables**.

   d. Click **New** under System Variables.

   e. Set the variable information as follows:

      Variable name: *JAVA_HOME*

      Variable value: *Java (JDK) install path*

3. Install the Content Search Engine software, using the appropriate values from your worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - Interactively

     i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-WIN.EXE file.

     ii. On the Select Autonomy K2 Server Type panel, choose Administration Server.

- Silently
  i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.
  ii. Launch the Content Search Engine installer by executing the following command:

```
P8CSE-4.5.0-WIN.EXE -f CSE_silent_install.txt -i silent
```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 6.2.1 Administration Server service.
- Verity K2 Administration Web Server service (Tomcat server).

To install additional Autonomy K2 Administration Servers on UNIX

1. Configure permissions for the host machine so that permissions of the user who runs the IBM FileNet Content Search Engine installation are root. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic. If *k2_os_user* can not have root privileges on this machine, you can set permissions as follows, to allow *k2_os_user* to run the installation:
   a. Access the host machine and log on as a user with root privileges.
   b. Enter the following commands to set the vspget program's sticky bit such that the service runs as root and is in the same group as *k2_os_user* (default path shown):

```
chown root /opt/verity/k2/_rs6k43/bin/vspget
chmod g+rs /opt/verity/k2/_rs6k43/bin/vspget
```

2. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.
3. Set the following environment variables and place the entries in the .profile file for *k2_os_user*.

   All UNIX types

```
JAVA_HOME=java_(JDK)_install_path/jdkversion
export JAVA_HOME
```

   HP-UX

```
PATH=$PATH:/verity_install_path/k2/_hpux/bin
export PATH
SHLIB_PATH=$SHLIB_PATH:/verity_install_path/k2/_hpux/bin
export SHLIB_PATH
```

AIX

```
PATH=$PATH:/verity_install_path/k2/_rs6k43/bin
export PATH
LIBPATH=$LIBPATH:/verity_install_path/k2/_rs6k43/bin
export LIBPATH
```

Solaris

```
PATH=$PATH:/verity_install_path/k2/_ssol26/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ssol26/bin
export LD_LIBRARY_PATH
```

Linux

```
PATH=$PATH:/verity_install_path/k2/_ilnx21/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ilnx21/bin
export LD_LIBRARY_PATH
```

4. Install the Content Search Engine software, using the appropriate values from your worksheet.

   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 the Content Search Engine installation program:

   – Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select CSE installer.

   – Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - Interactively

     i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-*UNIX TYPE*.BIN file.

     ii. On the appropriate panel, choose Administration Server as the type of installation.

   - Silently

     i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.

     ii. Execute the following command:

     ```
     ./P8CSE-4.5.0-UNIX_type.BIN -f CSE_silent_install.txt -i silent
     ```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 Administration Server service.
- Verity K2 Administration Server service (Tomcat server).

To start or stop the Autonomy K2 Services on UNIX

To manually start or stop the Autonomy K2 services, use the following commands, according to your environment:

HP-UX

Start Services:

```
nohup /verity_install_path/k2/_hpux/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_hpux/bin/k2adminstop
```

AIX

Start Services:

```
nohup /verity_install_path/k2/_rs6k43/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_rs6k43/bin/k2adminstop
```

Solaris

Start Services:

```
nohup /verity_install_path/k2/_ssol26/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ssol26/bin/k2adminstop
```

Linux

Start Services:

```
nohup /verity_install_path/k2/_ilnx21/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ilnx21/bin/k2adminstop
```

Configure Additional K2 Administration Servers

If you add additional K2 Index Servers or Search Servers to an existing configuration, you must enable them through Enterprise Manager to utilize them.

1. Log on to Enterprise Manager as the GCD Administrator and expand the Enterprise Manager tree.

2. Open the Index Area folder.
3. Right-click the index area that you want to add the new services to and select **Properties**.
4. Enable the new Search Servers as follows:
   a. Click **Edit Search Servers**.
   b. In the Search Servers Available pane, highlight any servers you want to enable for this index area and click **Add** to add the server to the Search Servers Selected list.
   c. Click **OK** to save the settings and enable the new servers.
5. Enable the new Index Servers as follows:
   a. Click **Edit Index Servers**.
   b. In the Index Servers Available pane, Highlight any servers you want to enable for this index area and click **Add** to add the server to the Index Servers Selected list.
   c. Click **OK** to save the settings and enable the new servers.

Move K2 Servers to new hardware

To move K2 Master Administration Server software to new hardware

To move an existing installation of the Autonomy K2 Master Administration server configuration, the installation must be imaged and then restored on the new hardware. If any details of the new configuration do not match the previous configuration, the configuration will not work.

1. Back up your Autonomy K2 Master Administration Server configuration.
2. Configure the new machine.
3. Restore your Autonomy K2 Master Administration Server installation.

To move K2 Administration Server software to new hardware

Follow this procedure to maintain your existing configuration on new hardware. Alternatively, you can remove the server from your configuration (see procedure below) and then create a new configuration after installing the K2 Administration Server software.

To move your existing configuration to new hardware, your new hardware must be configured the same as the previous machine. All drives and shares must be named the same and each server must have the same host name and IP addresses.

1. Take careful note of all aspects of your current Content Search Engine configuration. Record all share names, host names, server ports, and IP addresses. Record the server names and ports of all configured server services in the K2 Dashboard, and their locations.
2. If collections are stored on this machine, make a backup copy.
3. Set up your new hardware and ensure the machine host name and IP address match your previous configuration.

4. Follow the Add Additional K2 Administration Servers procedure above to install the K2 Administration Server software on the new machine.
5. Create the services for this machine in the K2 Master Administration Server dashboard and ensure all server services are named the same as your previous configuration and ensure all servers created in the Autonomy K2 Dashboard are named the same and configured for the same ports.

   NOTE You can change the ports, but if you do so you'll need to reboot the CE server(s). The names of the search servers, index servers and brokers may not be changed for any items referenced by an index area on the CE.
6. Restore the collections to the same Root Dir Path (as specified in the index area properties) as before. If the collections are no longer available, mark those collections as unavailable, and then run a reindexing job on those collections. Otherwise copy/restore those collections via whatever tool you used to save them.
7. Modify the appropriate directory path entries in the Verity.cfg file.

To remove K2 Administration Servers from your Content Search Engine configuration

1. If collections were stored on the machine you want to remove:
   a. Determine which K2 Server will hold the collections.On that machine, create a new collections directory and temp directory and assign the appropriate permissions. For details, see “Create a Content Search Engine collections directory” on page 79.

      NOTE See the index area properties (Root Dir Path) for the collections location for the appropriate index area.

   b. Modify the appropriate directory path entries in the verity.cfg file.

   c. Move the collections to the new location.
2. Modify the index area to use the index servers and search servers of the machine to which the collections were moved.
3. Access the K2 Dashboard and remove the appropriate K2 Administration server from the dashboard configuration.

Install the COM compatibility layer (CCL)

The option to install the COM compatibility layer is available as part of the IBM FileNet Content Engine installer.

To install the COM Compatibility Layer (CCL) from the IBM FileNet Content Engine installer

1. In the Choose Components dialog, select .NET Client, then click Next.
2. In the .NET API COM Compatibility Layer (CCL) Server URL dialog, enter a valid URL for the CCL (for example, `http://localhost:9080/wsi/FNCEWS40MTOM/`). Note that if you do not enter a valid URL, the CCL will not be installed.

If you do not install the CCL during the initial installation, you have the option of installing later by running the CE installer again. You can also install the CCL anytime by using the Configuration Manager tool.

Upgrade and configure IBM FileNet P8 Platform

This section contains the following tasks:

Upgrade IBM FileNet P8 documentation

Peform the tasks in this section to upgrade your IBM FileNet P8 documentation installation. See the following topics:

Task 1: Upgrade IBM FileNet P8 documentation

Overview

There are two methods available for updating an existing IBM FileNet P8 Documentation installation:

- "Refresh IBM FileNet P8 documentation without uninstalling" on page 200

  Use this option if you have IBM FileNet P8 3.5.*x* or 4.0.*x* documentation installed without any expansion product help.

- "Update IBM FileNet P8 documentation by uninstalling and reinstalling" on page 201

  Use this option if you have IBM FileNet P8 3.5.*x* or 4.0.*x* documentation installed with expansion product help included (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services, or Records Manager).

  Optionally, as of the IBM FileNet P8 4.5 documentation release, you can install and index web-posted auxiliary documentation, such as the IBM FileNet P8 Release Notes, to have it included in the same search functionality. For a complete list of such documentation, see "Gather auxiliary documentation" in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

NOTES

- As of the IBM FileNet P8 4.0 release, you can deploy the IBM FileNet P8 documentation on a Content Engine server because it, like the Application Engine server, is a J2EE application server. If you are an existing 3.5.*x* customer and wish to install the 4.5 documentation on a Content Engine server or a new application server, you must perform a new installation rather than an upgrade. Therefore, use the install instructions in the following topic:
  - "Install IBM FileNet P8 Platform documentation (All)" on page 22
- The refresh procedure assumes the following:
  - You wish to replace the 3.5.*x* or 4.0.*x* documentation at the existing location with updated files.
  - You are familiar with your application server's procedures for reinstalling or redeploying web site applications.
- Before updating any IBM FileNet P8 documentation, review the *IBM FileNet P8 Hardware and Software Requirements* for the required software versions and updates for third-party software. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- Although some versions of BEA WebLogic support deployment of WAR-file-based web applications, you cannot deploy ecm_help.war on this application server platform. You must instead use the fully expanded IBM FileNet P8 documentation directory (*ecm_help*) structure. Otherwise, the searches within the IBM FileNet P8 documentation will not work and you will receive null pointer errors.

- In environments where Windows NTFS is used, there is a 256 character limit on file and directory names (directory depth). When trying to delete an existing IBM FileNet P8 documentation web site, you might encounter access denied errors. See Microsoft Knowledge Base article http://support.microsoft.com/?kbid=320081 for more information.
- Before installing the latest IBM FileNet P8 documentation, be sure to back up your existing IBM FileNet P8 documentation web site according to your site and application server procedures. This precaution will allow you to restore the IBM FileNet P8 documentation web site quickly if for any reason you have to back out of or delay your IBM FileNet P8 software upgrade.
- If you install any IBM FileNet P8 expansion products as part of your upgrade (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services, or Records Manager), be aware that:
  - You must copy the associated help for all expansion products onto the updated IBM FileNet P8 documentation server.
  - You must update the index for the help Search feature as documented below. This action ensures that searches return all expected results. If you add help for other expansion products later, you must re-run the procedure below for updating the help Search index.
  - You must also update the help Search index if you update the installed IBM FileNet-based help that you have customized or translated as part of your own application development.
- Any time you update the documentation search index, a backup of the files in the existing Index/core directory will be automatically copied to the Index/IndexOld subdirectory. You can reapply these backed-up files to the core subdirectory (after first removing the new files created there) if you need to return to your previous indexed state.
- Depending on your operating system (Windows or UNIX) and application server version (WebSphere, WebLogic, or JBoss), some screens will be slightly different than those documented in the procedures listed below.

Refresh IBM FileNet P8 documentation without uninstalling

You can refresh an existing IBM FileNet 3.5.*x* or 4.0.*x* P8 documentation installation by simply copying the newer documentation files over the existing files, and then reindexing for Search.

Use this option only if you have IBM FileNet 3.5.*x* or 4.0.*x* documentation installed without any expansion product help.

NOTES

- If you have FileNet 3.5.*x* or 4.0.*x* documentation installed with expansion product help included, use the procedure "Update IBM FileNet P8 documentation by uninstalling and reinstalling" on page 201.
- This refresh procedure requires that you copy the expanded ecm_help directory from the IBM FileNet P8 Documentation package over the existing documentation. This stipulation applies

to all application servers, even those, such as WebSphere, that required a WAR file for initial deployment of the existing 3.5.*x* or 4.0.*x* documentation.

To refresh the IBM FileNet P8 documentation without first uninstalling

1. Stop the application server (or IBM FileNet P8 documentation site) on which the existing documentation is deployed so that no processes can access the documentation.
2. On the IBM FileNet P8 documentation application server, locate the deployed IBM FileNet P8 documentation directory, and back up (or copy to a safe location) the files located in:

   UNIX

   deployment_path/ecm_help/search/index/core

   Windows

   deployment_path\ecm_help\search\index\core
3. Access the 4.*x* IBM FileNet P8 Platform documentation package.
4. Copy the expanded ecm_help directory from the package over the deployed ecm_help directory.
5. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM P8 Platform and various expansion products. Optionally, also download any auxiliary documentation to the ecm_help/installation/web directory on the IBM FileNet P8 Platform documentation server. Check the documentation page on the IBM Information Management support page on www.ibm.com for the latest versions of these guides. See "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19 for details.
6. Continue as follows:
   - If you added expansion product documentation, you will need to update the search index. Go on to the procedure in the following topic, "Update help search index" on page 204.
   - If you have no further documentation to install, and you did not install any expansion products, then go on to the procedure in the topic "Start and verify IBM FileNet P8 documentation web site" on page 205.

Update IBM FileNet P8 documentation by uninstalling and reinstalling

Use this option if you have IBM FileNet P8 3.5.x or P8 4.0.x documentation installed with expansion product help included, for example: Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services, or Records Manager.

Use the following procedures to completely remove an existing IBM FileNet P8 documentation site from the application server before updating the site. Follow the procedure for you application server type:

- "To update the IBM FileNet P8 documentation on WebSphere application servers" on page 202
- "To update the IBM FileNet P8 documentation on WebLogic application servers" on page 202

- "To update the IBM FileNet P8 documentation on JBoss application servers" on page 203

To update the IBM FileNet P8 documentation on WebSphere application servers

1. Verify the WebSphere application server is running, and then start the WebSphere administrative console.
2. From the WebSphere administrative console, **Uninstall** the existing documentation site.
3. From the initial install location, delete the IBM FileNet P8 documentation installed files, leaving the ecm_help directory in place.
4. Copy the IBM FileNet P8 ecm_help.war file from the IBM FileNet P8 documentation package to an appropriate location on the local hard drive.
5. From the WebSphere administrative console, deploy the *ecm_help.war* file using **ecm_help** as the Context Root.
6. For each expansion product, copy the expanded ecm_help directory from the package over the deployed ecm_help directory.

   NOTE Repeat this step for each of your expansion products. You can copy more than one expansion product documentation set to the documentation application server (for example, FileNet P8 eForms and FileNet Content Federation Services). When you are finished, you end up with one ecm_help directory structure containing the core documentation set and one or more sets of expansion product documentation files.
7. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM P8 Platform and various expansion products. Optionally, also download any auxiliary documentation to the ecm_help/installation/web directory on the IBM FileNet P8 Platform documentation server. Check the documentation page on the IBM Information Management support page for the latest versions of these guides. See "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19 for details.
8. Restart the IBM FileNet P8 documentation site.
9. Continue as follows:
   - If you have added expansion products you will need to update the search index add on documentation for expansion products, go on to the procedure in the following topic, "Update help search index" on page 204.
   - If you have no further documentation to install, then go on to the procedure in the topic "Start and verify IBM FileNet P8 documentation web site" on page 205.

To update the IBM FileNet P8 documentation on WebLogic application servers

1. Verify the WebLogic application server is running, and then start the WebLogic Administration Console.
2. From the WebLogic Administration Console, **Stop** the existing documentation site and **Delete** the IBM FileNet P8 documentation files, leaving the ecm_help directory in place.
3. Copy the *ecm_help* directory structure from the IBM FileNet P8 documentation package to the location of the original IBM FileNet P8 documentation.

4. For each expansion product, copy the expanded ecm_help directory from the package over the deployed ecm_help directory.

   NOTE Repeat this step for each of your expansion products. You can copy more than one expansion product documentation set to the documentation application server (for example, FileNet P8 eForms and FileNet Content Federation Services). When you are finished, you end up with one ecm_help directory structure containing the core documentation set and one or more sets of expansion product documentation files.

5. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM FileNet P8 Platform and various expansion products. Optionally, also download any auxiliary documentation to the ecm_help/installation/web directory on the IBM FileNet P8 Platform documentation server. Check the documentation page on the IBM Information Management support page for the latest versions of these guides. See "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19 for details.

6. From the WebLogic Administration Console, restart the IBM FileNet P8 documentation site.

7. Continue as follows:

   - If you have added expansion product documentation, you will need to update the search index. Go on to the procedure in the following topic, "Update help search index" on page 204.
   - If you have no further documentation to install, then go on to the procedure in the topic "Start and verify IBM FileNet P8 documentation web site" on page 205.

To update the IBM FileNet P8 documentation on JBoss application servers

1. Shut down the JBoss application server.

2. From the initial install location, delete the IBM FileNet P8 documentation installed files, leaving the ecm_help directory in place.

3. Delete the temporary working directory for the IBM FileNet P8 Platform documentation from the *JBoss_home*/work/MainEngine/localhost directory.

4. Access the IBM FileNet P8 Platform documentation package.

5. Copy the ecm_help directory from the package to the existing location on the application server (from where you just removed the old version of the directory).

6. For each expansion product, copy the expanded ecm_help directory from the package over the deployed ecm_help directory.

   NOTE Repeat this step for each of your expansion products. You can copy more than one expansion product documentation set to the documentation application server before continuing (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services, or Records Manager) so you end up with one ecm_help directory containing multiple sets of expansion product files added to it.

7. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM FileNet P8 Platform and various expansion products. Optionally, also download any auxiliary documentation to the ecm_help/installation/web directory on the IBM FileNet P8 Platform documentation server. Check the documentation page on the IBM Information Management

support page on www.ibm.com for the latest versions of these guides. See "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19 for details.

8. Continue as follows:
   - If you have added expansion products you will need to update the search index. Go on to the procedure in the following topic, "Update help search index" on page 204.
   - If you have no further documentation to install, then go on to the procedure in the topic "Start and verify IBM FileNet P8 documentation web site" on page 205.

Update help search index

Perform this procedure to update the Search index after you have installed all the IBM FileNet P8 Platform and expansion product documentation on a supported application server.

To update the help Search index

NOTE Perform this procedure only if you refreshed the core documentation, or you have installed expansion product (or customized application) help or auxiliary documentation onto your IBM FileNet P8 documentation application server. Otherwise, skip to "Start and verify IBM FileNet P8 documentation web site" on page 205.

1. Make sure that the server on which the documentation is deployed/installed is stopped, and that no processes are accessing the documentation.
2. Make sure you have copied the help for all your various expansion products, or optionally downloaded updated documentation and auxiliary documentation, to a designated application server location containing the IBM FileNet P8 help. Otherwise, you will have to repeat this procedure if you add new help later.
3. Open a command prompt or terminal on the application server.
4. From the command line, navigate to the search subdirectory under your ecm_help root directory.
5. Using a text editor, open the search-indexing script file that is appropriate to your application server operating system:

   UNIX

   indexFiles.sh

   Windows

   indexFiles.bat
6. If necessary, set the JAVA_HOME variable in the script file with the path to your JRE installation. Save your changes and close the text editor.
7. If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (chmod 755) to the indexFiles.sh file.

8. Run the updated search-indexing script file (UNIX: indexFiles.sh or Windows: indesFiles.bat.

   By default, the script file backs up the existing index files to indexOld, and then re-indexes all the help files starting from the root ecm_help directory.

   NOTE As you run the search-indexing script, you might notice periodic *Parse Abort* errors. You can ignore these error conditions, as they are benign and do not adversely affect the overall indexing process.

9. Go on to the procedure in the next topic "Start and verify IBM FileNet P8 documentation web site" on page 205.

Start and verify IBM FileNet P8 documentation web site

Perform this procedure after you have updated (and, if necessary, reindexed) the IBM FileNet P8 documentation on a supported application server.

To verify the IBM FileNet P8 documentation web site

1. Start the IBM FileNet P8 documentation server. Use the appropriate instructions provided with your application server.
2. Verify that the application server and the new IBM FileNet P8 documentation web site are running. From your web browser, access the URL for your web environment, using your documentation server name, and port number, as in these examples:

   WebSphere

   http://yourdocserver:9080/ecm_help/

   WebLogic

   http://yourdocserver:7001/ecm_help/

   JBoss

   http://yourdocserver:8080/ecm_help/

   The help system should open.

   NOTE Apply the URL for your application server when it is time to configure the online help location for the various IBM FileNet P8 components, either while running installation programs or later via site preferences settings (for example, in Workplace or Workplace XT).

Upgrade and configure Content Engine and Content Search Engine

Perform the following upgrade tasks that apply to your existing version (3.5.x or 4.0.x) of Content Engine and Content Search Engine.

For upgrades from 3.5.x

1. To complete the pre-upgrade configuration, perform Task 1 on page 208.
2. To upgrade ECM Centera SDK library files, perform Task 2 on page 213 only if Centera fixed content devices will be in your upgraded IBM FileNet P8 environment, and either of the following conditions apply:
   - Centera fixed content devices are not in your existing environment.
   - Centera fixed devices are in your existing environment, your existing version of Content Engine is 4.0, and you have not installed any fix packs.
3. To install and upgrade Content Engine software and configure the Content Engine instance, perform Task 3a on page 218.
4. To install Content Engine software updates that have become available since the initial release of version 4.5, perform Task 5 on page 238.
5. To install the latest Process Engine Client files on Content Engine servers, perform Task 6 on page 239.
6. To deploy the Content Engine instances, perform Task 7 on page 241.
7. To configure storage devices, perform Task 9 on page 245.
8. To upgrade standalone instances (if any) of FileNet Enterprise Manager, perform Task 10 on page 248.
9. To establish the FileNet P8 domain and Global Configuration Data (GCD), perform Task 11 on page 249.
10. To upgrade Content Search Engine software from version 3.5.x, perform Task 12a on page 251.
11. To upgrade Content Engine data, perform Task 13 on page 275.
12. To complete the upgrade of Content Search Engine software from version 3.5.x, do Task 14 on page 289.

For upgrades from 4.0.x

1. To verify that all P8 activity has completed, perform Task 1 on page 208
2. To upgrade ECM Centera SDK library files, perform Task 2 on page 213 only if Centera fixed content devices will be in your upgraded IBM FileNet P8 environment, and either of the following conditions apply:
   - Centera fixed content devices are not in your existing environment.

- Centera fixed devices are in your existing environment, your existing version of Content Engine is 4.0, and you have not installed any fix packs.

3. To upgrade Content Engine software, perform Task 3b on page 230.
4. To configure upgraded Content Engine instances, perform Task 4 on page 233.
5. To install Content Engine software updates that have become available since the initial release of version 4.5, perform Task 5 on page 238.
6. To install the latest Process Engine Client files on Content Engine servers, perform Task 6 on page 239.
7. To deploy upgraded Content Engine instances, perform Task 7 on page 241.
8. Perform Task 8 on page 244 if you are staging your IBM FileNet P8 upgrade over a period of time and not upgrading one or more of the other IBM FileNet P8 server components at this time.
9. To upgrade standalone instances (if any) of FileNet Enterprise Manager, perform Task 10 on page 248.
10. To upgrade Content Search Engine software from version 4.0.x, perform Task 12b on page 268.
11. To upgrade Content Engine data, perform Task 13 on page 275.

Task 1: Complete pre-upgrade Content Engine configuration

The procedures in this topic prepare Content Engine for an upgrade from its current version: 3.5.x or 4.0.x. Each procedure title indicates whether it applies to upgrades from 3.5.x, 4.0.x, or both.

To verify that all in-progress event actions have finished (3.5.x or 4.0.x)

1. Start Enterprise Manager if it is not already running.
2. In each object store to be upgraded, perform the following substeps:
   a. In the list view, under the object store icon, right-click the **Search Results** folder, and choose **New Search**.
   b. In the Content Engine Query Builder dialog box, choose **QueueItem** from the Select From Table list.
   c. Retain all default settings and click **OK**. (Click **Yes** at the prompt for a WHERE clause.)
3. If any event items do remain in the queue, you will see them in the Query Status dialog box. To remove unwanted items, perform the following substeps:
   a. Set up the same search as in Step 2; but this time select the **Delete** check box in the **Action** tab of the Search dialog box, before clicking **OK**.
   b. Click **OK** again to confirm the deletion.

To verify that publishing requests have been completed (3.5.x or 4.0.x)

1. Start Enterprise Manager if it is not already running.
2. In each object store to be upgraded, perform the following steps:
   a. Expand the Publishing folder and click **Queue** and choose **View > All** requests.
   b. Verify that the queue (right pane) of publishing requests is empty. If the queue is not empty, do the following:
      i. Wait until all publish requests in the In Queue state or In Work state are processed.
      ii. If any publish requests are in the In Error state, contact your publishing administrator for the appropriate action to take (such as retrying after correcting the error or just deleting the item).
3. In the list view, under the object store icon, right-click the **Search Results** folder, and choose **New Search**.
4. In the Content Engine Query Builder dialog box, choose **PublishRequest** from the Select From Table list.
5. Retain all default settings and click **OK**, and then click **Yes** at the prompt for a WHERE clause.
6. Manually delete any publishing requests displayed in the Query Status window.

To delete custom subclasses of the PublishRequest object (3.5.x only)

Before Upgrader Tool can upgrade an object store, you must delete any custom subclasses of the PublishRequest class, along with any instances of such subclasses, as shown in the following steps:

1. Start version 3.5.x of Enterprise Manager.
2. In the left pane of Enterprise Manager, open an object store, and navigate to **Other Classes > Publish Request**.
3. In the right pane, delete any custom subclasses of Publish Request and any instances of such subclasses.
4. Repeat Step 2 and Step 3 for each object store.
5. Exit from Enterprise Manager.

To verify that all in-progress Content Engine transactions involving content files have finished (3.5.x only)

Before upgrading Content Engine, you must verify that all pending content transactions, indexing requests, and fixed-content migration requests have completed. Use the Content Resource Manager utility, as shown in the following steps.

CAUTION All transactions and requests must have completed before you upgrade Content Engine; otherwise the upgrade will fail.

1. On each machine where Content Engine 3.5.x is installed, complete the following steps:
   a. Start File Store Service if it is not already running.
   b. Stop Object Store Service and Content Cache Service.
2. On a machine where Content Engine 3.5.x is installed, start the *ce_install_path*\RMU.exe utility.
3. In the Content Resource Manager Utility window, complete the following steps:
   a. Navigate to **Statistics > Transactions** and note the values of *Total prepared phase 1*, *Total committed*, and *Total aborted*.
   b. Navigate to **Statistics > Can't Do Queue** and note the value of *Current items in the can't do queue*.
   c. Add the values of *Total committed* and the *Total aborted*. If the sum equals the value of *Total prepared phase 1*, and the value of *Current items in the can't do queue* is zero, then all content transactions have completed.
   d. Navigate to **Indexing Service > Index Control**. If the value of *Current index queue files* is zero, then all pending indexing requests have completed.
   e. Navigate to **Fixed Content > Migration Queuing**. If the value of *Current number in queue* is zero, then all fixed-content migration requests have completed.
4. If all activity checked in Step 3 has completed, then stop File Store Service on each machine where Content Engine 3.5.x is installed, and continue at "To stop and disable all Content-Engine-

related services on all servers in the FileNet P8 domain (3.5.x only)" on page 210; otherwise, complete the following steps:

a. Exit from the Content Resource Manager utility.

b. Wait a few minutes, and then return to Step 1.

To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain (3.5.x only)

1. On each machine where version 3.5.x of Content Engine Server software is installed, log on with local administrator permissions.
2. Stop and disable the following services:
   - Apache2
   - Content Engine Content Cache Service
   - Content Engine File Store Service
   - Content Engine Object Store Service
   - FileNet Publishing HTML Plug-in Service
   - FileNet Publishing PDF Plug-in Service
   - Process Services Manager
   - Wasp Server for Java
3. Right-click the Apache services monitor in the Windows system tray and click **Exit**.

To verify the directory permissions on file stores (3.5.x only)

If the application server on which Content Engine Server 4.5 is to be deployed is running on Windows, then give the following users write permission to the root directory where a 3.5.x file store is located. (This procedure is necessary to allow conversion to a 4.5 file storage area during upgrade.)

- The user running the Upgrader tool (*config_mgr_user*)
- The user running the application server (*ce_appserver_admin*). For WebSphere on Windows, this user is a service log-on account)

For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8.*

To update the Vital Product Data (VPD) script (4.0.x and Solaris only)

Perform the following steps only if Content Engine Server is at version 4.0.1 on a Solaris machine, a non-root user installed the current version of Content Engine Server, and a different non-root user will upgrade Content Engine Server to version 4.5:

1. Navigate to the $HOME/InstallShield//Universal/FileNet/CE/Gen1/_vpddb/ directory.

2. Create a backup of the vpd.script file and then edit the vpd.script file as follows:
   a. Change all instances of $REG_CE_LOCATIONS$ to the same Content Engine Server installation directory. For example,

      `/opt/FileNet/ContentEngine`

   b. Change all the occurrences of $40Installeruser$ to the same Content Engine Server installation directory. For example,

      `/opt/FileNet/ContentEngine`

   c. Save your edits.

To create the webSphereDefaultIsolationLevel custom property (4.0.x and WebSphere only)

Perform this procedure only if all the following conditions apply to your IBM FileNet P8 environment:

- Content Engine is deployed on WebSphere.
- Your database type is Microsoft SQL Server.
- You intend to upgrade Content Engine from version 4.0.1.

This procedure creates the webSphereDefaultIsolationlevel custom property, which needs to exist for the GCD XA and non-XA data sources before you can upgrade Content Engine to version 4.5.

1. Log on to the WebSphere administrative console.
2. Navigate to **Resources > JDBC > Data sources**.
3. Click the **GCD non-XA data source** (for example, FNGCDDS)
4. In the Configuration tab, click **Custom properties** under Additional Properties.
5. Click **New**, and create a property as follows:
   a. Set the name to webSphereDefaultIsolationLevel.
   b. Set the value to 2.
   c. Set the type to java.lang.String.
6. Click **Apply** and save your changes to the Master configuration.
7. Perform Step 4 through Step 6 for the GCD XA data source (for example, FNGCDSXA).

To verify permissions on the Engine-DS.xml file (4.0.x and JBoss only)

If you deployed Content Engine Server 4.0.x on JBoss application server, perform the following steps:

1. Navigate to the directory containing the Engine-DS.xml file. The typical location for this file is:

   UNIX

   /opt/FileNet/ContentEngine/lib

Windows

C:\Program Files\FileNet\ContentEngine\lib

2. Ensure that the Engine-DS.xml file has both read and write permissions.

To back up the database (3.5.x or 4.0.x)

1. Use your existing database backup solution to back up the Content Engine object store databases.
2. Back up the GCD.
   - For upgrades from 4.0.x, backup the GCD database. FNGCDDB is the default name of the GCD database.
   - For upgrades from 3.5.x, make a backup copy of the 3.5.2 GCD file. The default location is C:\Program Files\FileNet\Content Engine\sysconfig\sysinit\sysinit.dat.

To shut down CFS Exporter (CFS 4.0 only)

If your system includes Content Federation Services (CFS) 4.0.x, you must shut down the CFS Exporter before upgrading Content Engine to 4.5. Leave the Exporter shut down until you upgrade CFS to 4.5. Refer to the *IBM FileNet Content Federation Services Installation and Upgrade Guide* for instructions.

Task 2: Install or update ECM Centera SDK library files

Perform the procedure in this task to update ECM Centera SDK library files only if Centera fixed content devices will be in your upgraded IBM FileNet P8 environment, and either of the following conditions apply:

- Centera fixed content devices are not in your existing environment.
- Centera fixed devices are in your existing environment, your existing version of Content Engine is 4.0, and you have not installed any fix packs.

To install EMC Centera SDK version 3.2 library files

1. Log on to the Content Engine Server machine with a user account that has the appropriate permissions to create folders and install files.
2. Back up or delete any existing EMC Centera SDK library files, which are located in the Default Destination Installation Location column of the following table:.

Operating System	Default Destination Installation Location
AIX	/usr/local/Centera_SDK
Solaris	/opt/Centera_SDK
Windows	C:\Centera_SDK
Linux	/usr/local/Centera_SDK
HP-UX	/opt/Centera_SDK
HP-UXi	/opt/Centera_SDK

3. The Centera directory in the Content Engine software package contains the Centera SDK version 3.2 installation files. As shown in the following table, copy the appropriate directory to a location on the Content Engine Server machine, such as /tmp (UNIX) or C:\Temp (Windows).

Operating System	Directory To Be Copied
AIX, Solaris, Windows	Copy the entire Centera directory
Linux	Depending on your version of gcc, copy one of the following directories: Centera/gcc3.3 Centera/gcc4
HP-UX 11i v1 (11.11)	Copy Centera/11i-v1-11.11
HP-UX 11i v2 (11.23)	Copy Centera/11iv1-11.23

4. On the Content Engine Server machine, navigate within the Centera directory (at its copied location) to the install subdirectory, which contains the installer script.
5. Run the installer script corresponding to the operating system on the Content Engine Server machine. On UNIX, the installer script will prompt you for the install directory. On Windows, specify the install directory, such as C:\Centera_SDK, on the command line.

   UNIX

   ```
   install.sh
   ```

   Windows

   ```
   install.bat C:\Centera_SDK
   ```
6. The installer script creates both 32-bit and 64-bit library directories, and puts them in a default installation directory, depending on your operating system (as shown in the following table). Accept or change the default when prompted by the script.

Operating System	Subdirectories of extracted EMC Centera SDK Directory	Description
AIX, Solaris, and Windows	lib	lib has the native library files that are to be installed.
Linux	../gcc3.3/lib ../gcc4/lib	
HP-UX	../11i-v1-11.11/lib	
HP-UXi	../11i-v1-11.13/lib	

To configure EMC Centera SDK environment variables for version 3.2

1. Locate the sample setup script on the Content Engine installation media. The file name of the sample setup script depends on your operating system:

   UNIX

   setCenteraLibPath.sh

   Windows

   setCenteraLibPath.bat
2. Modify the sample setup script as indicated in the following table:

   Note that the CENTERA_LIB_PATH variable needs to point to the sample script, not just the installation directory that contains the script.

For example, if you have a 64-bit AIX system, and you change the destination installation path (*install_path* in the table below) from:

/usr/local/Centera_SDK (the default)

to:

/usr/local/Centera/SDK3.2.607

then change the installation path of the AIX script to:

/usr/local/Centera/SDK3.2.607/lib/64

Note that the actual location is appended with either lib/32 or lib/64 because the installation script creates both 32-bit and 64-bit library directories, and places them inside the lib directory.

Operating System	Script Revisions

AIX	From:
	`CENTERA_LIB_PATH=/usr/local/Centera_SDK/lib/32`
	`LIBPATH=$LIBPATH:$CENTERA_LIB_PATH`
	`export LIBPATH`
	to:
	`CENTERA_LIB_PATH=`*`install_path`*`/lib/32`
	or:
	`CENTERA_LIB_PATH=`*`install_path`*`/lib/64`

Solaris	From:
	`CENTERA_LIB_PATH=/opt/Centera_SDK/lib/32`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`
	to:
	`CENTERA_LIB_PATH=`*`install_path`*`/lib/32`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`
	or:
	`CENTERA_LIB_PATH=`*`install_path`*`/lib/64`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`

Operating System	Script Revisions

Linux	From:
	`CENTERA_LIB_PATH=/usr/local/Centera_SDK/lib/32`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`
	to:
	`CENTERA_LIB_PATH=`*install_path*`/lib/32`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`
	or:
	`CENTERA_LIB_PATH=`*install_path*`/lib/64`
	`LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$CENTERA_LIB_PATH`
	`export LD_LIBRARY_PATH`

HP-UX
HP-UXi
	From:
	`CENTERA _LIB_PATH=/opt/Centera_SDK/lib/32`
	`SHLIB_PATH=$SHLIB_PATH:$CENTERA_LIB_PATH`
	`export SHLIB_PATH`
	to:
	`CENTERA _LIB_PATH=`*install_path*`/lib/32`
	`SHLIB_PATH=$SHLIB_PATH:$CENTERA_LIB_PATH`
	`export SHLIB_PATH`
	or:
	`CENTERA _LIB_PATH=`*install_path*`/lib/64`
	`SHLIB_PATH=$SHLIB_PATH:$CENTERA_LIB_PATH`
	`export SHLIB_PATH`

Windows	From:
	`set CENTERA _LIB_PATH=C:\Centera_SDK\lib\32`
	`set PATH=%PATH%;%CENTERA_LIB_PATH%`
	to:
	`set CENTERA_LIB_PATH=`*install_path*`\lib\32`
	`set PATH=%PATH%;%CENTERA_LIB_PATH`
	or:
	`set CENTERA_LIB_PATH=`*install_path*`\lib\64`
	`set PATH=%PATH%;%CENTERA_LIB_PATH`

3. Copy the modified script text into one of the application server startup scripts shown in the following table, or save the updated script and call it from the application server startup script.

Application Server	Startup Script (UNIX)	Startup Script (Windows)
WebSphere	setupCmdLine.sh	setupCmdLine.cmd
WebLogic	setDomainEnv.sh	setDomainEnv.cmd
JBoss	run.sh	run.cmd

4. Stop and start the application server.
5. Continue at one of the following topics, depending on your existing Content Engine version:
   - "Upgrade Content Engine software from 3.5.x" on page 218
   - "Upgrade Content Engine software from 4.0.x" on page 230

Task 3a: Upgrade Content Engine software from 3.5.x

If you are upgrading from Content Engine 3.5.x on Windows, you can upgrade and continue to use Windows, or you can migrate your Content Engine to UNIX. In either case, you must prepare the Content Engine application for deployment and then upgrade the object store data.

Migrating your Content Engine version 3.5.x on a Windows server to Content Engine version 4.5 on UNIX requires a new installation of Content Engine on your UNIX machine. You also need access to a Windows machine to install and run FileNet Enterprise Manager and the Upgrader tool.

The upgrade or migration process begins with the following steps, which are explained in this topic.

- "Install or upgrade Content Engine" on page 218
- "Configure Content Engine instances" on page 220

Install or upgrade Content Engine

Use a procedure in this topic to install or upgrade the following Content Engine components interactively or silently:

- Content Engine Server
- Configuration Manager
- .NET Clients (including FileNet Enterprise Manager)
- Content Engine Upgrader
- FileNet Deployment Manager

If the Content Engine installer detects an existing Content Engine installation, the software will be upgraded.

On UNIX-based or Windows-based application servers, you can install Content Engine Server, Configuration Manager, or both. The remaining components require a Windows-based machine.

Install the .NET Clients software only on machines where you intend to run either the FileNet Enterprise Manager administrative client or a customized client application.

The Content Engine Upgrader is required only to upgrade Content Engine and Content Search Engine data from an earlier version.

The FileNet Deployment Manager is an optional component that you can use to deploy test systems into full production. See the IBM FileNet P8 help topic Application Deployment > Get Started with FileNet Deployment Manager for details.

The application server must be the Deployment Manager node (WebSphere) or the Administrator node (WebLogic) if you are installing in a managed environment.

To install or upgrade Content Engine

1. Log on to the application server as *ce_install_user*. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8.*
2. Access the Content Engine software package.
3. Start the Content Engine installation. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8.*

   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 the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select PE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - To install interactively:

     i. Run one of the following commands in the software package, depending on your operating system, and then follow the instructions on the wizard screens:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN`
HPUX	`P8CE-4.5.0-HPUX.BIN`
HPUXi	`P8CE-4.5.0-HPUXI.BIN`
Linux	`P8CE-4.5.0-LINUX.BIN`
Solaris	`P8CE-4.5.0-SOL.BIN`
Windows	`P8CE-4.5.0-WIN.EXE`
zLinux	`P8CE-4.5.0-ZLINUX.BIN`

     ii. In the Choose Install Path screen, specify the install path, for example /opt/FileNet/Content Engine (UNIX) or C:\Program Files\FileNet\ContentEngine.

     iii. In the Choose Components screen, select the appropriate check boxes, depending on your operating system:

     UNIX

     Select Content Engine Server and Tools (to install Configuration Manager).

     Windows

     Select Content Engine Server to upgrade your existing Windows installation.

     Select .NET Clients software on machines where you intend to run either the FileNet Enterprise Manager administrative client or a customized client application.

Select Tools (to install Configuration Manager, Content Engine Upgrader, and FileNet Deployment Manager).

iv. At the Install Complete screen of the wizard, selecting the Launch Configuration Manager check box will automatically start the tool for configuring Content Engine.

- To install silently:

  i. Open the CE_silent_install.txt file in the software package for editing.

  ii. Set the parameter values in the CE_silent_install.txt file for your site.

  For information on the Content Engine parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8.*

  iii. Run one of the following commands in the software package, depending on your operating system:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN -f CE_silent_install.txt –i silent`
HPUX	`P8CE-4.5.0-HPUX.BIN -f CE_silent_install.txt –i silent`
HPUXi	`P8CE-4.5.0-HPUXI.BIN -f CE_silent_install.txt –i silent`
Linux	`P8CE-4.5.0-LINUX.BIN -f CE_silent_install.txt –i silent`
Solaris	`P8CE-4.5.0-SOL.BIN -f CE_silent_install.txt –i silent`
Windows	`P8-CE-4.5.0-WIN.EXE -f CE_silent_install.txt –i silent`
zLinux	`P8CE-4.5.0-ZLINUX.BIN -f CE_silent_install.txt –i silent`

4. Check for errors in the ce_install_log_4_5_0.txt Content Engine error log file.

5. If you are migrating to UNIX, repeat Step 1 through Step 4 on a Windows machine to install FileNet Enterprise Manager and the Upgrader tool. Both programs are required to complete the upgrade.

Configure Content Engine instances

To configure an instance of Content Engine, you will run Configuration Manager. There are two versions of the Configuration Manager: command line interface (CLI) and graphical user interface (GUI).

You can run either version to configure an instance. After configuring an instance, you will deploy it in "Deploy upgraded Content Engine instances" on page 241. You can configure all the instances before deploying any of them, or you can configure and deploy one instance at a time.

NOTE (WebSphere only) For best results, configure no more than one Content Engine instance in a profile.

First, perform the procedure in "Grant permissions to the Configuration Manager user" on page 221, and then perform the procedure in one of the following subtopics, depending on which version of the tool you want to run:

- "Configure instances using the graphical user interface" on page 222.
- "Configure instances using the command line interface" on page 224. If you need an accessible version of Configuration Manager, use the command line interface instead of the GUI.

Grant permissions to the Configuration Manager user

Perform the following procedure to grant the file and directory permissions required by *config_mgr_user*, the user who will run Configuration Manager.

To grant permissions to the Configuration Manager user

1. If you have not done so already, log on to the machine where you installed Content Engine as *ce_install_user*. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8.*
2. Navigate to *ce_install_path*/tools/configure, which contains both the command line and graphical user interfaces of Configuration Manager, where *ce_install_path* is the path in which you installed Content Engine.
3. Grant *config_mgr_user* execute permission (UNIX) or read & execute permission (Windows) on the executable file of the interface of Configuration Manager you intend to use:
   - To enable use of the graphical user interface, grant permission to one of the following files in the configuration/CMUI directory:

     UNIX

     cmui

     Windows

     cmui.exe
   - To enable use of the command line interface, grant permission to one of the following files in the configuration directory:

     UNIX

     configmgr.sh

     Windows

     configmgr.bat
4. Grant write permission to the directory where you want Configuration Manager to place the configuration XML files it will generate.

   If you are not going to specify this directory when you run Configuration Manager, grant write permission on the default directory, *ce_install_path*/tools/configure/profiles.
5. Log off the machine hosting the application server and log back on as *config_mgr_user*, the Configuration Manager user.

6. Continue at one of the following subtopics:
   - "Configure instances using the graphical user interface" on page 222.
   - "Configure instances using the command line interface" on page 224.If you need an accessible version of Configuration Manager, use the command line interface instead of the GUI.

Configure instances using the graphical user interface

In this subtopic you will configure a Content Engine Server instance on an application server using the graphical user interface version of Configuration Manager. Use the information in your worksheet to specify the values for the parameters required to configure Content Engine. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

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 you must specify for Configuration Manager:

- Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select Configuration Manager.
- Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

Refer to the appendix "Configuration Manager user interface reference" on page 418 for complete information on using the graphical user interface.

To create a configuration profile

1. Navigate to the *ce_install_path*/tools/configure/CMUI directory, which contains the graphical user interface of Configuration Manager.
2. Start Configuration Manager by running one of the following commands, depending on the operating system that runs on the machine where you installed Content Engine:

   UNIX

   ```
   cmui
   ```

   Windows

   ```
   cmui.exe
   ```

3. Start the Create New Configuration Profile wizard, by selecting **File > Create New Configuration Profile** or by clicking the icon in the tool bar.
4. Enter a name for the profile. The name must be a valid directory name for your operating system.
5. Specify the path for the profile. Either type in the full path to the profile directory or click **Browse** to locate the directory. The default path is *ce_install_path*/tools/configure/profiles, where *ce_install_path* is the location where Content Engine is installed.
6. Click **Next**.
7. Select the application server type for this profile.

8. Click **Next**.
9. Enter the property values for the application server. The specific properties displayed depend on the server type that you selected in Step 7.
10. Click **Next**.
11. Select the tasks that you want to include in this profile. For a new Content Engine installation, you need to complete all four configuration tasks: Configure JDBC Data Sources, Configure LDAP, Configure Bootstrap, and Deploy Application.
12. Click **Finish** to create the profile.

    The profile you created will be displayed as an icon in the left-hand pane, along with three icons for the configure tasks, and one icon for the deploy task:

    - Configure JDBC Data Sources
    - Configure LDAP
    - Configure Bootstrap Properties
    - Deploy Application

To configure a Content Engine instance using the Configuration Manager graphical user interface

You will perform the configure tasks in this procedure, and the deploy task in "Deploy upgraded Content Engine instances" on page 241. Performing a task means editing it, running it, and then checking its completion status.

You can perform the configure tasks in any order, and you do not need to complete work on one configure task before starting another.

1. (Optional) By default, Configuration Manager does not save passwords for the tasks you will edit. That is, each time you start Configuration Manager and open a saved profile, you will need to specify any passwords required by the tasks and for the application server properties. For more information, see "Handling passwords in Configuration Manager" on page 417.
2. Edit a configure task by double-clicking one of the configure icons, or by right-clicking the icon and choosing **Edit Selected Task**. Then specify the parameter values in the right pane of the window, using the appropriate information from your installation and upgrade worksheet. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

   After specifying all the required values, choose **File > Save** to save your work, or save your work if you need to exit Configuration Manager before completing your edits. If you do not save your edits, Configuration Manager prompts you to save your changes before running the task.
3. To run a configure task, select the **Task enabled** check box, right-click the configure icon, and choose **Run Task**. (Running the configuration task may take a few minutes.)
4. To check the completion status of a configure task, right-click the configure icon, and choose **Check Task Status**.
5. If any configure tasks remain for this profile, return to Step 2 and choose another configure task.

6. Continue at "Install Content Engine software updates" on page 238.

Configure instances using the command line interface

In this subtopic you will configure a Content Engine instance on a given application server using the command-line version of Configuration Manager. Configuring a Content Engine instance involves four major steps (repeat the steps to configure another instance):

- Generate the configuration XML files that contain the properties and values used to perform various tasks required to configure the environment for a Content Engine instance. See "Generate the configuration XML files for a Content Engine instance" on page 224.
- Edit the configuration XML files by inserting your site's properties and values. See "Edit the configuration XML files for a Content Engine instance" on page 226.
- Execute the configuration XML files you edited. See "Execute the configuration XML files for a Content Engine instance" on page 227.
- Check that the configuration XML files that you executed have resulted in a correct configuration of the Content Engine instance. See "Check the completion status of Content Engine configuration tasks" on page 228.

If you are deploying multiple Content Engine instances on the same machine, you will need to generate, edit, and deploy a complete set of configuration files for each instance. Store the configuration files for each instance in a separate directory.

You can navigate through the steps above by generating all the configuration XML files before editing, executing, or verifying any of them; or you can generate, edit, execute, and verify one file at a time.

NOTE This subtopic refers to Configuration Manager as `configmgr`. When you run the tool, substitute `configmgr.sh` (on UNIX) or `configmgr.bat` (on Windows) in place of `configmgr`, depending on your operating system.

Generate the configuration XML files for a Content Engine instance

The following table lists the configuration XML files that you will generate using Configuration Manager:

File Name	Description
configurejdbc.xml	Settings for configuring JDBC connections to the databases used by Content Engine
configureldap.xml	Settings for connecting to and searching within a directory server
configurebootstrap.xml	Settings for creating the Global Configuration Data (GCD) and starting Content Engine
deployapplication.xml	Settings for deploying a Content Engine instance
configureapplicationserver.xml	Parent file of the above listed configuration XML files.
applicationserver.xml	Settings for the application server, including the location of the application server software and the name of the server.

You can generate the configuration XML files, in either of two ways:

- Run the tool multiple times, generating one configuration XML file at a time.
- Run the tool once to generate a "parent" file, configureapplicationserver.xml, and automatically generate all the configuration XML files. The parent file points to the individual configuration XML files.

Running the tool also generates the applicationserver.xml file (used in several configuration tasks). Subsequent executions of the tool will not overwrite applicationserver.xml.

To generate configuration XML files

1. Set the current directory to *ce_install_path*/tools/configure.
2. At the command prompt, run Configuration Manager to generate the configuration XML files all at once (Step a) or one at a time (Step b).

   See "Configuration Manager command-line reference" on page 429 for details on Configuration Manager syntax.

   a. To generate all the configuration files at once with a parent configuration file, run the following command. Do not type any line breaks when you type the command.

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureApplicationServer -path mypath
   ```

   where:

   app_server_type is WebSphere, WebLogic, or JBoss

   -path *mypath* is optional and specifies the path to where the generated configuration XML files will be placed. If you are deploying multiple Content Engine instances on the same machine, you will need to specify a separate directory for each instance.

   Continue at "Edit the configuration XML files for a Content Engine instance" on page 226.

   b. To generate a single configuration XML file, run the command in one of the following substeps:

   - To generate the configurejdbc.xml file:

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureJDBC -path mypath
   ```

   - To generate the configureldap.xml file:

   ```
   configmgr generateConfig -appserver app_server_type -task ConfigureLDAP
   -path mypath
   ```

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureLDAP -path mypath
   ```

   - To generate the configurebootstrap.xml file:

   ```
   configmgr generateConfig -appserver app_server_type -task ConfigureBootstrap
   -path mypath
   ```

   ```
   configmgr generateConfig -appserver app_server_type -db db_type
   -ldap ldap_type -deploy deploy_type -task ConfigureBootstrap [-path mypath]
   ```

where *mypath* is optional and specifies the path to where the generated configuration XML files will be placed. If you are deploying multiple Content Engine instances on the same machine, you will need to specify a separate directory for each instance.

Repeat Step b to generate one of the other configuration XML files, or continue at "Edit the configuration XML files for a Content Engine instance" on page 226.

You will eventually need to generate each of the configuration XML files to configure a Content Engine instance.

Edit the configuration XML files for a Content Engine instance

Perform the following procedure for each file you generated in "Generate the configuration XML files for a Content Engine instance" on page 224 to insert your site's properties and values.

To edit values in the configuration XML files

1. Use a text editor or XML editor to open one of the following configuration XML files that you generated in "Generate the configuration XML files for a Content Engine instance" on page 224:
   - configurejdbc.xml
   - configureldap.xml
   - configurebootstrap.xml
   - applicationserver.xml

   NOTE If you generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 224, you will also have generated the deployapplication.xml file. You will open this file for editing in "Deploy upgraded Content Engine instances" on page 241.

2. In the configuration XML file, replace each occurrence of `****INSERT VALUE****` with a value appropriate for your site. Refer to the descriptions in the file for more information.

   Set the <TaskEnabled> value to `true` in any configuration XML file you edit if you want to run the configuration task in "Execute the configuration XML files for a Content Engine instance" on page 227.

   NOTE If you previously specified values in the configureldap.xml file to add a realm to a federated repository, and want to put an additional realm in the repository, replace the previous values with the values for the additional realm.

3. (Optional, WebSphere only) If you have previously created XA and non-XA data sources that you want to use for the Global Configuration Data (GCD) database, make the following edits:

   a. In the configurejdbc.xml file, set the <TaskEnabled> value to `false` to avoid creating another pair (XA and non-XA) of data sources.

   b. In the configurebootstrap.xml file, set the <JDBCDataSourceXAFileName> and <JDBCDataSourceFileName> values to the XA and non-XA JNDI names, respectively, associated with the GCD database.

4. (Optional) Encrypt any passwords that you need to insert into the file by running the password encryption utility (see "To encrypt a password for Configuration Manager" on page 453), and then

copy the encrypted value into the file. It is a best practice to encrypt the passwords for the following accounts:

- The application server administrator account used in the applicationserver.xml file.
- The database administrator account used in the configurejdbc.xml file.
- The LDAP provider service principal account used in the configureldap.xml file.
- The master key—a word or phrase for encrypting sensitive FileNet P8 Global Configuration Data (GCD) entries—used in the configurebootstrap.xml file.

CAUTION Any password you do not encrypt will be stored and sent as clear text.

5. Save your edits.
6. Perform one of the following:
   - Repeat Step 1 through Step 5 of this procedure for any other configuration XML file that you have not yet edited.
   - Continue at "Execute the configuration XML files for a Content Engine instance" on page 227 to execute a configuration XML file(s) you have edited.
   - Return to "Generate the configuration XML files for a Content Engine instance" on page 224 to generate additional configuration XML files.

Execute the configuration XML files for a Content Engine instance

Perform the procedure in this subtopic to execute the settings in the configuration XML files you have generated and edited for a Content Engine instance.

NOTES

- Any task with the <TaskEnabled> element value set to false will not run (see Step 2 on page 226 in "To edit values in the configuration XML files" on page 226).
- If you are executing tasks for a profile that was created or edited in the Configuration Manager GUI, verify that the XML files contain values for the required passwords. See "Handling passwords in Configuration Manager" on page 417 for more information.

To execute the configuration XML file settings

1. (WebSphere only) Start the application server if it is not already running.
2. Set the current directory to *ce_install_path*/tools/configure.
3. At the command prompt, run Configuration Manager to execute the configuration XML files all at once (Step a) or one at a time (Step b). See "Configuration Manager command-line reference" on page 429 for command syntax details.

When running the tool, the `-path mypath` parameter is optional and specifies where you placed the configuration XML file. If you do not specify a path, the file must be in *ce_install_path*/tools/configure/profiles.

a. To execute all the configuration files at once with a parent configuration file, run the following command.

```
configmgr execute -task ConfigureApplicationServer -path mypath
```

NOTES

- To execute all tasks at once, you must have generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 224.
- If you did not enable the deployment task in the deployapplication.xml file, the tool will display an informational message indicating that the deployment did not occur. In which case you will perform the deployment in "Deploy upgraded Content Engine instances" on page 241.

b. To execute a single configuration XML file, type and run the command in one of the following substeps:

- To execute the configurejdbc.xml file:

```
configmgr execute -task ConfigureJDBC -path mypath
```

- To execute the configureldap.xml file:

```
configmgr execute -task ConfigureLDAP -path mypath
```

- To execute the configurebootstrap.xml file:

```
configmgr execute -task ConfigureBootstrap -path mypath
```

Repeat Step b to execute one of the other configuration XML files, or continue at "Check the completion status of Content Engine configuration tasks" on page 228.

You will eventually need to execute each of the configuration XML files to complete the configuration of a Content Engine instance.

4. Continue at "Check the completion status of Content Engine configuration tasks" on page 228.

Check the completion status of Content Engine configuration tasks

Use the procedure in this subtopic to verify that one or more Content Engine configuration tasks that you executed in "Execute the configuration XML files for a Content Engine instance" on page 227 have completed. You can check the status of all the tasks you executed or just check individual tasks.

NOTE Checking the completion status does not validate the information in the XML files.

To check the status of a Content Engine configuration

1. At the command prompt, run Configuration Manager to check the status of the configuration tasks all at once (Step a) or one at a time (Step b).

When running the tool, the `-path mypath` parameter is optional and specifies where you placed the configuration XML file. If you do not specify a path, the file must be in *ce_install_path*/tools/configure/profiles.

See "Configuration Manager command-line reference" on page 429 for command syntax details.

a. To check the status of all the configuration tasks at once, run the following command:

```
configmgr checkStatus -task ConfigureApplicationServer -path mypath
```

NOTE To check the status of all tasks at once, you must have generated all the files at once with the ConfigureApplicationServer task in "Generate the configuration XML files for a Content Engine instance" on page 224.

Continue at Step 3.

b. To check the completion status of a single configuration task, run the tool in one of the following substeps.

- To check the status of the ConfigureJDBC task using the configuration file in the specified path:

```
configmgr checkStatus -task ConfigureJDBC -path mypath
```

- To check the status of the ConfigureLDAP task using the configuration file in the specified path:

```
configmgr checkStatus -task ConfigureLDAP -path mypath
```

- To check the status of the ConfigureBoostrap task using the configuration file in the specified path:

```
configmgr checkStatus -task ConfigureBootstrap -path mypath
```

If you performed Step b, repeat the step to check the status of any other configuration tasks you have executed; otherwise, continue at Step 2.

2. Continue at one of the following procedures, as needed; otherwise, continue at Step 3.

- If you have any other configuration tasks to execute, continue at "Execute the configuration XML files for a Content Engine instance" on page 227.
- If you have any other configuration XML files to generate, continue at "Generate the configuration XML files for a Content Engine instance" on page 224.
- If you want to add a realm to a federated repository, continue at "Edit the configuration XML files for a Content Engine instance" on page 226.

3. Continue at "Install Content Engine software updates" on page 238.

Task 3b: Upgrade Content Engine software from 4.0.*x*

Use the procedure in this topic to upgrade Content Engine version 4.0.*x* to version 4.5. On UNIX-based application servers, you will upgrade Content Engine Server and install Configuration Manager. On Windows-based application servers, you will upgrade or install anew some or all of the following components.

- Content Engine Server
- Configuration Manager
- .NET Clients (including FileNet Enterprise Manager)
- Content Engine Upgrader
- FileNet Deployment Manager

NOTE FileNet Deployment Manager is used to migrate system data from a test environment to a production environment. See the IBM FileNet P8 help topic Application Deployment > Get Started with FileNet Deployment Manager for details.

CAUTION You must upgrade components that are in your existing installation, and you can install new components as well.

To upgrade Content Engine software from 4.0.*x*

1. Log on to the application server as *ce_upgrade_user*, the user who installed the current version of Content Engine.

   If the user who installed the current version will is different than the user who install the new version, make sure that you have completed the steps in "Assign directory permissions for Content Engine upgrade for 4.0.x to 4.5 on UNIX" on page 177 in *Plan and Prepare Your Environment for IBM FileNet P8*. That topic also contains information on changing the installation user.

   Continue with one of the following steps:

   - For AIX, where the root user installed CE 4.0.x, and *ce_upgrade_user* is non-root, continue at Step 2.
   - All others, continue at Step 5.

2. Shut down the application server.

3. (AIX only) Start the Content Engine uninstall.

   - To uninstall interactively

     i. Navigate to the _uninst2 directory in the current version of Content Engine and run the following command to uninstall it:

     ```
     uninstaller.bin
     ```

     ii. Wait until the uninstall completes before continuing with Step 5.

- To uninstall silently:

  i. Navigate to the _uninst2 directory in the current Content Engine installation.

  ii. Either edit the CE_silent_uninstall.txt sample response file or record your own reponse file. To record a response file, run the following command:

  CAUTION Do not specify a path to the response file within the Content Engine install directory; if you do this, the file will be deleted.

  ```
  ./uninstaller.bin -options -record path_to_response_file
  ```

  iii. Run the uninstall program.

  ```
  ./uninstaller.bin -options path_to_response_file -silent
  ```

  iv. Wait until the uninstall completes before continuing with Step 5.

4. Uninstall Content Search Engine Client Updater.

5. Make a back up copy of the *ce_install_path*/servers.xml file in a temporary location, such as /tmp.

6. (Windows only) Navigate to the ContentEngine directory and give Full Control permission to the user who will install the new version of Content Engine.

7. Log off the Content Engine Server machine, and log back on as *ce_upgrade_user*, the user who will install the new version.

8. Copy the new version of the Content Engine software package to a temporary location on the the application server, such as /tmp (UNIX) or C:\Temp (Windows) and set the current directory to that location.

9. Start the Content Engine upgrade (which will uninstall the current version if needed). For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8.*

   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 the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select CE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Upgrade."

- To upgrade interactively:

  i. Run one of the following commands in the software package, depending on your operating system, and then follow the instructions on the wizard screens:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN`
HPUX	`P8CE-4.5.0-HPUX.BIN`
HPUXi	`P8CE-4.5.0-HPUXi.BIN`
Linux	`P8CE-4.5.0-LINUX.BIN`

Platform	Command
Solaris	`P8CE-4.5.0-SOL.BIN`
Windows	`P8CE-4.5.0-WIN.EXE`
zLinux	`P8CE-4.5.0-ZLINUX.BIN`

ii. At the Choose Components screen of the wizard, choose components to upgrade/install.

iii. At the Install Complete screen of the wizard, selecting the Launch Configuration Manager check box will automatically start the tool for configuring Content Engine.

Do not select the check box if version 4.0.*x* of Content Engine is a cluster deployment, or a non-cluster Network Deployment (WebSphere only). In this case you must edit the servers.xml file first (see "Edit the servers.xml file" on page 234), and then manually start Configuration Manager.

- To upgrade silently:

i. Run one of the following commands in the software package, depending on your operating system:

Platform	Command
AIX	`P8CE-4.5.0-AIX.BIN -f CE_silent_install.txt –i silent`
HPUX	`P8CE-4.5.0-HPUX.BIN -f CE_silent_install.txt –i silent`
HPUXi	`P8CE-4.5.0-HPUXI.BIN -f CE_silent_install.txt –i silent`
Linux	`P8CE-4.5.0-LINUX.BIN -f CE_silent_install.txt –i silent`
Solaris	`P8CE-4.5.0-SOL.BIN -f CE_silent_install.txt –i silent`
Windows	`P8-CE-4.5.0-WIN.EXE -f CE_silent_install.txt –i silent`
zLinux	`P8CE-4.5.0-ZLINUX.BIN -f CE_silent_install.txt –i silent`

10. Check for errors in the ce_install_log_4_5_0.txt Content Engine error log file.
11. Copy the servers.xml file from its temporary location back to its previous location if you performed Step 2 on page 230
12. Continue at "Configure Content Engine instances upgraded from 4.0.x" on page 233.

Task 4: Configure Content Engine instances upgraded from 4.0.x

To configure the upgraded content engine instances, complete the procedures in this topic.

- "Grant directory permissions to the Configuration Manager user" on page 233
- "Edit the servers.xml file" on page 234
- Configure one or more upgraded instances of Content Engine on an application server, using one of these methods:
  - "Configure upgraded instances using a graphical user interface" on page 235
  - "Configure upgraded instances using a command line" on page 236

NOTE On a machine that runs Novell SUSE Linux Enterprise 9, you can configure Content Engine only from a command line.

Grant directory permissions to the Configuration Manager user

Perform the following procedure to grant the file and directory permissions required by *config_mgr_user*, the user who will run Configuration Manager. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8*.

To grant directory permissions to the Configuration Manager tool user

1. If you have not done so already, log on to the Content Engine Server machine as the user who ran the Content Engine install program, *ce_install_user.*
2. Grant the following permissions to *config_mgr_user,* the Configuration Manager user:
   - Execute permission on the following file:

     UNIX

     ce_install_path/tools/configure/configmgr.sh

     Windows

     ce_install_path/tools/configure/configmgr.bat

     where *ce_install_path* is the path where you installed Content Engine.
   - Write permission to the directory where you want the Configuration Manager tool to place the configuration XML files it will generate.

     NOTE If you are not going to specify this directory when you run the tool, grant write permission on the default directory, *ce_install_path*/tools/configure/profiles.

3. Grant *config_mgr_user* execute permission (UNIX) or Read & Execute permission (Windows) on the executable file of the interface of Configuration Manager you intend to use:
   - To enable use of the graphical user interface, grant the permission to one of the following files:

     UNIX

     cmui

     Windows

     cmui.exe
   - To enable use of the command line interface, grant permission to one of the following files:

     UNIX

     configmgr.sh

     Windows

     configmgr.bat
4. Grant write permission to the directory where you want Configuration Manager to place the configuration XML files it will generate.

   If you are not going to specify this directory when you run Configuration Manager, grant write permission on the default directory, *ce_install_path*/tools/configure/profiles.
5. Log off the machine hosting the application server and log back on as *config_mgr_user*, the Configuration Manager user. For details on required accounts and related permissions, see "Accounts for Content Engine" on page 66 in *Plan and Prepare Your Environment for IBM FileNet P8*.

Edit the servers.xml file

If you upgraded Content Engine from version 4.0.x of a cluster deployment, or a non-cluster Network Deployment (WebSphere only), perform the procedure in this subtopic to edit the servers.xml file; otherwise, skip this subtopic.

To edit the servers.xml file

1. On the Content Engine Server machine, navigate to *ce_install_path*/FileNet/ContentEngine, which contains the servers.xml file.
2. Use the following syntax to edit the settings for the Content Engine server:

   ```
   <server name="WS/Deployment_Manager_profile_name/cell_name/cluster_name"
   deployment="Cluster" version="4.0.x-xxxx"
   ear="path_to_existing_40x_boostrapped_EAR_file\Engine-ws.ear">
   ```

   where you would:
   - Add the attribute `deployment="Cluster"`.
   - Change the name attribute by entering your cluster name for the *cluster_name* variable (replacing the server name).

- Verify that the path to the EAR file is valid. The path should point to the EAR file deployed to the cluster.

3. Save the servers.xml file.

Configure upgraded instances using a graphical user interface

In this subtopic you will configure a Content Engine Server instance on an application server using the graphical user interface version of Configuration Manager.

Refer to the appendix "Configuration Manager user interface reference" on page 418 for complete information on using the graphical user interface.

NOTE If you need an accessible version of Configuration Manager, use the command line interface instead of the GUI. See "Configure upgraded instances using a command line" on page 236.

To create an upgrade configuration profile

1. Navigate to the *ce_install_path*/tools/configure/CMUI directory, which contains the graphical user interface of Configuration Manager.
2. Start Configuration Manager by running one of the following commands, depending on the operating system that runs on the machine where you installed Content Engine:

   UNIX

   cmui

   Windows

   cmui.exe

3. Start the Create Upgrade Configuration Profile wizard, by selecting **File > Create Upgrade Configuration Profile** or by clicking the wizard icon in the tool bar, and use the information in your worksheet to specify the values for the parameters in the wizard screens. For more information, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.
4. In the Upgrade Profile Path screen, specify any name for the profile. Configuration Manager will create a directory with the name you specify.

   If more than one instance of Content Engine Server is installed on the machine, Configuration Manager will create one parent directory and a subdirectory of the parent for each instance. The profile you create will be displayed as an icon in the left-hand pane, along with icons for Upgrade Bootstrap and Deploy Application.

NOTE The values you have specified for the profile apply to only one Content Engine instance on the application server. You must repeat this procedure for any other instances.

To configure a Content Engine instance using the Configuration Manager graphical user interface

NOTE For information on the values you will specify when configuring a Content Engine instance, see "Installation and upgrade worksheet" on page 216 of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

1. Edit the Upgrade Bootstrap task by double-clicking its icon, or by right-clicking the icon and choosing **Edit Selected Task**. Then specify the parameter values in the right pane of the window, using the appropriate information from your installation and upgrade worksheet.

   After specifying all the required values, choose **File > Save** to save your work, or save your work if you need to exit Configuration Manager before completing your edits. If you do not save your edits, Configuration Manager prompts you to save your changes before running the task.

2. To run the Upgrade Bootstrap task, select the **Task enabled** check box in the Upgrade Bootstrap tab, right-click the configure icon, and choose **Run Task**. (Running the configuration task may take a few minutes.)

3. To check the completion status of the Upgrade Bootstrap task, right-click the configure icon, and choose **Check Task Status**.

Configure upgraded instances using a command line

In this subtopic you will configure an upgraded Content Engine instance on a given application server using the command-line version of the Configuration Manager tool. Configuring an upgraded instance involves the following major steps (repeat the steps to configure another instance):

1. Generate the upgradebootstrap.xml file that contain the properties and values used to configure the environment for an upgraded Content Engine instance.

2. Execute the upgradebootstrap.xml file.

3. Verify that the executed upgradebootstrap.xml file has resulted in a correct configuration of the Content Engine instance.

You can navigate through the steps above by generating all the upgradebootstrap.xml files, one for each Content Engine instance, before editing, executing, or verifying any of them; or you can generate, edit, execute, and verify one upgradebootstrap.xml file at a time.

To configure a Content Engine instance using the Configuration Manager command line interface

1. Log on to the Content Engine Server machine as *config_mgr_user*, the Configuration Manager user. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

2. Set the current directory to *ce_install_path*/tools/configure, where:

   ce_install_path is the path where you installed Content Engine.

   As an example, *ce_install_path* might be /opt/FileNet/ContentEngine.

3. At a command prompt, run the following Configuration Manager command interactively, or silently (without prompts) by appending the flag `-silent`, to generate the upgradebootstrap.xml file.

   ```
   configmgr generateConfig -task upgrade -path mypath
   ```

   The *mypath* parameter in the command specifies the path where the tool will place the configuration XML files it generates. See "Configuration Manager command-line reference" on page 429 for details.

4. If your application server type is WebSphere or WebLogic, perform the following substeps:

   a. Navigate to the profile directory you created in "To create an upgrade configuration profile" on page 235, and open the applicationserver.xml file for editing.

   b. In the ApplicationServerAdminUsername element, set the <value> element content to the user name of the application server console administrator, as follows:

   ```
   <property name="ApplicationServerAdminUsername">
      <value>administrator_user_name</value>
   </property>
   ```

   c. In the ApplicationServerAdminPassword element, set the <value> element content to the password of the application server console administrator, as follows:

   ```
   <property name="ApplicationServerAdminPassword">
      <value>administrator_password</value>
   </property>
   ```

   d. Save your edits.

5. At a command prompt, run Configuration Manager interactively, or silently (without prompts) by appending the flag `-silent` to the command, to execute the upgradebootstrap.xml file.

   ```
   configmgr execute -task upgrade -path mypath
   ```

   The *mypath* parameter in the command specifies the path where the tool will find the upgradebootstrap.xml file.

6. At a command prompt, run Configuration Manager interactively, or silently (without prompts) by appending the flag `-silent` to the command, to check the completion status of the execute task.

   ```
   configmgr checkStatus -task upgrade -path mypath
   ```

Task 5: Install Content Engine software updates

Perform the procedure in this topic for each Content Engine instance to install software updates, fix packs or interim fixes.

If no Content Engine software updates are available, skip to "Install the latest Process Engine Client files on Content Engine servers" on page 44.

To install the Content Engine software updates

1. For instructions on how to obtain the latest Content Engine software updates, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
2. Open the readmes for the Content Engine software updates and perform the installation procedures in the readmes on each Content Engine instance.

   If you installed an instance into a managed environment, perform the procedure on the Deployment Manager node (WebSphere) or the Administrator node (WebLogic).
3. Continue at one of the following topics:
   - If you are upgrading from version 3.5.x or if your IBM FileNet P8 installation includes Process Engine, continue at "Install the latest Process Engine Client files on Content Engine servers" on page 239.
   - If you are upgrading from version 4.0.x and do not need to upgrade your Process Engine client files to 4.5, continue with "Configure Content Engine instances upgraded from 4.0.x" on page 233.

Task 6: Install the latest Process Engine Client files on Content Engine servers

To install the Process Engine Client files, perform the following steps on all application server machines where Content Engine Server is to be deployed.

To install the Process Engine Client files

1. On the machine where Content Engine is to be deployed, log on as *ce_upgrade_user*.
2. Access the Process Engine Client install software from the Process Engine installation software. The version of the install software must match the version of Process Engine.
3. Expand the (TAR or ZIP) Process Engine Client install software.
4. Shut down all instances of Enterprise Manager (EM) and any other Content Engine client applications, such as Application Engine.
5. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Content Engine will be deployed. Run the program either interactively (using the install wizard) or silently.

   NOTE The Process Engine Client installation program creates a backup of the previous Process Engine Client files. For example, on Windows you can find this backup at C:\Program Files\FileNet\PEClient\Backup\CE.

   To run the program interactively, run one of the commands in the table below, depending on the operating system and follow the wizard instructions:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN`

To run the program silently, perform the following steps:

a. In the expanded install software, open the file PEClient_silent_install.txt and edit it as follows:

   i. Change the `Variable_CheckboxCE` line to the following:

   ```
   -V Variable_CheckboxCE="true"
   ```

   ii. Save your edit.

b. Run one of the commands in the following table to perform the silent install:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN -silent -options "PEClient_silent_install.txt"`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN -silent -options "PEClient_silent_install.txt"`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN -silent -options "PEClient_silent_install.txt"`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN -silent -options "PEClient_silent_install.txt"`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN -silent -options "PEClient_silent_install.txt"`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE -silent -options "PEClient_silent_install.txt"`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN -silent -options "PEClient_silent_install.txt"`

Task 7: Deploy upgraded Content Engine instances

To deploy an upgraded instance of Content Engine, perform the: procedures in this topic after you have performed the configure tasks in one of the following topics

- For upgrades from 3.5.x, "Configure Content Engine instances" on page 220
- For upgrades from 4.0.x, "Configure Content Engine instances upgraded from 4.0.x" on page 233.

If you have configured all your instances, you can deploy all of them in this topic. Alternatively, if you configured only one of your instances, you can deploy it in this topic and then return to "Configure Content Engine instances" on page 220 or "Configure Content Engine instances upgraded from 4.0.x" on page 233 to configure another instance.

To deploy a Content Engine instance, perform the procedures in one of the following subtopics:

- "Deploy upgraded instances using a graphical user interface" on page 241
- "Deploy upgraded instances using a command line" on page 242

NOTE On a machine that runs Novell SUSE Linux Enterprise 9, you can deploy Content Engine only from a command line.

Deploy upgraded instances using a graphical user interface

In this subtopic you will deploy an upgraded Content Engine Server instance on an application server using the graphical user interface version of Configuration Manager.

Refer to the appendix "Configuration Manager user interface reference" on page 418 for complete information on using the graphical user interface.

To deploy a Content Engine instance using the Configuration Manager graphical user interface

NOTE For information on the values you will specify when deploying a Content Engine instance, see "Installation and upgrade worksheet" on page 216 of *Plan and Prepare Your Environment for IBM FileNet P8*.

1. If you haven't already done so, log on to the application server machine as *config_mgr_user*, the Configuration Manager user. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8.*
2. (WebSphere and WebLogic only) Stop the application server.
3. Navigate to the *ce_install_path*/tools/configure/CMUI directory, which contains the graphical user interface of Configuration Manager.
4. Start Configuration Manager by running one of the following commands, depending on the operating system that runs on the machine where you installed Content Engine:

   UNIX

   cmui

Windows

cmui.exe

5. If you upgraded your application server and then redeployed Content Engine before upgrading it to version 4.5 before upgrading Content Engine from version 4.0.x to 4.5, click the Update Application Server Properties icon and correct any obsolete parameter values for your application server type.
6. Edit the Deploy Application task by double-clicking its icon, or by right-clicking the icon and choosing **Edit Selected Task**. Then specify the parameter values in the right pane of the window, using the appropriate information from your installation and upgrade worksheet.
7. Select the **Task enabled** check box in the Upgrade Bootstrap tab.
8. Choose **File > Save** to save your work, or save your work if you need to exit Configuration Manager before completing your edits. If you do not save your edits, Configuration Manager prompts you to save your changes before running the task.
9. To run the Deploy Application task, right-click the configure icon, and choose **Run Task**. (Running the configuration task may take a few minutes.)
10. To check the completion status of the Deploy Application task, right-click the Deploy Application icon, and choose **Check Task Status**.
11. (WebSphere and WebLogic only) Start the application server.
12. (JBoss only) Stop and start the application server.

Deploy upgraded instances using a command line

In this subtopic you will deploy an upgraded Content Engine instance on a given application server using the command-line version of Configuration Manager. Deploying an upgraded instance involves the following steps (repeat the steps to deploy another instance):

1. Generate the deployapplication.xml file.
2. Execute the deployapplication.xml file you edited.
3. Verify that the deployapplication.xml file that you executed has run to completion.

You can navigate through the steps above by generating all the deployapplication.xml files before executing or verifying any of them; or you can generate, execute, and verify one file at a time.

To deploy a Content Engine instance using the Configuration Manager command line interface

1. If you haven't already done so, log on to the application server machine as *config_mgr_user*, the Configuration Manager user. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 in *Plan and Prepare Your Environment for IBM FileNet P8.*
2. (WebSphere and WebLogic only) Stop the application server.

3. Set the current directory to *ce_install_path*/tools/configure, where:

   ce_install_path is the path where you installed Content Engine.

   As an example, *ce_install_path* might be /opt/FileNet/ContentEngine.

4. If you upgraded your application server and then redeployed Content Engine before upgrading it to version 4.5 before upgrading Content Engine from version 4.0.x to 4.5, perform the following substeps:

   a. Open the applicationserver.xml file for editing.

   b. Correct any obsolete parameter values for your application server type.

   c. Save your edits.

5. Edit the deployapplication.xml file as appropriate for your site. Refer to the descriptions in the file for more information.

6. At a command prompt, execute the deployapplication.xml file by running Configuration Manager either interactively, or silently (without prompts) by appending the flag `-silent` to the command:

   ```
   configmgr execute -task DeployApplication -path MyPath
   ```

   The *mypath* parameter in the command specifies the path where the tool will place the upgrade.xml file it generates. See "Configuration Manager command-line reference" on page 429 for details.

7. At the command prompt, check the completion status of the deployapplication.xml file by running Configuration Manager either interactively, or silently (without prompts) by appending the flag `-silent` to the command:

   ```
   configmgr checkStatus -task DeployApplication -path mypath
   ```

8. (WebSphere and WebLogic only) Start the application server.

9. (JBoss only) Stop and start the application server.

Task 8: Install the latest Content Engine Client files on other IBM FileNet P8 servers (for staged upgrades)

When you upgrade Content Engine Server software, you must also upgrade the associated Content Engine Client files installed on other machines running IBM FileNet P8 components, for example, on Process Engine and Application Engine servers. In a standard upgrade of your IBM FileNet P8 system, you would install these client files in the course of upgrading each respective component.

However, if you are staging your IBM FileNet P8 upgrade over a period of time and not upgrading one or more of the other IBM FileNet P8 components at this time, you must still perform the Content Engine Client install on the other component machines. In this case, use the following topics in the sections for the other components:

- "Upgrade the Content Engine Client files on Process Engine servers" on page 310
- "Install the latest Content Engine Client files on Application Engine servers" on page 346

NOTES

- You will have this same client-installation requirement for any expansion products that use Content Engine Client files, such as:
  - Workplace XT
  - Records Manager
  - Business Process Framework
- You must redeploy any components that are web-application-server based, such as Application Engine, Workplace XT, and Records Manager, after you upgrade the Content Engine Client files on the associated machines.

Task 9: Configure storage devices for upgrades from 3.5.*x*

Perform the procedures in this topic to prepare your storage devices when upgrading from Content Engine version 3.5.*x*.

To prepare file store devices when migrating to UNIX

Install and configure your choice of software to permit the Content Engine UNIX server to mount your existing version 3.5.*x* file store devices.

To prepare NetApp SnapLock Volumes for the upgrade

If your existing FileNet P8 domain includes a SnapLock fixed content device, complete the following procedure. Perform this procedure for each of your NetApp filers (network-attached appliances for data storage) to allow access to your fixed content via NFS instead of, or in addition to, CIFS.

1. Check the prerequisites.

   a. Check the *IBM FileNet P8 Hardware and Software Requirements* to verify your NetApp filers use the version of the Data ONTAP operating system that is supported in version 4.5 of IBM FileNet P8. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

      For information on accessing and configuring your NetApp filers, consult the following Data ONTAP manuals:

      - System Administration Guide
      - Software Setup Guide
      - File Access and Protocols Management Guide

   b. Make sure you are licensed to use NFS to access your NetApp filers.

2. Set the security style of a NetApp storage volume to enable support for NFS clients.

   Each qtree (virtual subvolume of a storage volume) has exactly one of the security styles (scheme for setting security on files and directories in the qtree) shown in the following table:

Security Style	Description
UNIX	UNIX file permission attributes. Only NFS clients can create files and directories in a UNIX qtree.
NTFS	Windows access control lists. Only CIFS clients can create files and directories in an NTFS qtree.
Mixed	Both UNIX and NTFS security styles. Only one security style at a time is allowed. The current style is that of the last client to modify it.

Since all pre-4.0.0 versions of IBM FileNet P8 support only CIFS, all NetApp storage volumes used by Content Engine Server use NTFS security style, which you must change to UNIX or Mixed. Specify the Mixed style for qtrees that must service requests for both NFS and CIFS clients during the upgrade process; otherwise, specify UNIX style.

For each qtree, perform the following steps to specify the security style:

a. Access the Data ONTAP administrative console (see the ONTAP System Administration Guide for information on administrative access methods).

b. Run the qtree command as in the following example, which sets the UNIX security style on the qtree /vol/vol1/sa1 of NetApp filer *NAFiler* the security style to UNIX:

```
telnet NAFiler
qtree security /vol/vol1/sa1 unix
```

After the qtree command executes, all files created by Content Engine Server on a UNIX platform will have UNIX security attributes.

3. Map UNIX users and groups to Window equivalents.

Because version 3.5.*x* of Content Engine Server support CIFS, rather than NFS, all existing files on a NetApp volume have NTFS security attributes; and the users and groups with access rights to these files are defined by Windows.

To allow version 3.5.*x* files to remain accessible, you must create a mapping between the new UNIX account for version 4.5 of Content Engine Server and the old Windows account for version 3.5.*x*.

Each NetApp filer has its own configuration file, /etc/usermap.cfg, to map between Windows user names and equivalent UNIX user names. A UNIX user attempting to access a file having NTFS security attributes uses usermap.cfg to determine if a mapping exists between the UNIX account and an equivalent Windows account. If the mapping exists, the access checks on the target file will use the Windows account.

Each usermap.cfg entry has the following format:

[*IP_qualifier*:] *Windows_name* [*direction*] [*IP_qualifier*:] *UNIX_name*

The meaning of each element in the entry is shown in the following table:

Element	Meaning
IP_qualifier	Qualifies the name according to the source address of the requester
Windows_name	The name of the Window user or group in domain name format (for example, *DomainName**UserName*). The Windows name must be in the Windows domain that the NetApp filer is configured to use when authenticating Windows users.

Element	Meaning
UNIX_name	The name of a UNIX user or group. The name must be defined in the file or directory service that the NetApp filer uses to authenticate UNIX users. In many cases this will be the local /etc/passwd (for users) or /etc/group (for groups). If it is a group, it may be necessary to also define it in an NIS repository or an LDAP directory server, depending on how the filer is configured. In either case the UID (UNIX user ID) must be identical to the UID of the user under which the Content Engine Server is executing.
Direction	The direction of the mapping, either <= or =>. <=: Maps *UNIX_name* to *Windows_name* =>: Maps *Windows_name* to *UNIX_name*

For example, the following steps define a mapping on a NetApp filer between FNCE_OS_User (the Windows user account under which version 3.5.*x* of Content Engine Server executes) and FNCE_UNIX_User (the UNIX user account under which version 4.5 of Content Engine Server executes).

a. Log on to the machine where version 3.5.*x* of the Content Engine is installed.

b. Connect to the root volume on the target NetApp filer using the Administrator account.

   By default, the NetApp filer root volume is accessible from a Windows client as a CIFS share named *C$*), as in the following example (where *NAFiler* is the host name of the target NetApp filer at your site.

```
C\> net use n: \\NAFiler\C$ /user:NAFiler1\Administrator
```

c. Edit /etc/usermap.cfg by adding the following stanza:

```
CEDomain\FNCE_OS_User => FNCE_UNIX_User
```

d. Edit /etc/passwd by adding the following stanza:

```
FNCE_UNIX_User:CEServers:205:7100::/home/FNCE_UNIX_User:
```

Continue at "Upgrade FileNet Enterprise Manager" on page 248.

Task 10: Upgrade FileNet Enterprise Manager

If you did not upgrade FileNet Enterprise Manager as part of a Content Engine upgrade, perform the following procedure to upgrade FileNet Enterprise Manager interactively or silently.

To upgrade FileNet Enterprise Manager

1. On the machine where you will upgrade FileNet Enterprise Manager, log on as a member of the Local Administrators group or the Power Users group.
2. Access the Content Engine software package.
3. Start the FileNet Enterprise Manager installation. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Content Engine installation program:

   – Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select CE installer.

   – Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Upgrade."

   The wizard will uninstall the current version of the software and install the new version in the same location.

   - To install interactively:

     i. Run the following command in the software package:

     P8CE-4.5.0-Win.exe

     ii. When prompted, specify the following components to be installed:

     – .NET Clients (required for FileNet Enterprise Manager)

     – FileNet Enterprise Manager

     iii. Specify the install path, for example C:\Program Files\FileNet\Content Engine.

   - To install silently:

     i. Open the silent_installer.properties file in the software package for editing.

     ii. Set the parameter values in the silent_installer.properties file for your site. Be sure to set the CHOSEN_INSTALL_FEATURE_LIST parameter value to:

     ```
     DotNetClients,AdminTools
     ```

     iii. Save your edits.

     iv. Run the following command in the software package:

     ```
     P8-CE-4.5.0-WIN.EXE -f silent_installer.properties
     ```

Task 11: Establish the FileNet P8 domain and Global Configuration Data (GCD) for 3.5.*x* upgrade

Perform the procedures in this topic only if you are upgrading from Content Engine version 3.5.*x*.

With Content Engine installed and deployed, you will use Enterprise Manager to create a FileNet P8 domain.

To create a FileNet P8 domain

NOTE If you run Enterprise Manager as a limited user account, you cannot update the *Base URL for the FileNet P8 Platform help files* field in the General tab of the Enterprise Manager properties dialog box.

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn 4.5 on the desktop, or by choosing **Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn 4.5**.
2. In the FileNet P8 Logon dialog box, click **Add**, to create a FileNet P8 domain configuration.

   NOTE For subsequent logons to Enterprise Manager, you can access an existing FileNet P8 domain by clicking **Connect**.
3. In the Add Domain Configuration dialog box, specify the required information and then click **OK**. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   In a non-managed application server environment (such as a JBoss cluster or a standalone WebSphere or WebLogic), set *ce_server_name* and *port* for just one of the application servers. You can configure multiple Enterprise Manager connections to the other non-managed application servers if you wish.
4. Continue at "To configure directory service authentication" on page 249.

To configure directory service authentication

1. In the FileNet P8 Logon dialog box, click **Connect**.
2. Specify the required values in the Create a Directory Configuration wizard screens. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   The configuration parameters required by the Create a Directory Configuration wizard are in many cases the same as those you provided to Configuration Manager when you configured LDAP in "Configure Content Engine instances" on page 220.

   Refer also to the topic for your directory service within the IBM FileNet P8 help topic System Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers.

   NOTE For multi-realm authorization, run the Directory Configuration Wizard once for each realm. Refer to FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms.

3. In the Configure New Domain Permissions message box, click **OK** to acknowledge that the directory configuration is complete but remains in restricted mode. The Configure New Domain Permissions wizard automatically starts. Continue at "To configure permissions for a FileNet P8 domain" on page 250.

To configure permissions for a FileNet P8 domain

1. Specify the required values in the Configure New Domain Permissions wizard screens. For information on the parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   NOTE To edit the list of accounts having administrative access to the FileNet P8 domain, refer to the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to... > Add or remove a GCD administrator.

2. In the Configure New Domain Permissions message box, click **OK**.

To set the Statement Cache Size value for the GCD database

If you are using Microsoft SQL Server 2005 JDBC Driver or Oracle JDBC Driver, you need to set the Statement Cache Size parameter value to 0 for each data source you created to access the GCD database, as shown in the following steps:

1. Access the page containing the Statement Cache Size parameter:
   - (WebSphere) Navigate to the WebSphere administrative console page containing the field Statement Cache Size property. For example, in WebSphere 6.1.x, navigate to Resources > JDBC Providers > JDBC_provider > Data sources > *data_source* > WebSphere Application Server data source properties.
   - (WebLogic) Navigate in the tree view of WebLogic Administration Console to *FNCEDomain* > Services > JDBC > Data Sources > *Data_Source_Name* > Connection Pool
2. Set the Statement Cache Size to 0 and save your change.

Continue at "Upgrade Content Search Engine software from 3.5.x" on page 251.

Task 12a: Upgrade Content Search Engine software from 3.5.x

Follow this task to upgrade the Content Search Engine software and related data from 3.5.x to 4.5.x.

If you used the full-text search feature (CBR) in 3.5.x, you must install and configure the new 4.5 Content-Search Engine (Autonomy K2) software. The 3.5.x content-search indexes must be upgraded to 4.5.x index areas, now called K2 collections.

As of the 4.0 release of P8 Platform, the full-text search feature was removed as a built-in function of the Content Engine. Instead, the remote, Autonomy K2 Content Search Engine software is installed and configured through Enterprise Manager to enable full-text searching.

To upgrade Content Search Engine from version 3.5.x to 4.5.x, you will install and configure the Autonomy K2 product and then upgrade the existing indexes in conjunction with the upgrade of Content Engine data as part of the overall upgrade of the Content Engine.

Production environment upgrades

Upgrade of Content Search Engine and related Content Engine data must occur as an in-place upgrade. As such, the data is migrated to 4.5.x and is no longer usable in the previous version of the P8 Platform.

The 3.5.x index areas are upgraded at the same time as the Content Engine data is upgraded, through the Content Engine upgrader tool. When the upgrader tool runs, specific changes are made to the GCD, in relation to Content-Based Retrieval, to facilitate the installation of the new Autonomy K2 product. The tool also uses the existing Verity information to create the Verity Domain Configuration, including new locale requirements, for the Autonomy K2 product.

To perform an upgrade in a migration scenario, create a duplicate of your P8 system by performing a backup and then a restore. Using the restored system, complete the migration from 3.5.x to 4.5.x. This will leave the original production environment untouched while testing is performed on the new system.

When you’re ready to migrate the production environment to the current version, back up the system and then perform the in-place upgrade on the production environment.

Content Search Engine software upgrade

The Autonomy K2 software must be installed on each machine that is part of your Content Search Engine configuration. You must designate one machine in the configuration as the Master Administration Server. The Master Administration Server can be used as a standalone Content Search Engine, or additional K2 Administration Servers can be added. All K2 Administration servers are configured and controlled through the K2 Dashboard of the Master Administration server.

The Autonomy K2 software that underlies IBM FileNet P8 Content Search Engine has many inherent features that you might want to configure that are not discussed in the IBM FileNet

documentation. For details, see the Autonomy documentation that is installed with the Autonomy K2 Master Administration Server located at:

http://*Master Administration Server hostname*:9990/verity_docs/

The Autonomy documentation set is not searchable from the IBM FileNet P8 Help but it does have its own internal index and search functionality.

CAUTION Although the K2 Dashboard provides you with documentation for, and direct interfaces to, the K2 collections, IBM FileNet requires that you use Enterprise Manager to manage collections associated with Content Search Engine index areas (for example, to add and remove index areas).

NOTES

- Autonomy K2 was previously known as Verity, and you will see Verity still used in many of the interfaces described in the following procedures.
- The machines in your Content Search Engine configuration must run the same operating system as your Content Engine. For UNIX, it to doesn't have to be the same type.
- Where machine name variables are required, IP addresses will not validate. A valid name must be entered.
- If you unimport the style set, the original files will be deleted from your system. In this scenario, if you wish to re-import the style set, you will need to recover it from your installation package. In order to avoid this situation, you can either enter a unique name for the Style Set Alias during the initial Content Search Engine (Autonomy K2) installation, or make a backup copy of the original style set. If you enter a unique name for the style set during installation, ensure you use that name when you configure Content Engine for Content-Based Retrieval.
- Stop word files can be used to increase performance by about 30%. You can put a file named style.stp into the stylefiles directory to list words you do not want full-text indexed (for example, short words such as a, the, and). However, using a stop word file also prevents searching on these words. See the K2 documentation for more details. To create a stop word file you can typically copy a file named vdk30.stp from either the main K2 install directory or the foreign language locales package over to the main stylefile directory, and then rename it to style.stp. You must do this copy operation before you create collections.

  Indexing is case-sensitive, so the style.stp file should include capitalized versions of words in the stop word list in addition to the lower-case version. For example, use "the", "The" and "THE" if you think all of these would be encountered on a search.

Install Content Search Engine on Windows

Use this procedure to install and configure IBM FileNet P8 Content Search Engine, an optional component based on the Autonomy K2 product.

Install the Autonomy K2 software on each machine in your Content Search Engine configuration. You must complete a K2 Master Administration Server software installation before you can add any additional K2 Administration Server installations. The Master Administration Server's dashboard is the central control point for configuring all additional Administration Servers that are part of the Content Search Engine configuration.

You must designate each installation as either a Master Administration Server software installation or an Administration Server software installation.

CAUTION Only one K2 Master Administration Server can be installed for each Content Engine Domain.

To install Autonomy K2 on Windows

1. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.

   NOTE Ensure *k2_os_user* has administrator privileges on this machine.

2. Set the JAVA_HOME environment variable as follows:

   a. Open the System control panel.

   b. Click the **Advanced** tab.

   c. Click **Environment Variables**.

   d. Click **New** under System Variables.

   e. Set the variable information as follows:

      Variable name: *JAVA_HOME*

      Variable value: *Java (JDK) install path*

   NOTE The installer will not allow you proceed with the installation until the JAVA_HOME environment variable is set.

3. Install the Content Search Engine software, using the appropriate values from your worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - Interactively

     i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-WIN.EXE file.

     ii. On the Select Autonomy K2 Server Type panel, choose whether to install the Master Administration Server software or the Administration Server software.

   - Silently

     i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.

ii. Launch the Content Search Engine installer by executing the following command:

```
P8CSE-4.5.0-WIN.EXE -f CSE_silent_install.txt -i silent
```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 6.2.1 Administration Server service.
- Verity K2 Administration Web Server service (Tomcat server).

Install Content Search Engine on UNIX

You must install the K2 Master Administration Server software first. The Master Administration Server's dashboard is the central control point for configuring Content Search Engine for single-server or multi-server installations.

You must designate each installation as either a Master Administration Server software installation or an Administration Server software installation.

CAUTION Only one K2 Master Administration Server can be installed for each Content Engine Domain.

To install Content Search Engine on UNIX

1. Configure permissions for the host machine so that permissions of the user who runs the IBM FileNet Content Search Engine installation are root. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic. If *k2_os_user* can not have root privileges on this machine, you can set permissions as follows, to allow *k2_os_user* to run the installation:

   a. Access the host machine and log on as a user with root privileges.

   b. Enter the following commands to set the vspget program's sticky bit such that the service runs as root and is in the same group as *k2_os_user* (default path shown):

   ```
   chown root /opt/verity/k2/_rs6k43/bin/vspget
   chmod g+rs /opt/verity/k2/_rs6k43/bin/vspget
   ```

2. Access the host machine and log on as *k2_os_user*. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.

3. Set the following environment variables and place the entries in the .profile file for *k2_os_user*.

   All UNIX types

   ```
   JAVA_HOME=java_(JDK)_install_path/jdkversion
   export JAVA_HOME
   ```

HP-UX

```
PATH=$PATH:/verity_install_path/k2/_hpux/bin
export PATH
SHLIB_PATH=$SHLIB_PATH:/verity_install_path/k2/_hpux/bin
export SHLIB_PATH
```

AIX

```
PATH=$PATH:/verity_install_path/k2/_rs6k43/bin
export PATH
LIBPATH=$LIBPATH:/verity_install_path/k2/_rs6k43/bin
export LIBPATH
```

Solaris

```
PATH=$PATH:/verity_install_path/k2/_ssol26/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ssol26/bin
export LD_LIBRARY_PATH
```

Linux

```
PATH=$PATH:/verity_install_path/k2/_ilnx21/bin
export PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/verity_install_path/k2/_ilnx21/bin
export LD_LIBRARY_PATH
```

4. Install the Content Search Engine software, using the appropriate values from your worksheet.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

- Interactively

  i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-*UNIX TYPE*.BIN file.

  ii. On the appropriate panel, choose whether to install the Master Administration Server software or the Administration Server software.

- Silently

  i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation.

ii. Execute the following command:

```
./P8CSE-4.5.0-UNIX_type.BIN -f CSE_silent_install.txt -i silent
```

NOTE The following two Autonomy K2 services will be installed and started on your machine after the installation is complete:

- Verity K2 6.2.1 Administration Server service.
- Verity K2 Administration Web Server service (Tomcat server).

To start or stop the Autonomy K2 Services on UNIX

To manually start or stop the Autonomy K2 services, use the following commands, according to your environment:

HP-UX

Start Services:

```
nohup /verity_install_path/k2/_hpux/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_hpux/bin/k2adminstop
```

AIX

Start Services:

```
nohup /verity_install_path/k2/_rs6k43/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_rs6k43/bin/k2adminstop
```

Solaris

Start Services:

```
nohup /verity_install_path/k2/_ssol26/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ssol26/bin/k2adminstop
```

Linux

Start Services:

```
nohup /verity_install_path/k2/_ilnx21/bin/k2adminstart &
```

Stop Services:

```
/verity_install_path/k2/_ilnx21/bin/k2adminstop
```

Configure Content Search Engine

Use this procedure to configure services required on the K2 Master Administration Server, and on additional Administration Servers you may install for IBM FileNet P8 Content Search Engine. All servers are configured through the Master Administration Server Dashboard.

NOTES

- When naming particular servers you create with this procedure, it is a good idea to indicate the type of server you've created. Otherwise, when you configure Content Engine through the Enterprise Manager, determining which server is which could be confusing. For example, use *servername*_broker to indicate that this is a Broker Server service.
- Ensure you carefully record the server names, ports and settings that you define. Much of the following information will be required later when you configure the IBM FileNet P8 Content Engine for Content-Based Retrieval later in this guide.
- A range of ports is recommended in the Verity K2 Dashboard for each service you create. You do not have to choose a port number from within that range.
- Repeat the related step in the procedure below to add additional services. Some guidelines must be adhered to when adding additional services:
  - Multiple Brokers can be assigned, so that if one goes down the others will be used. However, each broker must have all K2 Servers (search servers) attached that are needed to access collections (index areas). The Content Engine Server will not call multiple brokers and merge the results.
  - If you add additional Index Servers and K2 Servers (search servers), they will not be activated until you enable them through Enterprise Manager. See *Configure Content Engine for Content Based Retrieval* for details.
  - Each K2 Administration Server must contain a Ticket Server for Content Engine.
  - For good stability and performance, Broker Servers must be attached to local Ticket Servers for security on each machine.

To configure Content Search Engine

1. Configure the Autonomy K2 Dashboard to use SSL security. The Autonomy K2 Dashboard web application, by default, uses a non-SSL web site and sends username and password information in plain text. For information on how to modify your Tomcat web applications to use SSL, access the following address:

   http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html

2. Access the K2 Dashboard by launching your browser and entering:

   http://*Master Administration Server hostname*:9990/verity_dashboard/main.jsp

3. Create a K2 Index Server, as follows:
   a. Click **K2 Index Servers** under System View**.**
   b. Click **Add a K2 Index Server** on the K2 Index Server Summary page.

c. Enter the following information on the Configure basic settings for the new K2 Index Server page:
   - Service Alias: *server_name*_index_server
   - Port: 9960 - 9979 (suggested range)

d. Click **Next** to continue with the installation.

e. Enter the following information on the Configure threads for the K2 Index Server page:
   - Synchronous Threads: 25
   - Asynchronous Threads: 3
   - Access Type: **Authorized Administrator**

f. Click **Finish** to continue with the installation.

4. Set the Index Server logging properties:
   a. Click the Index Server, Under System View, that you want to adjust.
   b. Click **Edit Properties** under Actions.
   c. Click the **Logging** tab.
   d. For Status Log Nominal Size, enter the following value:

      9000 kilobytes

   e. Click **Modify**.
5. Create a K2 Broker Server:
   a. Click **K2 Brokers** under System View**.**
   b. Click **Add a K2 Broker** on the K2 Broker Summary page.
   c. Enter the following information on the Configure basic settings for the new K2 Broker page:
      - Service Alias: *servername*_broker
      - Port: 9900 - 9909 (suggested range)
   d. Click **Finish**.
6. Create a K2 Server (search server) and attach the Broker.
   a. Click **K2 Servers** under System View.
   b. Click **Add a K2 Server** under Actions on the K2 Server Summary page.
   c. Enter the following information on the Configure basic settings for the new K2 Server page:
      - Service Alias: *server_name*_search_server
      - Port: 9920 - 9949 (suggested range)
   d. Click **Next**.
   e. Click **Next** on the *Set security options for this service* page.

f. Enter the following information on the Attach to K2 Brokers page:

- Select the K2 Brokers that will provide access to this service: Select the K2 Broker you created in step 3 from the drop-down menu, *servername*_broker.

g. Click **Finish**.

7. Import the IBM FileNet Styleset.

NOTE The K2 Dashboard displays a notification that the StyleSet Editor web application cannot be accessed. This message can be ignored as it relates to a function that is not utilized by IBM FileNet Content Search Engine.

a. Click **Collections** under System View.

b. Click **Manage Style Sets** under Actions on the Collection Summary page.

c. Click **Import** on the Manage Style Sets page (top right).

d. Enter the following information on the Import page:

- Style Set Alias: FileNet_FileSystem_PushAPI
- Gateway Type: --Auto-detect--
- Source Administration Server (if multiple servers are installed). Choose which server.
- Source Path (default listed):

  Windows

  C:\Program Files\filenet\contentengine\verity\data\stylesets\FileNet_FileSystem_PushAPI

  UNIX

  /opt/verity/data/stylesets/FileNet_FileSystem_PushAPI

e. Click **Import**.

8. Create a K2 Ticket Server.

a. Click **K2 Ticket Servers** under System View.

b. Click **Add a K2 Ticket Server** under Actions on the K2 Ticket Server Summary page.

c. Enter the following information on the Configure basic settings for the new K2 Ticket Server page:

- Service Alias: *server_name*_ticket_server
- Port: 9910 - 9919 (recommended range)

d. Click **Next**.

e. Enter the following information on the Configure the login module to use with this K2 Ticket Server page:

- Select which Login Module type to use with this K2 Ticket Server:
  - Windows
  - UNIX

  NOTE LDAP Ticket Servers are not currently supported.
- Default Domain (Windows only): Enter the domain on which this K2 Server is authenticated.

f. Click **Next**.

g. Enter the following information on the Configure the persistent store module to use with this K2 Ticket Server page:

- Select the Persistent Store Module type to use with this K2 Ticket Server: Choose File and Memory.

h. Click **Finish**.

i. (Windows only) Specify local login settings:

i. Click **Edit Properties**.

ii. Click Windows Login Module.

iii. Check **Use Local Credentials:**.

iv. Check Enable Built-in Groups.

v. Click **Modify**.

9. Set Autonomy K2 Administration Security.

a. Click the K2 Ticket Server you created.

b. Click **Manage Administration Security** under Actions.

c. Enter the following information on the Manage Administration Security page:

- Select a K2 Ticket Server to configure for administration security: From the drop-down menu, select the K2 Ticket Server you just created.
- User Name: Enter the *k2_os_user*. For UNIX installs, this is the user you logged in as to run the installation. For details on required accounts and related permissions, see the *Specify IBM FileNet P8 Accounts* topic.
- Password: Enter the authentication password.
- Default Domain (Windows only): Enter the domain on which this user and K2 Server are authenticated.

d. Click **Modify**.

K2 will authenticate the user based on the information you entered. If the check fails, an error message will indicate what failed and request that you re-enter the information.

If administrator access is successful, Autonomy K2 will close the Dashboard and require that you log on again as the Dashboard Administrator to complete the configuration.

10. Launch the K2 Dashboard and log on.
11. Restart K2 services:
    a. Under *Notifications* on the K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You must perform either a **Quick Restart** or a **Full Restart**.

       Click Home in the top-left corner of the page after each restart to view the remaining notifications. Repeat the process until there are no notifications remaining.
12. Enable additional K2 Admin Users (optional).
    a. From the K2 Dashboard home page, click **Administration Servers**.
    b. Click **Manage K2 Administrative Users**.
    c. Click **Add User** on the Manage K2 Administrative Users page.
    d. Enter the name of an authenticated user on the directory service that you want to make a K2 Administrator and click **Add**.
13. Enable security on the K2 services you have created.
    a. From the K2 Dashboard home page, click **K2 Ticket Servers**.
    b. Click your ticket server ***server_name*_ticket_server**.
    c. Click **Manage K2 Broker/K2 Server Security** in the Services Secured by this K2 Ticket Server section at the bottom of the page.
    d. Click the **K2 Servers** button on the Manage K2 Broker/K2 Server Security page.
    e. Click the service you just created which is listed in the window on the right to enable security.

       CAUTION If you have a multi-server configuration, numerous services, installed on other machines, will be listed also. Select only the service to which you want to attach a broker. Brokers must be attached to local ticket servers for Content Search Engine.
    f. Click the **K2 Brokers** button on the Manage K2 Broker/K2 Server Security page.
    g. Click the Broker in the window on the right that you want to attach to the local K2 ticket server you selected above.
    h. Click **Modify** to save your changes.

14. Restart K2 services, as follows:
    a. Click **Home** in the top left corner of the page.
    b. Under Notifications on the Verity K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You must perform either a **Quick Restart** or a **Full Restart**.

       Click **Home** in the top left corner of the page after each restart to view the remaining notifications. Repeat the process until there are no notifications remaining.

Install additional locales (optional)

Complete this procedure only if you require locales other than English.

NOTES

- The panels in the Locales installer indicate that the installer is only compatible with Autonomy K2 version 2.6.0. However, the Locales installer is compatible with version 2.6.1, the version you installed with IBM FileNet Content Search Engine 4.5.
- On one of the installer panels, you will be asked to select your installed version of Autonomy K2 from a dropdown menu. Select K2 6.2.0 as the valid entry.

CAUTION Although the Locales installer offers the option of installing to a location other than the Verity install location, do not choose this option. For Content Search Engine, you must install the locales to the k2 directory of the Autonomy K2 installation path.

Windows

1. Access the Autonomy K2 Master Administration Server machine and log on as *k2_os_user*.
2. Copy the P8CSE450WIN.zip file from the installation package to the machine and extract the contents to a local directory.
3. Stop the Autonomy K2 services.
    a. Access Component Services.
    b. Stop the Verity K2 Administration Server service and the Verity K2 Administration Web Server service.
4. Create a directory in the following specific location (default drive shown):

    C:\Program Files\Common Files\InstallShield\Universal*WinVersion*\x86*host_name*\Gen1_vpddb

    a. Replace *WinVersion* with one of the following valid options:
        - Windows XP
        - Windows 2000
        - Windows 2003
        - Windows 2008

b. Replace *host_name* with the name of the Master Administration Server machine.

For example:

C:\Program Files\Common Files\InstallShield\Universal\Windows 2003\x86\myMachine\Gen1_vpddb

5. Locate the vpd.script file in the K2 install directory and open it in a text editor. Make the following modifications in the file:
   - Replace the instance of <*K2InstallDir*> with the K2 install directory. For example:

     C:\Program Files\FileNet\ContentEngine\verity
   - Replace the instance of <*myHostName*> with the Master Administration Server machine name.
   - Replace the instance of <*WinVersion*> with the same option you chose in step 4 above.
6. Copy the vpd.script file to the directory location you created in step 4 above.
7. Navigate to the decompressed Locales installer location and execute the setupwin32.exe file.The installer will locate the K2 installation and start, based on the settings you completed above.

   NOTE The following license key is required:

   2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS
8. Start the Verity K2 services after the installation is complete.

UNIX

1. Access the Autonomy K2 Master Administration Server machine and log on as *k2_os_user*.
2. Copy the appropriate compressed file for your platform from the installation package to the machine and extract the contents to a local directory. For example:

   P8CSE450AIX.tar.gz
3. Stop the Autonomy K2 services.

   a. Access opt/verity/appserver/bin

   b. Use the following command, according to your environment:

   HP-UX

   ```
   /verity_install_directory/k2/_hpux/bin/k2adminstop
   ```

   AIX

   ```
   /verity_install_directory/k2/_rs6k43/bin/k2adminstop
   ```

   Solaris

   ```
   /verity_install_directory/k2/_ssol26/bin/k2adminstop
   ```

   Linux

   ```
   /verity_install_directory/k2/_ilnx21/bin/k2adminstop
   ```

4. Prepare your system with the following environment variable:

   ```
   VERITY_CFG=/opt/verity/k2/common/verity.cfg
   ```

   ```
   export VERITY_CFG
   ```

5. Navigate to the decompressed Locales installer location and execute the setup*UNIX_type*.bin. For example:

   setupaix.bin

   NOTE The following license key is required:

   2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS

6. Start the Verity K2 services after the installation is complete.

To configure services on Administration Servers

Use this procedure to create and configure services on specific Content Search Engine machines (Administration Servers), other than the Master Administration Server. Services for all machines in your Content Search Engine configuration are configured through the Master Administration Server Dashboard.

1. Click **Administration Servers** under System View**.**
2. Click the Administration Server to which you want to add services.
3. Click **Add a Service** under K2 Services on this Administration Server.
4. Select the service you want to add.
5. Follow the instructions and guidelines for the appropriate service in the *To configure Content Search Engine* procedure above to complete the service addition.

Configure Autonomy K2 to upgrade 3.5.x indexes

To create a collections directory

For performance reasons, it is recommended that you create one collections directory for each index area you create in IBM FileNet P8 Content Engine. Each collections directory you create must be set to provide proper security access. The path to both the collections directory and collections temp directory must be entered in the index area properties when you create them.

Security and communication between Autonomy K2, Content Engine, and the collections directory is handled through the user accounts and permissions provided to those accounts. For details on required accounts and related permissions, see "Accounts for Content Search Engine" on page 87 and the related worksheet. For detailed information on security, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Security for integrated components and third-party products > Autonomy K2 Server > Security for Autonomy K2 Server.

Collections directory requirements for Content Search Engine

CAUTION Contrary to information outlined in the Autonomy-supplied documentation set, remote collections are not supported for use with IBM FileNet Content Search Engine. Collections must be written locally to the Autonomy K2 server. Using a remote-mounted disk that is accessed via the network (NFS, PCNFS, or CIFS) will cause stability problems under load and corrupt your collections. Any existing configurations that utilize non-local collections directories must be re-configured.

On a Windows system, the requirement to write Verity collections to local disks means that you are writing to a path such as D:\collections, and you cannot use a UNC path such as: *servername*\collections.

You can configure a second server that just reads these collections (note that you cannot have a second server that writes the collections). You must map a drive on the machine that reads collections to the file system on the machine that writes collections.

To do this (using the example above), share out the D: drive as some name other than D$, since you can't set permissions on D$. So, for example, set it as “DDrive”, and then map the D: drive on the reader box to *servername*\DDrive. Now D:\collections on the reader box references the same file system as the D: drive on the writer box.

For Windows machines, Autonomy K2 runs as a Windows service by default. Windows services do not mount mapped drives. One solutions is to run the Verity Administration service from a command line instead, so that it doesn’t run as a Windows service. To launch the service manually from a command line on a default installation, execute the following command:

```
C:\Program Files\verity\k2_61\k2\_nti40\bin\k2admin.exe" -cfg "C:\Program
Files\verity\k2_61\k2\common\verity.cfg
```

NOTE Only the Verity K2 6.1.1 Administration Server service needs to be started this way. The Verity Administration Web Server may be left as a Windows service.

Another solution is to use a tool like srvany.exe (supplied as part of the Windows Resource kit) to run a .cmd file that first maps the drives, and then issues the command above to start the k2admin.exe. The command to map drives, using the above example, is:

```
net use D: \\servername\DDrive
```

There are also third-party products available that function the same as the svrany.exe program.

The one drawback to using svrany is that although it will start the service correctly, it will not shut it down. You must use the Verity rcadmin command line tool to stop the service, or else use TaskManager to end all the processes that start with the prefix k2.

Relevant `rcadmin` commands from a command line:

- `rcadmin` to start the program
- `login` *`k2_os_user`* to log on

- `adminsignal` to initiate the shutdown and then one of the following in response to the `Type of Signal` then type one of the following:
  - `2 - Shutdown`
  - `3 - WS Refresh`
  - `4 - Restart all servers`

The service can now be restarted with the Windows Services tool.

1. Create a directory in which you will store collections (VerityIndexArea.RootDirectoryPath). This directory must be located on a disk that is local to the Verity server. Set permissions to allow access to the *k2_os_user*.

   NOTE This path must be local to the index server that will be assigned to write collections.

2. Create a temp directory (VerityIndexArea.TempDirectoryPath) which will be used by the K2 Index Server and Content Engine Server during operations.

   NOTE This path must be visible to both the Content Engine and the Autonomy K2 servers. This means that if the K2 Administration Server and Content Engine are not installed on the same machine, they both must be on a network mounted file system.

3. Provide read/write access to the collections directory for Autonomy K2 by entering the full path to the location and record the temp collections path.

   Windows

   a. Open the following K2 configuration file in a text editor (default path shown):

      C:\Program Files\filenet\contentengine\verity\k2\common\verity.cfg

   b. Modify the next available alias settings by adding the collections path, where new collections will be written. For example, change alias6, mapping6, and dirmode6 to the following:

      `alias6=path1`

      `mapping6=C:\`*`Collections_Directory`*

      `dirmode6=wr`

   c. Modify the next available alias settings by adding an entry for each 3.5.x collections path, expressed as UNC, that you want to upgrade. For example, change alias7, mapping7, and dirmode7 to the following:

      `alias7=path2`

      `mapping7=\\CE35_File_storage_share\FS_`*`object_store`*`\index\fs0`

      `dirmode7=wr`

   UNIX

   a. Open the following K2 configuration file in a text editor (default path shown):

      /opt/verity/k2/common/verity.cfg

b. Modify the next available alias settings by adding the collections path. For example, change alias6, mapping6, and dirmode6 to the following:

```
alias6=path1
mapping6=/Collections_Directory_Path
dirmode6=wr
```

NOTE The Collections_Directory_Path must be a local path and not a mount point.

4. Set file store access. Each file store that will be full text indexed must be accessible by the Autonomy K2 server that will perform the full text indexing. Permissions on the file store must be set the same as the permissions on the collections directories, allowing both the Content Engine Operating System User and the Verity Operating System User to access them. The names of the file store directories must also be the same on each server that access the directories.

   NOTE Although we are mapping remote collections, which are not supported, indexing will not be run on these indexes. The mapping is used to close the collections.

Completing the upgrade

The remaining steps required to complete an upgrade of Content Search Engine must be completed as part of the overall Content Engine upgrade and are outlined at the appropriate locations later in this guide.

Task 12b: Upgrade Content Search Engine software from 4.0.x

This task describes how to upgrade the IBM FileNet Content Search Engine software and related data from 4.0.x to 4.5.x.

If you used the full-text search feature (CBR) in 4.0.x, to upgrade you must export your current K2 server configurations on each machine, uninstall Autonomy K2, install the new version of Content-Search Engine (Autonomy K2) software and then import your saved configuration.

Production environment upgrades

Upgrading of Content Search Engine and related Content Engine data must occur as an in-place upgrade. As such, the data is migrated to 4.5.x and is no longer usable in the previous version of the P8 Platform.

The 4.0.x index areas are upgraded at the same time as the Content Engine data is upgraded, through the Content Engine upgrader tool.

To perform an upgrade in a migration scenario, create a duplicate of your P8 system by performing a backup and then a restore. Using the restored system, complete the migration from 4.0.x to 4.5.x. This will leave the original production environment untouched while testing is performed on the new system.

When you're ready to migrate the production environment to the current version, back up the system and then perform the in-place upgrade process on the production environment.

Content Search Engine Upgrade

To upgrade Autonomy K2 Server software on Windows

To upgrade Content Search Engine, perform the following upgrade procedure on all machines running Autonomy K2 software. Upgrade the K2 Master Administration Server first, before upgrading the K2 Administration Servers.

NOTE If you have a multi-server configuration, do not start any K2 services until after all machines have been upgraded.

1. Access the Autonomy K2 machine and log on as *k2_os_user*.
2. Access Windows Component Services and stop the two Verity K2 services.
3. Copy the IBM FileNet Content Search Engine installation package to the Autonomy K2 machine.
4. (Master Administration Server only) Open a command window and change directory to the location of the IBM FileNet Content Search Engine installation package. Execute the following command:

   ```
   backup_k2.bat "verity_install_path" "temporary_location_path"
   ```

   Replace *"verity_install_path"* with the path to the verity folder in the current installation path.

Replace *"temporary_location_path"* with the location in which you want to temporarily back up your Autonomy K2 configuration files, outside of the current installation path. Files from this backup will be used to restore your existing configuration after you complete the software upgrade.

For example: `backup k2.bat "C:\Program Files\contentengine\verity" "C:\temp"`

5. Access the Windows control panel and run, in reverse order of installation, the uninstall for each Autonomy-related entry that is listed in the Add or Remove Programs control panel. Ensure each has been executed before continuing.

   CAUTION Order of uninstall is crucial. Some of the entries do not contain any obvious version identification and you must click the **Change/Remove** button and wait for the Autonomy uninstaller to launch in order to identify the version number.

   NOTE This will remove only the added Patch, Service Pack and Fix Pack updates and will not uninstall the base IBM FileNet Content Search Engine software.

6. Complete the manual uninstall procedure for IBM FileNet Content Search Engine 4.0.x as follows:

   a. Remove the K2 Dashboard service:

      i. From a command window, access the following directory:

         C:\Progam Files\Filenet\ContentEngine\verity\appserver\bin

      ii. Enter the following command:

         ```
         service remove k2
         ```

   b. Remove the Autonomy K2 installation:

      i. From a command window, access the following directory:

         C:\Progam Files\Filenet\ContentEngine\Verity

      ii. Enter the following command:

         ```
         k2\_nti40\bin\vconfig -cfg "C:\Program
         Files\filenet\contentengine\verity\config.vcnf" -dir "C:\Program
         Files\filenet\contentengine\verity" -verbose -log log.txt -uninstall
         ```

         The Autonomy K2 Administration Server service and Tomcat application server will be uninstalled at the completion of the vconfig command.

7. If collections are stored anywhere in the installation path, move the *collections directory* and the *temp directory* to a temporary location outside the installation path.

8. Delete the Autonomy K2 software installation directory.

9. Access the IBM FileNet Content Search Engine installation package and Install the Content Search Engine software, using the appropriate values from your worksheet:

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

- Interactively

  i. Execute the P8CSE-4.5.0-WIN.EXE file.

  ii. On the Select Autonomy K2 Server Type panel, choose whether to install the Master Administration Server software or the Administration Server software.

- Silently

  i. Access the IBM FileNet Content Search Engine installation package and edit the CSE_silent_install.txt file to reflect the appropriate responses for your installation. Refer to your installation worksheet.

  ii. Launch the Autonomy K2 Master Administration server installer by executing the following command:

```
P8CSE-4.5.0-WIN.EXE -f CSE_silent_install.txt -i silent
```

CAUTION Ensure you install to the same path as your previous installation.

10. Access Windows Component Services and stop the two running Verity K2 services.

11. (Master Administration Server only) Open a command window and change directory to the location of the IBM FileNet Content Search Engine installation package. Execute the following command:

    `restore_k2.bat “temporary_location_path” “verity install path”`

    Replace *“temporary_location-path”* with the temporary path you entered earlier in this procedure when you ran `backup k2.bat`.

    Replace *“verity install path”* with the path to the verity folder in the Autonomy K2 installation path. This action will restore your configuration from the backup files.

    For example: `restore k2.bat “C:\temp” “C:\Program Files\contentengine\verity”`

12. If you moved collections in step 6 above, restore your *collections directory* and *temp directory* to their original locations.

13. For a single-machine configuration, Access Windows Component Services and start the two K2 services. For multi-machine configurations, don't start the services until all machines have been upgraded.

14. Access the K2 Dashboard and verify that the services from your previous installation are restored and running. For Master Administration Servers, delete the temporary directory of the backup you ran earlier (*“temporary_location_path”)*.

To upgrade Autonomy K2 Server software on UNIX

To upgrade Content Search Engine, perform the following upgrade procedure on all machines running Autonomy K2 software. Upgrade the K2 Master Administration Server first, before upgrading the K2 Administration Servers.

NOTE If you have a multi-server configuration, do not start any K2 services until after all machines have been upgraded.

1. Access the Autonomy K2 machine and log on as *k2_os_user*.

   NOTE Ensure *k2_os_user* has root privileges on this machine.

2. Stop the Autonomy K2 Administration Server service and the Tomcat application server. Use the following command, according to your environment:

   HP-UX

   `/verity_install_directory/k2/_hpux/bin/k2adminstop`

   AIX

   `/verity_install_directory/k2/_rs6k43/bin/k2adminstop`

   Solaris

   `/verity_install_directory/k2/_ssol26/bin/k2adminstop`

   Linux

   `/verity_install_directory/k2/_ilnx21/bin/k2adminstop`

3. Copy the IBM FileNet Content Search Engine installation package to the machine.

4. (Master Administration Server only) Change directory to the location of the IBM FileNet Content Search Engine installation package. Execute the following command:

   `backup_k2.sh "verity install path" "temporary location path"`

   Replace *"verity_install_path"* with the path to the verity folder in the current installation path.

   Replace *"temporary_location_path"* with the location in which you want to temporarily back up your Autonomy K2 configuration files, outside of the current installation path. Files from this backup will be used to restore your existing configuration after you complete the software upgrade.

   For example, `backup k2.sh "/opt/verity" "/opt/temp"`.

5. Uninstall all Autonomy K2 Service Packs and Fix Packs.

   a. Change directory to *verity install path*/verity/patches.

   b. Execute the patch uninstaller.

      - CLI

        Execute the following command:

        `./K2TK_patchUninstall.bin -console`

- Interactively (X terminal must be installed)

  Launch the K2TK_patchUninstall.bin file and complete the uninstall screens.

  NOTE This will remove only the added Patch, Service Pack and Fix Pack updates and will not uninstall the base IBM FileNet Content Search Engine software.

6. Complete the manual uninstall procedure for IBM FileNet Content Search Engine 4.0.x as follows:

   a. Access the following folder:

      verity_install_path/verity/

   b. Enter the following command (default install path shown):

      ```
      k2/platform/bin/vconfig -cfg "/opt/verity/config.vcnf" -dir "/opt/verity" -
      verbose -log log.txt -uninstall
      ```

      Substitute one of the following for *platform*:

      - _ssol26 (Solaris 8.0, 9.0 or 10.0)
      - _hpux (HP-UX 11i with -AA compiler flag)
      - _rs6k43 (AIX 5.2 and 5.3)
      - _ilnx21 (Red Hat Advanced Server 3.0 and 4.0, SUSE 8 and 9)

   The Autonomy K2 Administration Server service and Tomcat will be uninstalled at the completion of the vconfig command.

7. If collections are stored anywhere in the installation path, move the *collections directory* and the *temp directory* to a temporary location outside the installation path.

8. Delete the install directory.

9. Access the IBM FileNet Content Search Engine installation package and Install the Content Search Engine software, using the appropriate values from your worksheet:

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   i. Click the AutoFilter drop-down arrow in the “Installation or Configuration Program" column header and select CSE installer.

   ii. Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   - Silently

     i. Access the IBM FileNet Content Search Engine installation package and edit CSE_silent_install.txt to reflect the appropriate responses for your installation. Refer to your installation worksheet.

     ii. Execute the following command:

     ```
     ./P8CSE-4.5.0-UNIX_type.BIN -f CSE_silent_install.txt -i silent
     ```

- Interactively (X terminal must be installed)
  i. Access the IBM FileNet Content Search Engine installation package and execute the P8CSE-4.5.0-*UNIX TYPE*.BIN file.
  ii. On the appropriate panel, choose whether to install the Master Administration Server software or the Administration Server software.

CAUTION Ensure you install to the same path as your previous installation.

10. Stop the Verity K2 Administration Server service and the Tomcat service which were automatically started by the installer. See step 2 above for instructions.
11. (Master Administration Server only) Change directory to the location of the IBM FileNet Content Search Engine installation package. Execute the following command:

    `restore_k2.sh "temporary_location_path" "verity_install_path"`

    Replace `"temporary_location_path"` with the temporary path you entered earlier in this procedure when you ran the `backup k2.sh` file.

    Replace `"verity install path"` with the path to the verity folder in the Autonomy K2 installation path. This action will restore your configuration from the backup files.

    For example, `restore k2.bat "/opt/temp" "/opt/verity"`.
12. If you moved collections in step 6 above, restore your *collections directory* and *temp directory* to their original locations.
13. For a single-machine configuration, start the K2 services. For multi-machine configurations, don't start the services until all machines have been upgraded.

    To start the Verity K2 Administration Server service and the Tomcat service, use the follow commands, according to your environment:

    HP-UX

    `/verity_install_directory/k2/_hpux/bin/k2adminstart`

    AIX

    `/verity_install_directory/k2/_rs6k43/bin/k2adminstart`

    Solaris

    `/verity_install_directory/k2/_ssol26/bin/k2adminstart`

    Linux

    `/verity_install_directory/k2/_ilnx21/bin/k2adminstart`

14. Access the K2 Dashboard and verify that the services from your previous installation are restored and running. Delete the temporary directory of the backup you ran earlier (*"temporary_location_path"*).

NOTE The upgrade of Content Search Engine must be completed as part of the overall Content Engine upgrade. The steps are outlined at the appropriate locations later in this guide.

To complete the upgrade, you will enter Autonomy K2 Master Administration Server configuration and collections path details into the upgrader tool during the Content Engine upgrade.

Once the Content Engine upgrader tool has been run, the upgrade of Content Search Engine is complete.

Task 13: Upgrade Content Engine Data

Upgrading Content Engine data from version 3.5.*x* or 4.0.*x* to 4.5 involves the following major steps (detailed later in this task topic), to be done in the order shown:

1. If you did not install Content Engine Upgrader when you upgraded the Content Engine software, install the Upgrader tool now. Run the Content Engine installer on a Windows machine, and select the **Tools** option in the Choose Components screen.

   For more information on running the Content Engine installer, review the following procedures:

   - "To install or upgrade Content Engine" on page 219, for upgrades from 3.5.x
   - "To upgrade Content Engine software from 4.0.x" on page 230

2. Upgrade Content Engine data.

   a. If you will run the Content Engine Upgrader tool on a computer that does not have Content Engine installed, copy the JDBC drivers for your database to the computer where the Upgrader tool is installed. See "To install the JDBC drivers" on page 276.

   b. Edit the Upgrader tool utility file CE450Upgrader.bat. See "To edit the upgrader utility file" on page 276.

   c. Run the Upgrader tool from a command line or a graphical interface to upgrade Content Engine 3.5.*x* or 4.0.*x* items (including the GCD, object stores, file storage areas, etc.) to version 4.5.

      Even if you plan to run the Upgrader tool using the graphical user interface method, it is helpful to first read the command line interface method. Both methods involve the same basic steps:

      i. Create an XML upgrade status file.

      ii. Run the Upgrader tool, driving it from the XML upgrade status file.

      See "Command Line Interface to the Upgrader tool" on page 284 or "Upgrading data using the graphical user interface" on page 278.

3. (Upgrades from 3.5.*x* only) Complete post-upgrade Content Engine configuration.

   a. Clear the read-only attribute for NTFS file storage areas.

   b. (Optional) Move file storage areas and Content Search areas (collections) from Windows machines to UNIX machines.

NOTES

- You do not need to check in checked-out documents before running the Upgrader tool.
- The Upgrader tool must run on a Windows machine with at least 1.5 GB of available memory (as indicated by Windows Task Manager).
- If the maximum heap size of the JVM is 1 GB or more, do not run the Upgrader tool on the same machine where any other major application is running (such as the database used by Content Engine) unless the machine has at least 2 GB of RAM.

- To upgrade Content Engine 3.5.*x* or 4.0.*x* data—object stores, addons, file stores, fixed content devices (FCDs), etc.—to version 4.5, you will use the Upgrader tool, which you can run interactively, via a graphical user interface (GUI), or silently, via a command line interface (CLI). You can also switch between the two methods during the upgrade.
- If you are upgrading an IBM FileNet P8 environment containing a large number of object stores, use the following approach:
  - When using the Upgrader tool, upgrade at most 20 object stores at a time.
  - If the application server where version 4.5 of Content Engine is deployed is running on a 32-bit JVM, you should restrict each 4.5 FileNet P8 domain to no more than 50 object stores. If the 3.5.*x* or 4.0.*x* FileNet P8 domains contain more than 50 object stores, consider partitioning them into multiple 4.5 FileNet P8 domains during the upgrade process.
  - Ensure that 3.5.*x* and 4.0.*x* object stores having the same basic set of system objects (for file stores, fixed file stores, etc.) are upgraded into the same 4.5 FileNet P8 domain.
- For non-English support relating to collections, refer to *IBM FileNet P8 Platform Installing Non-English Environments Technical Notice*. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

To install the JDBC drivers

Perform this procedure only if you have installed the Upgrader tool on a computer that does not have Content Engine installed. For example, if you are upgrading a UNIX system, you must run the Upgrader tool on a Windows system to upgrade the Content Engine data. You can install the Upgrader tool on Windows without installing all of the Content Engine components.

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 object stores you will be upgrading. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
2. Depending on your database types, copy the appropriate JDBC driver file to a directory on the system that has the Upgrader tool installed. For example. copy the file to the C:\jars directory.

   CAUTION Do not copy the file to the ...WebSphere\AppServer\lib\ext directory.

To edit the upgrader utility file

1. Edit the *ce_install_path*/tools/upgrade/CE450Upgrader.bat file to increase the minimum and maximum JVM heap sizes. The recommended minimum and maximum JVM heap sizes (-Xms and -Xmx) are 512 MB and 1024 MB. If you are upgrading a system that has only a few object stores that do not have many custom objects, you can skip this step and use the default settings for the minimum and maximum heap size arguments in these commands (256 MB and 512 MB, respectively).

To increase the minimum and maximum JVM heap sizes, make one of the following edits:

Command line interface (CLI)
Change the line immediately after `CLI:` to the following (without carriage returns):

```
java -Xms512m -Xmx1024m -cp "%CLASSPATH%" %JAVA_OPTIONS% -Dwasp.location=".\wasp"
com.filenet.upgrader.ui.UpgradeUtility %*
```

Graphical user interface (GUI)
Change the line immediately after `GUI:` to the following (without carriage returns):

```
java -Xms512m -Xmx1024m -cp "%CLASSPATH%" %JAVA_OPTIONS% -Dwasp.location=".\wasp"
com.filenet.upgrader.ui.MainFrame
```

2. Edit the *ce_install_path*/tools/upgrade/CE450Upgrader.bat file to specify the database JDBC driver file settings.

   a. Uncomment the `set JDBCPATH` command for your database type by removing `rem` from the beginning of the statement line.

   b. Change *`YOURDRIVERPATH`* to the full path where you installed the appropriate JDBC driver in the procedure "To install the JDBC drivers" on page 276. For example, if *`YOURDRIVERPATH`* is C:/jars, use one of the following commands:

   DB2 for Windows, Linux, or UNIX
   ```
   set JDBCPATH=c:/jars/db2jcc.jar;c:jars/db2jcc_license_cu.jar
   ```

   DB2 for z/OS
   ```
   set JDBCPATH=c:/jars/db2jcc.jar;c:jars/db2jcc_license_cu.jar;c:jars/
   db2jcc_license_cisuz.jar
   ```

   Oracle
   ```
   set JDBCPATH=c:/jars/ojdbc14.jar
   ```

   SQL Server
   ```
   set JDBCPATH=c:/jars/sqljdbc.jar
   ```

3. Continue at one of the following methods:

   - "Upgrading data using the graphical user interface" on page 278
   - "Command Line Interface to the Upgrader tool" on page 284

Upgrading data using the graphical user interface

The GUI version of the Upgrader tool follows the same task sequence as the CLI version, but provides an interactive interface, in which each step is tied to the next, and in which you can check the accuracy of configuration parameters before attempting the actual upgrade.

To run the Upgrader tool

1. Log on to the machine where the Upgrader tool is installed as *ce_upgrader_tool_user*. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 of *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Double-click the icon for CE450Upgrader.bat to start the Upgrader tool.

   The first time the Upgrader tool is run, the Begin Upgrade screen is displayed. On subsequent starts, the tool displays the last screen that you were using.

To create a new XML upgrade status file

1. If the Begin Upgrade screen is not displayed, select **File > New**.
2. Select the Upgrade type.
3. If you are upgrading from Content Engine 3.5.*x*, enter the full path to the GCD file sysinit.dat. The default version 3.5.*x* location is *ce_install_path*\sysconfig\sysinit. If needed, click **Browse** to locate the file.
4. Enter the XML Upgrade Status File full path and file name. If needed, click **Browse** to select the location.
5. Click **Begin**.

   The XML upgrade status file is created in the location you specified, and the Provide Content Engine Authentication Information screen is displayed.

To configure authentication, FCD, and object store parameters

1. Specify the Content Engine authentication information:
   a. In the Specify Content Engine Authentication Information screen, specify the following values:
      - GCD Administrator user name
      - GCD Administrator password
      - Version 4.5 FileNet P8 domain name
      - Server URL. The default is http://*server_name*:*port_number*/wsi/FNCEWS40DIME.

b. Click **Test** to verify the authentication information. A confirmation message is displayed. Perform one of the following:

- If the Upgrader tool can connect using the values you supplied, a success message is displayed. Click **OK** to close the message box.
- If the connection validation fails, an error message is displayed. Click **OK** to close the message box, and then edit the field values to correct the problem. Repeat Step b until the validation test passes.

c. Click **Next**.

The Upgrader screen is displayed, with the Upgrade tab selected. The Upgrade tree in the left pane lists the items selected for upgrade. The status of each item is indicated with a colored ball. The Upgrade key legend at the bottom of the left pane shows the possible status states: Needs Info, Ready, Processing, Done, Failed, or Unsupported. You must supply data for each item that is in Needs Info state before the upgrade can complete.

2. Specify the property values for any fixed content devices that are in Needs Info state.

a. In the Upgrader Tree pane, expand **UpgradeItems > CommonGCD > Fixed Content Devices**. Click the **device name** to select a fixed content device that is in Needs Info state. The properties and values for the FCD are displayed in the right pane in the FCS Settings tab.

b. Enter the values for the FCD settings properties. Values with an * (asterick) are required values.

NOTES

- (CFS-IS) Leave the value of the CSMCache parameter blank. You can specify a value after the upgrade completes.

  For descriptions of these parameters, refer to IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > FileNet P8 Domain > How to... > View/ modify P8 domain properties > Fixed Content Devices (General tab).

- For SnapLock Fixed Devices where CE 4.5 is on UNIX, ensure you have prepared the UNIX NFS mount to the SnapLock device as instructed in “Configure storage devices for upgrades from 3.5.x” on page 245. The SnapLock share must be an NTFS share, must be mountable as an NFS device, requires proper user mapping. The Root Directory Path, SnapLock User Name, and Device Root Directory must reference UNIX information.
- For Centera devices and where CE 4.5 is installed on a UNIX Server, ensure that new Centera shared libraries have been properly installed on the CE 4.5 server as instructed in “Install or update ECM Centera SDK library files” on page 213.

c. Click **Save**.

d. Repeat Step a through Step c for each fixed content devices that is in Needs Info state.

3. Specify the property values for any object stores that are in Needs Info state.

   a. In the Upgrader Tree pane, expand **UpgradeItems > Object Stores**. Click the **object store name** to select an item that is in Needs Info state.

      Each object store may have one or more items in Needs Info state. You must supply the required values for the Database, GCDObjectStore, P8CoreObjectStoreAddons, FileStores, FixedFileStores, and DBStorageArea as needed for each object store. FileStores, FixedFileStores, and DBStorageArea are displayed only for object stores that have these features enabled. If the object store has CBR (Content Based Retrieval) enabled, then selecting the DBStorageArea node presents a panel for entering CBR information.

   b. Specify the database properties for the selected object store.

      i. Click **Database**. The Database Settings tab opens in the right pane.

      ii. Specify the following values:

         - Database user name
         - Database password
         - Database type
         - Database class
         - Database URL

         The following table shows examples of URLs for each database type:

Database Type	URL
DB2	jdbc:db2://*db_hostname*:*port*/*database_name* The default port is 50000.
MS SQL Server	jdbc:sqlserver://*db_hostname*:*port* The default port is 1433.
Oracle	jdbc:oracle:thin:@*db_hostname*:*port*:*Oracle_instance_SID* The default port is 1521.

      iii. Click **Test** to verify the connection to the database. A confirmation message is displayed. Perform one of the following:

         - If the Upgrader tool can connect using the values you supplied, a success message is displayed. Click **OK** to close the message box.
         - If the connection validation fails, an error message is displayed. Click **OK** to close the message box, and then edit the field values to correct the problem. Repeat Step iii until the validation test passes.

c. Specify the Java Naming and Directory Interface (JNDI) properties for the selected object store. Click **GCDObjectStore** to enter the Java naming and Directory Interface (JNDI) properties for the selected object store. Specify the following values:

   - Local JNDI name (the non-XA value)
   - Global JNDI name (the XA value)

d. Specify the File store settings for the selected object store. Click **FileStores** to expand the list of file stores.

   i. Click the **name** of a file store that is in Needs Info state. The File Store Path tab opens in the right pane.

   ii. Set the Root path to its UNIX equivalent. For example, `/opt/CE_35_filestore_mount/FS_One`.

   iii. Click **Save** if you made any changes.

   iv. Repeat Step i through Step iii for each file store.

e. Specify the DBStorageArea content-based retrieval properties for the selected object store.

   i. Click **DBStorageArea**. The Autonomy K2 tab opens in the right pane.

   ii. Specify the following values:

      - K2 security user
      - K2 security user password
      - Style set alias
      - Index server names
      - K2 brokers
      - User domain
      - User group
      - Server host name for the K2 master administration server
      - Server port for the K2 master administration server
      - K2 server names

      NOTE To specify multiple Index Server names, K2 Broker Server names, or K2 Server names, press **Enter** after typing each value to bring the cursor to a new line. Do not use any other delimiters (such as commas or spaces) to separate your values.

iii. Click **Test** to verify the connection to the K2 server. A confirmation message is displayed. Perform one of the following:

- If the Upgrader tool can connect using the path you supplied, a success message is displayed. Click **OK** to close the message box.
- If the connection validation fails, an error message is displayed. Click **OK** to close the message box, and then edit the field values to correct the problem. Repeat Step iii until the validation test passes.

f. If an object store is not yet ready to upgrade, clear the check box for that object store. The selected object store must be in Ready state to be upgraded.

g. Take all object stores selected to upgrade offline. Click **Take Offline** in the Description/Online/Offline tab of the right pane.

h. Repeat Step 3 as needed until all desired object stores are either in Ready state or disabled for the upgrade.

Perform the Upgrade

For an item to be qualified for upgrading, two conditions must be met:

- The check box for the corresponding node in the Upgrade Tree pane must be selected.
- The item must be in Ready state.

Each selected check box maps to a *yes* value of `SelectedForUpgrade` within the XML upgrade status file for the corresponding item.

When you click **Start** to initiate an upgrade, the icon next to each item in the Upgrade Tree pane changes color as its state changes. The color code for these states is indicated in the Upgrade Key at the bottom of the Upgrade Tree pane. The icon for any item whose upgrade fails turns red.

As the item upgrade proceeds, its corresponding Status value in the XML upgrade status file changes to reflect the current state of the upgrade. The contents of the Log pane also show how the upgrade progresses.

If you click **Stop**, the upgrade will stop after completing the current step for the item being upgraded. If an error happens during the upgrade, the corresponding node will be in UpgradeFailed state. Before trying again to upgrade an item whose previous upgrade attempt failed, perform the following:

1. Fix the error in the item before retrying the upgrade.
2. With the item node selected, click **Reset** to reset the Status value of the item to ReadyToUpgrade.

You can also click **Stop** if the log indicates a condition that makes it pointless to try to upgrade any more of the items selected for upgrade.

To upgrade items

1. In the Upgrade Tree pane, select the check boxes for all the object stores to be upgraded.
2. Click **Start** to initiate the upgrade and wait until the log or the color of the icon indicates that the upgrade is complete.
3. In the Upgrade Tree pane, select any other object stores to be upgraded, and clear the check boxes for all those object stores already upgraded.
4. Click **Start** to initiate the upgrade. A confirmation message is displayed. Click **Yes** to proceed with the upgrade, or click **No** to cancel.

   As the upgrade proceeds, the success or failure of each item being upgraded is captured in a viewable report, which you can view in the **Report** tab.

   This report is the same as that generated by running the command-line version of the Upgrader tool.

   NOTE (DB2 only) If a Transaction Log Full exception (SQL ErrorCode -964) occurs during the upgrade, it is recommended that you increase LOGSECOND (the maximum number of secondary log files) for the database and then repeat this step.

   To change LOGSECOND, run the following command:

   ```
   UPDATE DATABASE CONFIGURATION FOR db_name USING LOGSECOND n
   ```

   where *n* is the new value of LOGSECOND.

   If you encounter the exception after the log file counts have been doubled, contact your IBM service representative.
5. To view the upgrade report, click the **Report** tab near the top left of the screen. To save the report as an HTML file, click **Save** in the Report pane.
6. To view the log, click the Log tab near the top left of the screen. To copy the log data to the clipboard, click **Copy**. To clear the log data, click **Clear**.
7. If the Upgrader tool successfully upgrades every item continue at "Complete Content Search Engine Upgrade from 3.5.x" on page 289. Otherwise, complete the following steps:

   a. Use the information (exception code and stack trace) in the command-line output or the log4j file to correct the error.

   b. Select the failed item and click **Reset** on the Description tab, which will change the status of the item from UpgradeFailed to ReadyToUpgrade.

c. If the object store add-on upgrade takes longer than 10 minutes and the upgrade fails with exceptions displayed in the upgrade console or log, you need to increase your application server timeout settings. Complete the following steps:

   i. Increase your application server transaction timeout value. Make a note of your existing values before you make changes. Depending on your hardware, you may need to increase the timeout value to as much as 6000 seconds.

      WebSphere - Navigate to **Application Servers > *servername* > Transaction Serivces > Runtime**, and increase the values for *Total transaction lifetime timeout* and *Maximum transaction timeout*.

      WebLogic - Increase the value for ***Domain* > Services > JTA > Timeout Seconds** and the value for **Servers > *CEServerInstanceName* > Configuration > Tuning > Stuck Thread Max Time**. Use the same value for both settings.

      JBoss - Edit the jboss-service.xml file to set the value for <attribute name="TransactionTimeout">.

   ii. Restart the application server.

   iii. Restart the upgrade for the object store add-ons.

   iv. Once the upgrade completes successfully, you can reset the application server timeout settings to the previous values.

d. Return to Step 1.

Command Line Interface to the Upgrader tool

In this section you will create an XML upgrade status file and run the Upgrader tool using the command-line interface (CLI) method.

To create an XML upgrade status file

This procedure interrogates the GCD of the Content Engine 3.5.x installation to produce an XML file containing representations of the items that can be upgraded.

The XML file contains placeholders, which you will need to manually edit, for system settings that cannot be derived from the GCD, such as authentication settings. Each placeholder is indicated by the character string '###'.

1. Log on to the machine where the Upgrader tool is installed as *ce_upgrader_tool_user*. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 of *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Navigate to the *ce_install_path*\tools\upgrade directory, which contains CE450Upgrader.bat.
3. Note the path *GCD_Path* to the GCD file sysinit.dat on the Content Engine server machine.
4. Designate a path *XML_Path* for the XML upgrade status file to be generated in Step 5.
5. Run the following command to create the XML upgrade status file upgrade.xml:

```
CE450Upgrader.bat -i"GCD_Path/sysinit.dat" -o"upgrade.xml"
```

6. Manually edit upgrade.xml as required for your site, as follows:
   - Replace each occurrence of the string '###' with information that is appropriate for your site.
   - The passwords you specify must be in plain text. The Upgrader tool will encrypt these passwords, as well as any other sensitive data, such as passwords in fixed content device parameters.
   - A 'yes' value of SelectedForUpgrade means you want to upgrade the item; a 'no' value means you do not want the item to be considered for upgrade.
   - The Status attribute values have the meanings shown in the following table:

Status Value	Description
NeedsInfo	Additional configuration information is needed. Set this field value to ReadyToUpgrade only after you have supplied the required information.
ReadyToUpgrade	The item will be upgraded if its SelectedForUpgrade value is *yes*.
UpgradeStarted	The Upgrader tool has started upgrading the item.
UpgradeFinished	The Upgrader tool has successfully upgraded the item.
UpgradeFailed	The Upgrader tool has failed to upgrade the item.
Unsupported	Devices (FCDs only) not supported by the Upgrader tool

   - An item will be upgraded only if its SelectedForUpgrade value is 'yes' and its Status value is 'ReadyToUpgrade'.
   - If you do not want to include Content Search Engine as part of your Content Engine upgrade, you need to first unconfigure CBR from the 3.5.*x* object store using the CE 3.5.*x* Enterprise Manager Administration, and then re-run the Upgrader tool.
7. If no file store or FCD contains full-text (CBR) indexes (Verity collections) to upgrade, then its preparation is complete; otherwise continue at Step 8.
8. Enter the information below for the K2 variables in the XML file. Autonomy K2 security account information is required.

Parameter	Description
UserName	K2 Security User
UserPassword	K2 Security User password
UserDomain	Domain on which the K2 services run
UserGroup	K2 Security Group
AdminServerHost Name	Name of the host on which the K2 Master Administration server is installed

Parameter	Description
AdminServerPort	K2 Master Administration Server port
StyleSetAlias	FileNet_FileSystem_PushAPIFileNet_FileSystem_PushAPI
K2ServerNames	K2 Server names
IndexServerNames	K2 Index Server names
Brokers	K2 Broker Server names

To specify multiple Index Server names, K2 Server names, or K2 Broker Server names, press **Enter** after typing each value to bring the cursor to a new line. Do not use any other delimiters (such as commas or spaces) to separate your values.

9. Locate the directories containing collections and set permissions to allow access to the following users:
   - Content Engine Operating System User
   - K2 Operating System User
10. For each directory containing 3.5.x content-search indexes (Verity collections) that you want to upgrade to 4.5 index areas (K2 collections), make collections readable by Verity by opening the Verity.cfg file and entering the following information:
    - alias: Path number that increments for each path you list.
    - mapping: full path to the collections directory.
    - dirmode: permission value that you must set to *wr* (write and read).

    a. Open the following Verity configuration file in a text editor:

       C:\ProgramFiles\filenet\contentengine\verity\k2\common\verity.cfg

    b. Modify the next available alias settings by entering the information listed above for each collections directory you will upgrade.

       For example, if the next available settings are number 6 and you want to upgrade collections on *myserver*, located in *FileStores\myfilestore\index*, you would change alias6, mapping6, and dirmode6 to the following:

       ```
       alias6=path1
       mapping6=\\myserver\FileStores\myfilestore\index
       dirmode6=wr
       ```

To add another directory, *myotherserver\collections\index* for example, you would modify settings for number 7 as follows:

```
alias7=path2
mapping7=\\myotherserver\collections\index
dirmode7=wr
```

NOTE If the the upgraded 4.5 Content Engine will be on UNIX, be sure to specify the UNIX mount to the 3.5.*x* content-search indexes. For example,

```
mapping6=/opt/CE_35_filestore_mount/index
```

11. Continue at "To run the Upgrader tool using the CLI" on page 287.

To run the Upgrader tool using the CLI

You will now run the Upgrader tool, CE450Upgrader.bat, with the file upgrade.xml you generated in "To create an XML upgrade status file" on page 284 as input to drive the actual upgrade.

Before upgrading an object store, the Upgrader tool takes the object store offline. After upgrading an object store, the Upgrader tool updates the corresponding Status value in upgrade.xml.

1. Log on to the machine where the Upgrader tool is installed as *ce_upgrader_tool_user*. For details on required accounts and related permissions, see "Accounts for Content Engine upgrade" on page 184 of *Plan and Prepare Your Environment for IBM FileNet P8*.
2. Navigate to the *ce_install_path*\tools\upgrade directory, which contains CE450Upgrader.bat.
3. (Optional) To see the available options, run the Upgrader tool from a command line, as follows:

   ```
   CE450Upgrader -h
   ```

   Notice from the command output that the Upgrader tool supports the following:

   - Specify the amount of time (in seconds) for it to wait after taking an object store offline before upgrading it (using the `-d` option).
   - Generate an HTML report of the upgrade (using the `-r` option).
4. Run the Upgrader tool, specifying options shown in Step 3.

   As it attempts to upgrade each item, the Upgrader tool sends a status message to the command line and to a log4j logging system. If it fails in upgrading an item, the Upgrader tool will halt.
5. If the Upgrader tool successfully upgrades every item in Step 4, then continue at "Complete Content Search Engine Upgrade from 3.5.x" on page 289. Otherwise, complete the following steps:

   a. Use the information (exception code and stack trace) in the command-line output or the log4j file to correct the error.

   b. Edit upgrade.xml by replacing any Status value of UpgradeStarted or UpgradeFailed to ReadyToUpgrade.

   c. If the failure occurred due to an error *after* the database upgrade has completed, then restore your database from backup.

6. If the Upgrader tool successfully upgrades every item continue at "Complete Content Search Engine Upgrade from 3.5.x" on page 289. Otherwise, complete the following steps:

   a. Use the information (exception code and stack trace) in the command-line output or the log4j file to correct the error.

   b. Select the failed item and click **Reset** on the Description tab, which will change the status of the item from UpgradeFailed to ReadyToUpgrade.

   c. If the object store add-on upgrade takes longer than 10 minutes and the upgrade fails with exceptions noted in the log, you need to increase your application server timeout settings. Complete the following steps:

      i. Increase your application server transaction timeout value. Make a note of your existing values before you make changes. Depending on your hardware, you may need to increase the timeout value to as much as 6000 seconds.

         WebSphere - Navigate to **Application Servers > *servername* > Transaction Serivces > Runtime**, and increase the values for *Total transaction lifetime timeout* and *Maximum transaction timeout*.

         WebLogic - Increase the value for ***Domain* > Services > JTA > Timeout Seconds** and the value for **Servers > *CEServerInstanceName* > Configuration > Tuning > Stuck Thread Max Time**. Use the same value for both settings.

         JBoss - Edit the jboss-service.xml file to set the value for <attribute name="TransactionTimeout">.

      ii. Restart the application server.

      iii. Restart the upgrade for the object store add-ons.

      iv. Once the upgrade completes successfully, you can reset the application server timeout settings to the previous values.

   d. Return to Step 4.

Task 14: Complete Content Search Engine Upgrade from 3.5.x

After upgrading Content Engine data and Content Search Engine data, you must create new collections and remove old index areas

This procedure covers the minimum setup and configuration steps to get CBR configured and running with Autonomy K2 for an upgrade from 3.5.x. For more detail on Content-Based Retrieval and Content Engine, see the IBM FileNet P8 help topic Configure CBR found at FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Configure CBR.

NOTE Where machine name variables are required, IP addresses will not validate. In these cases, you must enter the host name for the machine.

1. Launch Enterprise Manager and log on as the GCD Administrator.
2. Access each Index Area to be removed and set the status to **Closed**. For details, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > View/change index area status.

   NOTE The Index Areas must be closed first and deleted later because they cannot be deleted until after a reindex has completed. Even though remote collections have been mapped, no indexing will be run on those Index Areas.
3. Create a new index area for 4.5 Content Search Engine collections. Launch the Create an Index Area wizard:

   a. In the tree view, expand the **Object Stores** container.

   b. Right-click the object store to which you want add an index area and select **New > Index Area**.

   c. Work through the wizard screens. Refer to the *Installation and Upgrade Worksheet* in *Plan and Prepare your Environment for FileNet P8* for the parameter names and values you will specify when running the wizard. For additional details, click **Help** in the wizard screens.

   For more information on creating indexes, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Create Verity index area.
4. Enable CBR for class definitions by activating the CBR Enable flag of the class you want available for CBR, as follows:

   a. Right-click the class you want to configure in the Enterprise Manager tree and click **Properties**.

   b. Select CBR Enabled and click **OK**.
5. Enable CBR for the class properties you want available for CBR, as follows:

   a. Right-click the class you want to configure and click **Properties**.

   b. Click the **Property Definitions** tab.

   c. Click the string property you want to enable for CBR indexing and click **Edit**.

d. Check **CBR Enabled** and click **OK**.

6. Reindex to create new collections. For details, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Reindex.

7. Delete the old index areas and reclaim the space. For details, see the IBM FileNet P8 help topic FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Remove Verity Index Area.

8. If you are not updating the Process Engine to version 4.5 at this time, ensure you install the latest Content Engine 4.5.x client files on the Process Engine server. See "Upgrade the Content Engine Client files on Process Engine servers" on page 310.

9. If you're not updating Application Engine Servers to version 4.0.2 at this time, ensure you install the latest Content Engine 4.5. client files on the Application Engine server. See "Install the latest Content Engine Client files on Application Engine servers" on page 346.

Upgrade and configure Process Engine

Perform the following tasks to upgrade your Process Engine installation:

1. "Complete pre-upgrade Process Engine configuration" on page 292.
2. Perform the Process Engine upgrade. See one of the following tasks:
   - "Upgrade Process Engine Interactively" on page 300
   - "Upgrade Process Engine silently" on page 305
3. "Install Process Engine software updates" on page 309
4. "Upgrade the Content Engine Client files on Process Engine servers" on page 310
5. "Install the latest Process Engine Client files on other IBM FileNet P8 servers (for staged upgrades)" on page 314
6. "Complete post-upgrade Process Engine configuration" on page 315

Task 1: Complete pre-upgrade Process Engine configuration

A number of steps must be completed prior to installing the Process Engine software. The steps vary depending on what release you are upgrading from.

To execute pre-upgrade steps for upgrades from PE 3.5.x

NOTES

- Verify that you have reconciled the Process Engine user security information. See "To prepare Process Engine for upgrade" on page 179 of *Plan and Prepare Your Environment for IBM FileNet P8*.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to upgrade Process Engine but can be re-enabled after the upgrade is complete.
- Determine when to execute the upgrade SQL script for Oracle databases.This script must be run either:
  - manually, before running the Process Engine installation program.

    or
  - automatically, from the Process Engine installation program, allowing the installation program to prompt for the sys password for Oracle in an xterm window.

    or
  - 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 of *Plan and Prepare Your Environment for IBM FileNet P8* for detailed information on the script and modes of execution.
- (UNIX only) Verify that the raw partitions used for the SEC database have been expanded to 64MB. If the partitions have been expanded, Process Engine installation will automatically update the SEC database to use the expanded partitions.
- Purge the event logs and statistic records in the Process Engine database. Clear the event logs and statistics records that are no longer required for Process Engine tracker items or workflow milestones. Clearing these items reduces the amount of time required to upgrade the Process Engine database.
- If you are using a production Process Analyzer, verify that all Process Engine events have been transmitted to Process Analyzer and that the Process Analyzer events have been published. If the events have not been transmitted and published, either before purging from the Process Engine database, or before the Process Engine database upgrade, they will not be available to the Process Analyzer.

  WARNING Do not start the Process Task Manager until specifically told to do so. When you start Task Manager, the database is accessed, which should not be done until the database

has been upgraded using the procedures in "Complete post-upgrade Process Engine configuration" on page 315.

To create the Process Engine ODBC data source and test the connection (SQL Server only)

These steps are not necessary if you are already running a SQL Server 2005 database.

If SQL 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.

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 dropdown 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.
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 the 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 verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x

The steps in this procedure apply only to customers who are using a production Process Analyzer, they are not necessary for a simulation Process Analyzer. If you are not using Process Analyzer, proceed to "To purge event logs and statistics records" on page 298. If you are using Process Analyzer you must execute these procedures before proceeding to "To purge event logs and statistics records" on page 298.

The Process Analyzer gets its data from the Process Engine database. All generated events must be transmitted from Process Engine to Process Analyzer before the upgrade. The following steps must be taken to verify that transmission is complete. Some of these steps must be taken on the Process Engine database, some on the Process Engine server using the vwtool utility, and some on the Process Analyzer VMAEDM database.

You will query the Process Analyzer VMAEDM database for a date/time value. Date/time values are stored in the Process Engine and Process Analyzer databases in different formats. When a value has been acquired from the Process Analyzer database, you will use vwtool to convert the value to the appropriate format for the Process Engine database. You will then execute a query on the Process Engine database using the converted date/time value as one parameter in the SQL query.

1. Stop Process Engine applications. These applications include any applications that are generating events or running workflows.
2. Keep both Process Engine and Process Analyzer running until all the events from Process Engine are transmitted to Process Analyzer.
3. Execute the following sub-steps to acquire a date/time value from the Process Analyzer database and convert it to an appropriate format for the Process Engine database.
   a. On the Process Analyzer VMAEDM database, execute the following SQL query:

      ```
      select InstallDate from X_SchemaInfo
      ```

      This returns a date and time string, such as 09/11/2006 16:23:59. This string must be converted.

   b. On the Process Engine server, start the vwtool utility to convert the Process Analyzer data/time string to Process Engine format.

   c. At the vwtool prompt, type `convert`, as in:

      ```
      <vwtool:26>convert
      ```

      Then press **Enter**. The following choice list displays:

      t - Time number to string
      s - String to time number
      e - Error tuple to three part

p - Three part error to error tuple
l - Log event type number to string
i - User id to user name
n - User name to user id

d. At the Choice? prompt, enter:

   s

   and press **Enter**. This converts a string to a time number and returns the following information to indicate what the current date/time mask is, as in:

   Current System Mask: mm/dd.yyy hh:tt:ss

   Time Mask (CR=system mask):

e. Press **Enter** to accept the default mask.

f. When prompted to enter the time string (CR=''), type the value that was returned from the SQL query executed on Process Analyzer VMAEDM in step a. (for example, 09/11/2006 16:23:59). Your input must match the format of the current system mask from step d above. Then press **Enter**.

   A string value is returned for the date/time entered, which you must make note of for the next set of queries. For example:

   Time...[0x4505F00F].....**1158017039** => '09/11/2006 16:23:59'

g. Get a list of all regions on the disk by typing the following at a vwtool prompt:

   `regions`

h. When prompted, respond by typing:

   d

i. For every region, type the following at a vwtool prompt:

   `reg X`

   where *X* is the region number

j. Type the following at the vwtool prompt:

   `config`

k. Locate and make note of the physical table name associated with every event log.

l. On the Process Engine database, execute a SQL query (such as the following example) to verify that no untransmitted events remain in the Process Engine database. You will query for the number of records in every physical table associated with event logs, using the names you acquired in the previous step. Following is an example of the query syntax:

   `Select count(*) from f_sw.`*`physical table name`* `where F_AEXmitStat = 1 and F_TimeStamp >` *PAInstallDate*

   where :

   The *physical table name* was acquired in step k.

The *PAInstallDate* is the number returned in step f.

Note that the physical table name must be preceeded by f_sw. in the query.

The query result must be 0. If the queries do not return 0, not all events have been collected, in which case Process Engine and Process Analyzer must keep running until all the events are transmitted and the queries return 0.

To verify that all Process Analyzer 3.5.x events have been published

Process Analyzer must publish all the events in its VMAEDM database. Verify that all events have been published through the Microsoft Query Analyzer. Query for the number of rows in the F_Events table in the VMAEDM database with PAJobId = 0.

The following is an example of the query on the Process Analyzer database.

Using VMAEDM:

```
Select count(*) from F_Events where PAJobId = 0
```

The above query should return 0. If the query returns anything other than 0, then not all events have been published. In that case, you must leave Process Analyzer running until the query returns 0.

To purge event logs and statistics records

Before you upgrade Process Engine, use the vwlog utility to reduce the number of event log and statistics records in the database. This step is optional, but eliminating some of these records can significantly reduce the amount of time necessary for the upgrade to complete. Note that purging these records can take a significant amount of time, so plan this activity accordingly.

Execute the command by logging on as a member of the PE Configuration group.

The following are some examples of vwlog syntax. Do not use the -P option when you purge logging records if you are using Process Analyzer.

`vwlog -X -r` *`region number`* `(this command removes the statistics from the specified isolated)`

The following command will remove all log records from the isolated region. Use this only if all workflows have terminated and you no longer need tracking or milestone information.

`vwlog -L -r` *`region number`* `(this command removes all log records from the specified isolated region)`

There are multiple optional parameters for the vwlog utility, allowing selection of log records for deletion meeting a number of conditions, such as log records for terminated workflows, for tracker related records, and more. See the IBM FileNet P8 help topics under FileNet P8 Administration > Process Engine Administraton > Administrative tools > vwlog for additional information.

NOTE If Process Analyzer is installed, you must complete the steps detailed in "To create the Process Engine ODBC data source and test the connection (SQL Server only)" on page 293 and "To verify that all Process Analyzer 3.5.x events have been published" on page 296 before you purge event logs.

To stop all Process Engine-related services and applications

1. Complete the following procedure:

   UNIX

   Log on as root.

   Windows

   Log on as fnsw.

2. Set the PPM and any routers to manual startup if they are currently configured to autostart.
3. Stop the following components if they are running:
   - Process Simulator
   - Process Analyzer
   - Custom applications that require a router
   - Component Manager
   - Routers - for Application Engines, Content Engine and custom applications
   - Content Engine
   - Pooled Process Manager (PPM)
   - Process Service
   - Process Task Manager
4. Enter the following at a command prompt after the FileNet software is shut down.

   Windows

   `killfnsw -D -y -S`

   UNIX

   `killfnsw -DAyS` or `killfnsw -D -A -y -S`

5. (AIX only) Execute the following:

   `slibclean`

6. (UNIX only) Execute the following command to look for any java processes that are still running:

   `ps –ef | grep java | grep VW`

   Kill any P8-related running java processes.

   Proceed to either "Upgrade Process Engine Interactively" on page 300 or "Upgrade Process Engine silently" on page 305.

To execute pre-upgrade steps for upgrades from PE 4.0.x

- IBM recommends that you complete the upgrade of the entire P8 platform and verify functionality before upgrading database software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to upgrade Process Engine but can be re-enabled after the upgrade is complete.
- Determine when to execute the upgrade SQL script for Oracle databases.This script must be run in one of the following ways:
  - manually, before running the Process Engine installation program.

    or
  - automatically, from the Process Engine installation program, allowing the installation program to prompt for the sys password for Oracle in an xterm window.

    or
  - 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 of *Plan and Prepare Your Environment for IBM FileNet P8* for detailed information on the script and modes of execution.

To purge event logs and statistics records

Before you upgrade Process Engine, use the vwlog utility to reduce the number of event log and statistics records in the database. This step is optional, but eliminating some of these records can significantly reduce the amount of time necessary for the upgrade to complete. Note that purging these records can take a significant amount of time, so plan this activity accordingly.

Execute the command by logging on as a member of the PE Configuration group.

The following are some examples of vwlog syntax. Do not use the -P option when you purge logging records if you are using Process Analyzer.

```
vwlog -X -r region number (this command removes the statistics from the specified
isolated)
```

The following command will remove all log records from the isolated region. Use this only if all workflows have terminated and you no longer need tracking or milestone information.

```
vwlog -L -r region number (this command removes all log records from the specified
isolated region)
```

There are multiple optional parameters for the vwlog utility, allowing selection of log records for deletion meeting a number of conditions, such as log records for terminated workflows, for tracker related records, and more. See the IBM FileNet P8 help topics under FileNet P8 Administration > Process Engine Administraton > Administrative tools > vwlog for additional information.

To stop all Process Engine-related services and applications

1. Complete the following procedure:

   UNIX

   Log on as root.

   Windows

   Log on as fnsw.

2. Stop the following components if they are running:
   - Process Simulator
   - Process Analyzer
   - Component Manager
   - Content Engine
   - Process Service
   - Process Task Manager
3. Shut down any active windows displaying Process Engine log files.
4. Enter the following at a command prompt to stop the Process Engine software:

   ```
   initfnsw –y stop
   ```

5. Enter the following at a command prompt after the FileNet software is shut down.

   Windows

   ```
   killfnsw -D -y -S
   ```

   UNIX

   `killfnsw -DAyS` or `killfnsw -D -A -y -S`

6. (AIX only) Execute the following:

   ```
   slibclean
   ```

7. (UNIX only) Execute the following command to look for any java processes that are still running:

   ```
   ps –ef | grep java | grep VW
   ```

   Kill any P8-related running java processes.

   Proceed to either "Upgrade Process Engine Interactively" on page 300 or "Upgrade Process Engine silently" on page 305.

Task 2a: Upgrade Process Engine Interactively

Interactively upgrade Process Engine software by performing the procedures in this topic that are appropriate for your platform.

You will find references to logging on as the root and fnsw users within the following procedures. For all UNIX operating systems, the root user must run in the Bourne or Korn shell and the fnsw user must run in the Korn shell.

Before starting the Process Engine installation, verify that you have completed the steps in the “To configure your /etc/hosts file” procedure in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Several Process Engine upgrade screens will display only if the upgrade is from Process Engine version 3.5.*x*. Those screens are noted in the *Installation and Upgrade Worksheet*.

To verify the database connection

Verify the database connection between the Process Engine and the database. See “Verify the ability to connect to the database” on page 125 of *Plan and Prepare Your Environment for IBM FileNet P8* for details.

To upgrade the Process Engine software interactively (UNIX)

1. Log on to the server as the root user.

   NOTE This user does not need to be a database administrator unless you will be executing the SQL scripts from the Process Engine installation program.

2. Access the Process Engine software package.
3. From the console, launch the appropriate P8PE-4.5.0-*platform*.bin installation program.
4. Wait for files to finish unpacking.
5. Complete the Process Engine installation screens using the appropriate information from your installation worksheet. Installation program screens may indicate this is an installation, although this is an upgrade.

   HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

   a. Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select PE installer.

   b. Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Upgrade."

6. Monitor /fnsw/local/logs/wizard to check the progress of the upgrade since the installation program will run for several minutes,and its progress, though displayed, might not visibly advance for an extended period of time.
7. Reboot when prompted.

8. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	/fnsw/local/logs/PE (if the upgrade completes successfully) or /fnsw/tmp_installer (if the upgrade has errors)
IS mini-installer logs	/fnsw/local/logs/wizard and other subdirectories under /fnsw/local/logs

9. Log off as the root user and log on as fnsw (or the alias).
10. Proceed to "To set the f_maint and f_sw passwords" on page 302.

To upgrade the Process Engine software interactively (Windows)

1. Log on as a member of the local Administrators group or a user with equivalent permissions. If you plan to run the SQL scripts from the Process Engine installation program, the user you log on as must also be a database administrator. See "Specify IBM FileNet P8 accounts" on page 65 of *Plan and Prepare Your Environment for IBM FileNet P8* for information on requirements for logging on as a Windows domain user for Process Engine installation.

   NOTE This user does not need to be a database administrator unless you will be executing the SQL scripts from Process Engine installation.

2. Access the Process Engine software package, and start the P8PE-4.5.0-Win.exe installation program.

   NOTE To run the Process Engine installation from disk, you must copy the installation files to a disk volume where 8.3 name generation is enabled, or if 8.3 name generation is disabled, you must copy the installation to a path that uses only short (8.3) names.

   When running from disk, either interactively or silently, be aware that the Process Engine installation program has a 64-character path limitation when the path is expressed in 8.3 format. This limitation applies to the IMSInst subdirectory where the underlying Image Services (IS) mini-installer setup.exe is located. For example, the original path where the IS mini-installer resides is:

   \\server08\Software\InstallationDisks\FileNet\Release P8 4.5.0\ProcessEngine\Windows\IMSInst

   When expressed in 8.3 format the path might be:

   \\server08\Software\INSTAL~1\FileNet\RELEAS~1.0\PROCES~1\Windows\IMSInst

   This compressed path is 73 characters long, exceeding the 64-character limit.

3. Complete the Process Engine installation screens using the appropriate information from your installation worksheet. Installation program screens may indicate this is an installation, although this is an upgrade.

HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following steps to quickly see only the installation properties you must specify for the Process Engine installer:

a. Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select PE installer.

b. Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Upgrade."

4. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE450_setup.log C:\Program Files\FileNet\PE C:\FNSW
IS mini-installer logs	*Windir*\mini_installer.log, Windows Event logs, and log files under \FNSW_LOC\logs
Output files from SQL Script validation (only if the SQL scripts were executed from the Process Engine installation program).	C:\Program Files\FileNet\PE

5. Start the following services:
   - IMS ControlService
   - Process Engine Services Manager

To set the f_maint and f_sw passwords

For added security, Process Engine stores an encrypted version of the passwords for the f_sw and f_maint users, or their aliases, in a file called rdbup.bin. This is in addition to passwords for these users in the Oracle or SQL Server database, or the DB2 operating system user's password. The encrypted password and the database or operating system user's passwords must match.

To verify that the passwords match, use the following procedure to start the Xdbconnect utility. Xdbconnect works only if the passwords in the encrypted file and the database match.

Use the following procedure to change the passwords for the f_maint and f_sw users after upgrading the Process Engine software. For Oracle and SQL Server databases, both the encrypted file and the database passwords will be updated. For DB2, only the encrypted file will be updated.

1. Start the Database Server Connect application by executing the following:

   ```
   Xdbconnect -r
   ```

2. Log on as SysAdmin.
3. Change the primary password for the users f_sw and f_maint (or their alias) to match the database password (Oracle and SQL Server) or operating system user's password (DB2).
4. Exit the application.

To re-enable Oracle Password Complexity Verification

If, as directed earlier, you disabled the the Oracle Password Complexity Verification feature prior to upgrading Process Engine, you can re-enable it now.

To restore any custom modifications for root and fnsw users (UNIX)

Process Engine installation creates a new versions of a number of files. If the previous versions of these files contained any custom settings, edit the new files for the fnsw and root users accordingly. Saved files are in *.filename*.old.*nn*, where *nn* is a sequential number. The latest saved version will be in the highest numbered file. The following files are modified by the Process Engine installation program:

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

To edit the /etc/inittab file

By default, the Process Engine software starts automatically when you restart the server and needs its database started beforehand. If the database is not automatically started on server restart, edit the /etc/inittab file on the Process Engine machine to comment out the autostart of Process Engine. The following are examples of changes to make.

AIX

```
rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

```
#rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

HP

```
rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

#rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1

Solaris

```
fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

to

```
#fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

To edit the pe_start file (HP)

If the value for the maxdsiz kernel parameter is > 1GB, edit the pe_start file.

Change:

```
nohup /usr/ccs/lbin/dldd32 2>&1 >/dev/null
```

to

```
nohup /usr/ccs/lbin/dldd32 +a 0x70000000 2>&1 >/dev/null
```

Proceed to "Install Process Engine software updates" on page 309.

Task 2b: Upgrade Process Engine silently

Silently upgrade Process Engine software by performing one of the following procedures.

Before starting the Process Engine installation, verify that you have completed the steps in the “To configure your /etc/hosts file” procedure in the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

To verify the database connection

Verify the database connection between the Process Engine and the database. See “Verify the ability to connect to the database” on page 125 of *Plan and Prepare Your Environment for IBM FileNet P8* for details.

To upgrade the Process Engine software silently (UNIX)

Take the following steps to silently upgrade Process Engine.

1. Access the Process Engine software package, and copy the contents to a local temporary directory on the local disk.
2. Edit the PE_silent_install.txt file to reflect the appropriate responses for your system. See “Encrypt passwords” on page 453 for information on use of the password encryption tool.
3. Save the edited response file to your temporary directory.
4. Log on as the root user in the Korn shell.
5. Navigate to the temporary directory on the local disk.
6. Open a command prompt and execute:

   `P8PE-4.5.0-`*platform* `-silent -options PE_silent_install.txt`
7. Monitor /fnsw/local/logs/wizard to check the progress of the upgrade since installation will run for several minutes,and its progress, though displayed, might not visibly advance for an extended period of time.
8. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	/fnsw/local/logs/PE (if the upgrade completes successfully) or /fnsw/tmp_installer (if the upgrade has errors)
IS mini-installer logs	/fnsw/local/logs/wizard and other subdirectories under /fnsw/local/logs

9. Log off as the root user and log on as fnsw (or the alias).

To upgrade the Process Engine software silently (Windows)

1. Access the Process Engine software package, and copy its contents to a temporary directory on the local disk.
2. Edit the PE_silent_install.txt file to reflect the appropriate responses for your system. See "Encrypt passwords" on page 453 for information on use of the password encryption tool.
3. Save the edited response file to your temporary directory.
4. 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. See "Specify IBM FileNet P8 accounts" on page 65 of *Plan and Prepare Your Environment for IBM FileNet P8* for information on requirements for logging on as a Windows domain user for Process Engine installation.
5. Open a command prompt and navigate to the temporary directory. Execute:

   ```
   P8PE-4.5.0-Win.exe -silent -options PE_silent_install.txt
   ```

6. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE450_setup.log C:\Program Files\FileNet\PE C:\FNSW
IS mini-installer logs	*Windir*\mini_installer.log, Windows Event logs, and log files under \FNSW_LOC\logs
Output files from SQL Script validation (only if the SQL scripts were executed from the Process Engine installation program).	C:\Program Files\FileNet\PE

7. Start the following services:
   - IMS ControlService
   - Process Engine Services Manager

To set the f_maint and f_sw passwords

For added security, Process Engine stores an encrypted version of the passwords for the f_sw and f_maint users, or their aliases, in a file called rdbup.bin. This is in addition to passwords for these users in the Oracle or SQL Server database, or the DB2 operating system user's password. The encrypted password and the database or operating system user's passwords must match.

To verify that the passwords match, use the following procedure to start the Xdbconnect utility. Xdbconnect works only if the passwords in the encrypted file and the database match.

Use the following procedure to change the passwords for the f_maint and f_sw users after upgradinging the Process Engine software. For Oracle and SQL Server databases, both the encrypted file and the database passwords will be updated. For DB2, only the encrypted file will be updated.

1. Start the Database Server Connect application by executing the following:

   ```
   Xdbconnect -r
   ```

2. Log on as SysAdmin.
3. Change the primary password for the users f_sw and f_maint (or their alias) to match the database password (Oracle and SQL Server) or operating system user's password (DB2).
4. Exit the application.

To re-enable Oracle Password Complexity Verification

If, as directed earlier, you disabled the the Oracle Password Complexity Verification feature prior to upgrading Process Engine, you can re-enable it now.

To restore any custom modifications for root and fnsw users (UNIX)

Process Engine installation creates a new versions of a number of files. If the previous versions of these files contained any custom settings, edit the new files for the fnsw and root users accordingly. Saved files are in *.filename*.old.*nn*, where *nn* is a sequential number. The latest saved version will be in the highest numbered file. The following files are modified by the Process Engine installation program:

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

To edit the /etc/inittab file

By default, the Process Engine software starts automatically when you restart the server and needs its database started beforehand. If the database is not automatically started on server

restart, edit the /etc/inittab file on the Process Engine machine to comment out the autostart of Process Engine. The following are examples of changes to make.

AIX

```
rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

```
#rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

HP

```
rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1
```

to

#rcfn:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1

Solaris

```
fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

to

```
#fn:3:wait:/bin/sh /etc/rc.initfnsw </dev/console >/dev/console 2>&1
```

To edit the pe_start file (HP)

If the value for the maxdsiz kernel parameter is > 1 GB, edit the pe_start file.

Change:

```
nohup /usr/ccs/lbin/dldd32 2>&1 >/dev/null
```

to

```
nohup /usr/ccs/lbin/dldd32 +a 0x70000000 2>&1 >/dev/null
```

Proceed to "Install Process Engine software updates" on page 309.

Task 3: Install Process Engine software updates

Perform the procedure in this topic for the Process Engine to install software updates (fix pack or interim fix).

To install the Content Engine software updates

1. To obtain the latest Process Engine software updates, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the Process Engine software updates and perform the installation procedures in the readmes on the Process Engine.
3. Proceed to install latest Content Engine files on the Process Engine. See "Upgrade the Content Engine Client files on Process Engine servers" on page 310.

Task 4: Upgrade the Content Engine Client files on Process Engine servers

To upgrade the Content Engine Client files, perform the following steps on the machine where Process Engine Server is installed.

To stop Process Engine software

1. Log on as fnsw or the alias.
2. At a command prompt, enter the following:

   Windows

   ```
   killfnsw -D -y -S
   ```

   UNIX

   ```
   killfnsw -D -A -y -S
   ```

3. On AIX, execute:

   ```
   slibclean
   ```

To uninstall Content Engine Client 4.0.x instances

Perform the following procedure to identify and uninstall all Content Engine Client version 4.0.x files on a Process Engine Server machine. Repeat the procedure for all other Process Engine Server machines.

Depending on the Content Engine Client fix pack versions that you have installed, you may have more than one instance of the Content Engine Client 4.0.x installed. Additionally, the version 4.0.x Content Engine Client installer supported installing multiple instance of the Content Engine Client on the same server. You need to uninstall all instances before upgrading to the 4.5 version of Content Engine Client.

1. Locate and uninstall all instances of the Content Engine Client on the Process Engine server.

   Windows

   a. Select **Start > Settings > Control Panel > Add/Remove Programs**.

   b. Select **Content Engine Client Updater** or **IBM FileNet Content Engine Client Installer**, depending on the fix pack version of Content Engine Client that is installed.

      Prior to fix pack P8CE-4.0.1-006: The Content Engine Client is named Content Engine Client Updater.

      As of fix pack P8CE-4.0.1-006: The Content Engine Client is named IBM FileNet Content Engine Client Installer.

   c. Click **Change/Remove** to start the uninstaller.

   d. In the Uninstall IBM FileNet Client Installer screen, click **Next**.

   e. In the Select Components screen, select Complete Uninstall. Click **Next**.

f. In the Uninstall Complete screen, click **Done**.

g. Repeat Step b through Step f for each additional instance of **Content Engine Client Updater** or **IBM FileNet Content Engine Client Installer** listed in Add/Remove Programs.

UNIX

a. Run one of the following commands, depending on the fix pack version of Content Engine Client that is installed.

```
/ce_client_install_path/ClientUpdater/_uninst/uninstaller.bin
/ce_client_install_path/CEClient/_CEClientuninst/uninstaller
```

Prior to fix pack P8CE-4.0.1-006: The Content Engine Client files are in a directory called ClientUpdater, and the uninstallation program is .../ClientUpdater/_uninst/uninstaller.bin.

As of fix pack P8CE-4.0.1-006: The Content Engine Client files are in a directory called Content Engine Client and the uninstallation program is .../CEClient/_CEClientuninst/uninstaller.

b. In the Uninstall IBM FileNet Client Installer screen, click **Next**.

c. In the Select Components screen, select Complete Uninstall. Click **Next**.

d. In the Uninstall Complete screen, click **Done**.

e. Repeat Step a through Step d for each installed instance of Content Engine Client on this server.

To install the Content Engine Client files

Refer to the appropriate information from your installation worksheet while performing the following steps.

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. On the machine where Process Engine is installed, log on as fnsw, with these permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Process Engine is installed
   - Execute permission on the Content Engine Client install software
3. Verify that there is a current backup of Process Engine.
4. Copy the Content Engine Client install software from the Content Engine installation software to the temporary directory, and unzip the software package. The version of the installation software must match the version of Content Engine.
5. Start the Content Engine client installation process.
   - To install the Content Engine client interactively:
     i. Access the IBM FileNet Content Engine client update software.

ii. Run one of the commands in the table below, *CE_version* is the version of Content Engine you intend to run, for example, 4.5.0..

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN`

iii. Complete the installation program wizard.

- To install the Content Engine client files silently:

i. Make a back up copy of the input file listed below for your operating system:

Windows

`CEClient_silent_install_win.txt`

UNIX

`CEClient_silent_install_unix.txt`

ii. Open the silent input file in a text editor. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your update.

iii. Navigate to the path containing the Content Engine Client installation program, and run one of the commands in the following table to perform the silent installation, where:

CE_version is the version of Content Engine you intend to run, for example, 4.5.0.

path is the path that contains the installation program.

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN -f `*`path`*`/CECLIENT.AIX/ CEClient_silent_install.txt -i silent`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN -f `*`path`*`/CEClient.HPUX/ CEClient_silent_install.txt -i silent`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.EXE -f `*`path`*`/CEClient.HPUXI/ CEClient_silent_install.txt -i silent`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN -f `*`path`*`/CEClient.Linux/ CEClient_silent_install.txt -i silent`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN -f `*`path`*`/CEClient.Solaris/ CEClient_silent_install.txt -i silent`

Operating System	Install Program
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE -f`*`path`*`\CEClient.Windows\CEClient_silent_install.txt -i silent`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN -f`*`path`*`/CEClient.zLinux/CEClient_silent_install.txt -i silent`

Task 5: Install the latest Process Engine Client files on other IBM FileNet P8 servers (for staged upgrades)

When you upgrade Process Engine server software, you must also upgrade the associated Process Engine Client files installed on other machines running IBM FileNet P8 components, for example, on Content Engine and Application Engine servers. In a standard upgrade of your IBM FileNet P8 system, you would install these client files in the course of upgrading each respective component.

However, if you are staging your IBM FileNet P8 upgrade over a period of time and not upgrading one or more of the other IBM FileNet P8 components at this time, you must still perform the Process Engine Client install on the other component machines. In this case, use the following topics in the sections for the other components:

- "Install the latest Process Engine Client files on Content Engine servers" on page 239
- "Install the Latest Process Engine Client files on Application Engine servers" on page 350

NOTES

- You will have this same client-installation requirement for any expansion products that use Process Engine Client files, such as:
  - Workplace XT
  - Records Manager
  - Business Process Framework
- You must redeploy any components that are web-application-server based, such as Content Engine, Application Engine, Workplace XT, and Records Manager after you upgrade the Content Engine Client files on the associated machines.

Task 6: Complete post-upgrade Process Engine configuration

You must perform the following additional procedures to complete the upgrade of all Process Engine data and objects. Perform the procedures in one of the following two subtopics, depending on the starting point of your upgrade:

"Complete the upgrade from Process Engine 3.5.x" on page 315

"Complete the upgrade from Process Engine 4.0.2-001, 4.0.3, or higher" on page 327

Complete the upgrade from Process Engine 3.5.x

The format of the event log files changed in the 4.0 release. As a part of this upgrade, several SQL scripts must be executed for SQL Server and DB2 databases only.

CAUTION Throughout this procedure there are multiple software restarts. Execute all restarts as documented. Do not start or restart Process Task Manager or other IBM FileNet software unless specifically told to do so. When you restart Process Task Manager it accesses the database, which should not be done until the database has been upgraded.

To execute steps on a UNIX operating system, the terminal must support X Windows and the DISPLAY environment variable must be set.

SQL Server Client software is required on the Process Engine server to execute a number of SQL scripts documented in this topic if the database is a remote SQL Server. The SQL Server Client software can be removed from the Process Engine server after the Process Engine database has been successfully upgraded to the 4.5 release.

To update the Process Engine database objects

After you have updated the Process Engine software, you must update the Process Engine database objects.

1. **UNIX**

   Log on as fnsw.

   Windows

   Log on as a local administrator.

2. (Windows only) Ensure that the Process Engine services are started:
   - IMS ControlService
   - Process Engine Services Manager
3. (Windows only) Enable the redirection of log messages to the Image Services error log. This redirection logs messages to the Image Services error log as well as to the default Windows Event Log. When you enable this redirection, you can monitor the progress of the database object upgrade in a command window.

To enable the redirection, change the LogToFiles value from 0 to 1 for the following registry key.

HKEY_LOCAL_MACHINE > SOFTWARE > FileNET > IMS > CurrentVersion

4. Restart the Process Engine software by entering the following at a Windows command prompt or UNIX command line.

```
initfnsw -y restart
```

5. If you are using a SQL Server database, proceed to Step 6. If you are using a DB2 database, proceed to Step 7. If you are using an Oracle database, proceed to Step 8.

6. (SQL Server only) Edit and run the \fnsw\mssql\vwmssql35to40_pre1.bat file. Database schema changes will be made to existing event log database tables.

   a. Save the file to the same directory as vwmssql35to40_pre1a.bat.

   b. Change the values in the file as appropriate for your system. The content of the vwmssql35to40_pre1.bat file looks like this:

   ```
   osql /U sa /P /n /d VWdb /h-1 /i vwmssql35to40_pre1.sql
   ```

   Change the values for your system to:

   ```
   osql /D DSN /U sa /P sa /n /d VWdb /h-1 /i vwmssql35to40_pre1.sql
   /o pre1a.log
   ```

   where:

   `/D` indicates the following variable is your ODBC data source name (DSN).

   `/U` indicates the following variable is the administrator user name in the Process Engine database.

   `/P` indicates the following variable is the administrator user's password in the Process Engine database.

   `/d` indicates the following variable is the Process Engine database name.

   Optionally, you can add an output file /o pre1a.log. Otherwise, all output goes only to the screen.

   c. Run the vwmssql35to40_pre1a.bat file.

   d. Proceed to Step 8.

7. (DB2 only) Edit and run the \fnsw\DB2\vwdb2_35to40_pre1.bat file on Windows or \fnsw\DB2\vwdb2_35to40_pre1.sh on UNIX platforms. Database schema changes will be made to existing event log database tables.

   a. Copy the file to the same directory as vwdb2_35to40_pre1a.bat or vwdb2_35to40_pre1a.sh.

   b. Change the values in the file as appropriate for your system.

   Change the values for your system to:

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

where:

database_name is your Process Engine DB2 database name

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

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

c. Run the vwdb2_35to40_pre1a.bat under from the DB2 command line processor on Windows, or run vwdb2_35to40_pre1a.sh from a command prompt on UNIX.

d. Proceed to Step 8.

8. Initiate the database Process Engine schema changes by executing the following command:

vwtool

and then complete the following substeps:

a. Choose Yes when a message is presented indicating that an upgrade is required.

b. Reply No when prompted as to whether you want to initiate tracking to capture the changes make to a trace file.

NOTE A number of messages will scroll to the window. Do not choose X in response to the prompt as this will terminate the current vwtool command and return an error.

c. Type exit to end vwtool when the upgrade is complete.

When vwtool starts, it automatically checks the Process Engine database level, updates the schema accordingly, and creates two additional scripts.

SQL Server

fnsw\mssql\vwmssql35to40_post1.sql

fnsw\mssql\vwmssql35to40_post2.sql

DB2

fnsw\DB2\vwdb2_35to40_post1.sql

fnsw\DB2\vwmdb2_35to40_post2.sql

CAUTION If the upgrade fails at this point, you must restore the Process Engine database backup.

See the Image Services error log to monitor the progress of the updates and ensure that no errors occur.

Check the logs to verify that messages similar to the following are captured:

```
2006/10/17 16:23:43.261  <fnsw> VW/Process  (14952) ... [INFO]
VW: Database upgrade successful to version 46, please follow instructions to perform
the next step.

2006/10/17 16:23:43.303  <fnsw> VW/Process  (14952) ... [INFO]
VW: Must restart software to complete upgrade procedure

Please follow the upgrade documentation to continue with the upgrade procedures.
```

You must:

1. Restart the Process Engine software.
2. Configure the Process Engine connection to the Content Engine.
3. Run vwtool to continue the upgrade procedure.

Ignore messages designated as SERIOUS if they are in combination with a successful message for that upgrade, especially if all process IDs are the same for all the errors and INFO messages.

9. Exit vwtool when you get a message that the procedure is complete.
10. Do a backup of the Process Engine database. This backup can serve as a checkpoint, should an error occur later that requires a database restore.
11. Restart the Process Engine software. At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

To update the Process Engine security

Before updating the Process Engine security, ensure the following:

- Your directory server is running and correctly configured.
- Content Engine 4.0 or 4.5 is running.

NOTE When the Process Task Manager starts, a message will be presented indicating that routers must be migrated. Routers will be migrated in "To migrate routers and update isolated regions" on page 320 as a part of completing Process Engine configuration changes. The message can be ignored now.

1. Update the information on the Process Task Manager Security / General tab, as follows:

   a. Start Process Task Manager as follows, depending on your operating system:

      Windows

      Select **Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager**.

      UNIX

      Enter the following command on the command line:

      ```
      vwtaskman
      ```

   b. Verify that Process Engine is running. To start it, right-click your Process Engine server in the feature pane and choose Start from the Action menu.

   c. Select the Process Engine in the feature pane and select the Security tab to configure the General settings.

      Provide the *pe_service_user* and password, the *pe_admin_group* and the optional *pe_config_group*. See the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Configure the Process Engine > Security for details on the user and groups.

NOTE The service user name should be entered as a short name, not a distinguished name.

d. Click **Apply.**

NOTE If you get an error applying security settings, click **Close** on the message, correct the problem if noted, and repeat Step 1 on page 318.

Additional information is available in:

Windows

\fnsw_loc\logs\TM_daemon\PEDirectoryServerConnectionDebug.txt

UNIX

/fnsw/local/logs/TM_daemon/PEDirectoryServerConnectionDebug.txt

e. Click **OK** to close the dialog box indicating that you must run vwtool.

f. Exit Process Task Manager.

2. Run vwtool and choose **Yes** when a message is presented indicating that an upgrade is required.

This step moves all existing user environment records from the 3.5.x format to the 4.5 format.

a. Reply No when prompted as to whether you want to initiate tracking to capture the changes make to a trace file.

CAUTION Always choose No in response to the question to override until you have carefully evaluated all users who's environment records did not migrate properly. If you are certain that all unmigrated users are no longer valid Process Engine users, that is, they have no Process Engine work, then you can override these errors and complete this part of the upgrade. Once you choose to override these errors, there is no way to recover the user environment records for any users not migrated to 4.5. Any Process Engine work for unmigrated users is lost. Choose Yes to ignore errors and force the completion of the upgrade only after you have resolved any outstanding problems.

As user environment records are moved, information is logged to the Image Services error log. When vwtool finishes, it will display messages on the screen indicating whether or not the migration was successful, how many users were migrated, and how many users were not migrated.

If all users did not successfully migrate to the 4.5 format, you will need to look at the messages in the error log and resolve the problems. Examples of the types of resolution required could include the need to fix a problem with the configuration of the Content Engine and its application server's access to the directory server or the need to create users in the underlying directory server itself.

b. Address any errors that occurred execute vwtool again. You might be prompted to ignore issues related to the user environment record upgrades.

NOTE If vwtool fails with a shared memory error, follow Step 3 through Step 6 in "(Windows only) To configure contiguous free memory for Process Engine" on page 327 to set a hardcoded shared

memory address, setting the address to 0x122300000. After setting and verifying the address, execute Step 2 again.

3. Check the error log to verify that the database version number has been updated to 52. This update will happen only after either successful migration of all environment records or all errors have been intentionally overridden.

4. At a Windows command prompt, or UNIX command line, type the following to restart the Process Engine software:

```
initfnsw -y restart
```

To migrate routers and update isolated regions

Use the following required sub-procedures to convert all routers to connection points, and assign passwords to any existing isolated regions:

To remove existing routers in Process Task Manager

1. On the Process Engine and Application Engine servers, use Process Task Manager to view, and make note of, the general properties of each Process Router:

   - Process Router name
   - Process Engine
   - Isolated region

2. Delete each Process Router, as follows.

   NOTE You can do this immediately or wait until the corresponding connection points are created. The Process Routers displayed in Process Task Manager are not used and have no effect on the system other than to cause a warning message on Process Task Manager startup. Once all Process Routers have been deleted and Process Task Manager has been restarted, the Process Router node no longer appears.

   a. Select the Process Router you want to delete.

   b. Select Delete from the Action menu.

   c. Repeat Step a and Step b until you have deleted all Process Routers.

To configure new regions in Enterprise Manager

1. Start Enterprise Manager 4.5.

2. Navigate to the PE Region IDs node and start the wizard.

3. Specify an isolated region for each unique Process Engine / isolated region combination. (If you have more than one Process Router pointing to the same Process Engine / isolated region combination, you will identify only one Process Engine Region ID.)

For this Process Engine Region ID property...	Use the value from this Process Router property...
DNS name or IP address of the Process Engine	Process Engine
Region number	Isolated region

4. Assign a password for each region as you create it. Make note of the password you assign. You will need to enter that password in the steps that follow for assigning the password to regions in Process Task Manager. The passwords must match.
5. Navigate to the PE Connection Points node, start the wizard, and create new connection points for each region.
6. Close Enterprise Manager.

To configure new regions in Process Task Manager

1. Start Process Task Manager on the Process Engine server as follows, depending on your operating system:

   Windows

   Select **Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager.**

   UNIX

   Enter the followng command on the command line:

   ```
   vwtaskman
   ```

   The terminal must support X Windows and the DISPLAY environment variable must be set.
2. Select Process Engine in the feature pane.
3. Right-click on the **Regions** folder, select New to create a new region.
4. Select the **Security Settings** sub-tab to set a region password.

   NOTE The password must match the password that you entered when creating a Process Engine Region in Step 4 on page 321.

   After you have entered all parameters, click **Apply** and restart the Process Service when prompted. If errors are returned, additional information is available in the following file:

   Windows

   \fnsw_loc\logs\TM_daemon\PEDirectoryServerConnectionDebug.txt

   UNIX

   /fnsw/local/logs/TM_daemon/PEDirectoryServerConnectionDebug.txt

To update email notification

If you're using email notification, do the following to enable email notification:

1. Add a language pack for the Default Authoring Locale.

   See the IBM FileNet P8 help topic FileNet P8 Administration > Process Engine Administration > Workflow administration tasks > Coordinating workflow design > Enable Email notification for information on adding a language pack for the Default Authoring Locale.

   NOTE Process Engine supports localized email notification. For details on configuring this, see the *IBM FileNet P8 Non-English Support Guide*. To download this document from the IBM FileNet support Web site, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

2. Verify that the Default Authoring Locale is correct (it has defaulted to the operating system's locale).

To update all isolated regions

Use the following procedure to update isolated regions, namely to perform a transfer of the `upgrade.cdl` file.

1. Restart the Process Engine software by typing the following at a command line:

   ```
   initfnsw -y restart
   ```

2. Change directories to the following location of the upgrade.cdl file, depending on your operating system:

   Windows

   \fnsw_loc\sd

   UNIX

   /fnsw/local/sd

3. Initiate a transfer on all working isolated regions by entering the following command:

   `vwtfer -o upgrade.cdl -r` All -Y

   where:

   -r All indicates all regions

   At the prompt, log on as a user who is a member of the PEAdministrators group.

(SQL Server and DB2 only) To run optional post-upgrade scripts

If your Process Engine database is either SQL Server and DB2, you can use the following sub-procedures to run two optional post-upgrade scripts. For each database type, the first script copies records from archived event-log tables to the new 4.5 version of the tables, and the second script deletes the archived log tables, thereby saving database space.

Because no such scripts are available for Oracle databases, skip to the procedure "To back up the database and restart software" on page 324, which applies to all database types.

(SQL Server)To run optional post-upgrade scripts

1. Edit and run the \fnsw\mssql\vwmssql35to40_post1.bat file, as follows:
   a. Save the vwmssql35to40_post1.bat file to the same directory as vwmssql35to40_post1a.bat.
   b. Change the values in the vwmssql35to40_post1.sql file as appropriate for your system. The contents of the file look like this:

      ```
      osql /U sa /P /n /d VWdb /h-1 /i vwmssql35to40_post1.sql
      ```

      Change the values for your system to:

      ```
      osql /D DSN /U σα /P sa /n /d VWdb /h-1 /i vwmssql35to40_post1.sql /o post1a.log
      ```

      where:

      `/D` indicates the following variable is your ODBC data source name (DSN).

      `/U` indicates the following variable is the administrator user name in the Process Engine database.

      `/P` indicates the following variable is the administrator user's password in the Process Engine database.

      `/d` indicates the following variable is the Process Engine database name.

      `/o` indicates the following variable is the optional post1a.log output file. If you choose to eliminate this entry, output displays to the screen only.
   c. Run the vwmssql35to40_post1a.bat file.
2. Delete archived event log tables by editing and running vwmmql35to40_post2.bat.
   a. Change the values in the vwmssql35to40_post2a.sql file as appropriate for your system. The contents of the file looks like this:

      ```
      osql /U sa /P /n /d VWdb /h-1 /i vwmssql35to40_post2.sql
      ```

      Change the values for your system to:

      ```
      osql /D DSN /U sa /P sa /n /d VWdb /h-1 /i vwmssql35to40_post2.sql /o post2a.log
      ```

      where:

      `/D` indicates the following variable is your ODBC data source name (DSN).

      `/U` indicates the following variable is the administrator user name in the Process Engine database.

      `/P` indicates the following variable is the administrator user's password in the Process Engine database.

      `/d` indicates the following variable is the Process Engine database name.

      /o indicates the following variable is the optional post2a.log output file. If you choose to eliminate this entry, output displays to the screen only.

b. Run the vwmssql35to40_post2a.bat file.

c. Proceed to the procedure "To back up the database and restart software" on page 324.

(DB2) To run optional post-upgrade scripts

1. Connect to the DB2 database and run the \fnsw\DB2\vwdb2_35to40_post1.sql file.

   a. Start the DB2 command line processor and log on to the Process Engine database with the Process Engine runtime user (f_sw).

   b. Enter the following command in that window:

   `db2 connect to` *`database_name`* `user` *`PE runtime user`* `using` *`password`*

   where:

   `database_name` is your Process Engine DB2 database name.

   PE runtime user is the Process Engine runtime user (f_sw).

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

   c. Run the vwdb2_35to40_post1.sql file by executing the following in the DB2 command line processor:

   ```
   db2 -tvf vwdb2_35to40_post1.sql
   ```

2. Connect to the DB2 database and run the \fnsw\DB2\vwmdb2_35to40_post2.sql file, as follows:

   a. Start the DB2 command-line processor and log on to the Process Engine database with the Process Engine runtime user (f_sw),

   b. Enter the following command in that window:

   `db2 connect to` *`database_name`* `user` *`PE runtime user`* `using` *`password`*

   where:

   `database_name` is your Process Engine DB2 database name.

   PE runtime user is the Process Engine runtime user (f_sw).

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

   c. Run the vwmdb2_35to40_post2.sql file by executing the following in the DB2 command-line processor:

   ```
   db2 -tvf vwdb2_35to40_post2.sql
   ```

   d. Proceed to the procedure "To back up the database and restart software" on page 324.

To back up the database and restart software

1. Back up the Process Engine database. While this backup is not required, it is a best practice. It provides a checkpoint that can be used later if a restore is needed in the context of this upgrade.

2. (Windows only) Start the following services and set them back to automatic startup.
   - IMS ControlService
   - Process Engine Services Manager
3. (Windows only) Disable the redirection of log messages to the Windows Event Log by changing the LogToFiles value from 1 to 0 for the following registry key:

   HKEY_LOCAL_MACHINE > SOFTWARE > FileNET > IMS > CurrentVersion
4. Restart the Process Engine software, as follows:

   At a Windows command prompt, or UNIX command line, type the following:

   ```
   initfnsw -y restart
   ```

(Windows only) To configure contiguous free memory for Process Engine

Use the following procedure to configure the largest available contiguous free memory block. If you fail to perform this procedure, the system will not allocate shared memory at some point during normal execution and will cease to function correctly.

If you have already created the registry key, reset the value based on the information acquired here.

1. Start vwtool at a command prompt, and log on using the *pe_service_user*.
2. Use the processmap command to find the largest contiguous free memory area, as in:

   ```
   <vwtool:1>processmap
   ```

   This command returns the following:

   ```
   Process Id (CR=this vwtool process):
   ```

   Press **Return** (CR) to get the process map for this process, as in the following example, where the process ID is 2592:

   ```
   C:\FNSW\BIN\vwtool.exe (ID:2592)

   Address    Attrib         Size    Owner
   =======    ======         ====     =====

   00000000   Free          65536
   00010000   Private       12288
   00013000   Free          53248
   00020000   Private         4096
   ```

   (pages of memory addresses omitted here)

   ```
   7FFDE000   Private        4096
   7FFDF000   Private        4096
   7FFE0000   Private        65536

   C:\FNSW\BIN\vwtool.exe (ID:2592)

   Largest FREE block found: 453873664 bytes at address 0x4B577000
   Rounded up to a 64K boundary, free block address 0x4B580000
   ```

In this example, `0x4B580000` is the address you will need in the next step. In some cases you might see only the line referencing the largest free block because the value is already at a 64K boundary.

3. Start regedit from the Windows **Start > Run** command and perform the following steps to create a DWORD value for IS StartShmAddress, using the address noted in step 2 as follows:

   a. Navigate to the following regedit key:

   `HKEY_Local_Machine\Software\FileNet\IMS\CurrentVersion\`

   b. Create a new DWORD value named:

   `StartShmAddress`

   c. Enter or verify the following in the Edit DWORD Value Screen:

   Value name = `StartShmAddress`

   Value data = *address of largest free memory block*

   From the example above the value will be 4B580000.

   Base is hexadecimal.

   d. Click **OK.**

   e. Exit from regedit.

4. Restart the Process Engine software.

5. Verify the setting you just applied for the shared memory address by executing the following at a command prompt:

```
ipc_tool -A
```

The following is an example of the information that is returned.

```
Image Services software shared memory segment limit: 129 segments
Current configured segment size: 0x01000000 bytes (16 MB)
Before allocating shared memory for Image Services, the SysV library
performs a test to determine the system shared memory limit.  This test
can be used as a reference for performance tuning.  The test results vary
depending on the amount of memory in use by other processes.  The actual
amount of shared memory available during operation may be less.  The test
results are:

   Successfully attached to 27 segments
   Successfully obtained 432 MB of shared memory

The following table displays the number of shared memory segments currently
in use by Image Services.  Segment #0 (called the address manager) is small.
The other segment(s) contain the actual Image Services data.  Note that
running ipc_tool will force the creation of segments #0 and #1 even when no
other Image Services process is up.
Shared Memory Address Manager Information

    Address     Shm id          Creator

Enter <space> to continue, 'q' to quit:
```

```
  0 0x4b580000  FNSHM_464d0000  Shared address manager
  1 0x4c580000  FNSHM_464a0000  FileNet server software
Total Image Services shared memory allocated: 16 MB
(This does not include segment #0)
```

NOTE The First shared memory address above 0x4B580000 is the value you would check for this example.

6. Exit ipc_tool. If the shared memory address is correct, proceed to the next installation task. If the value is not correct, verify Step 1 through Step 4 above before proceeding.

Complete the upgrade from Process Engine 4.0.2-001, 4.0.3, or higher

Perform the following steps to update a Process Engine database to version 4.5. A database update will take place automatically by restarting the software.

To update the database to version 4.5

1. Run the following command to restart the Process Engine:

   ```
   initfnsw -y restart
   ```

2. Check the Image Services error logs to verify that messages similar to the following are captured:

   2008/04/17 16:23:43.261 *fnsw* VW/Process (14952) ... [INFO]

   VW: Database upgrade successful to version 52 please follow instructions to perform the next step

   2008/04/17 16:23:43.303 *fnsw* VW/Process (14952) ... [INFO]

   VW: Must restart software to complete upgrade procedure

   Ignore messages designated as SERIOUS if they are in combination with a message indicating a successful upgrade, especially if all process IDs are the same for all errors and INFO messages.

3. Proceed to "(Windows only) To configure contiguous free memory for Process Engine" on page 327.

(Windows only) To configure contiguous free memory for Process Engine

Use the following procedure to configure the largest available contiguous free memory block. If you fail to perform this procedure, the system will not allocate shared memory at some point during normal execution and will cease to function correctly.

If you have already created the registry key, reset the value based on the information acquired here.

1. Start vwtool at a command prompt, and log on using the *pe_service_user*.

2. Use the processmap command to find the largest contiguous free memory area, as in:

   ```
   <vwtool:1>processmap
   ```

This command returns the following:

```
Process Id (CR=this vwtool process):
```

Press **Return** (CR) to get the process map for this process, as in the following example, where the process ID is 2592:

```
C:\FNSW\BIN\vwtool.exe (ID:2592)

Address    Attrib         Size    Owner
=======    ======         ====     =====

00000000   Free          65536
00010000   Private       12288
00013000   Free          53248
00020000    Private          4096
```

(pages of memory addresses omitted here)

```
7FFDE000   Private        4096
7FFDF000   Private        4096
7FFE0000    Private           65536

C:\FNSW\BIN\vwtool.exe (ID:2592)

Largest FREE block found: 453873664 bytes at address 0x4B577000
Rounded up to a 64K boundary, free block address 0x4B580000
```

In this example, `0x4B580000` is the address you will need in the next step. In some cases you might see only the line referencing the largest free block because the value is already at the 64 Kb boundary.

3. Start regedit from the Windows **Start > Run** command and perform the following steps to create a DWORD value for IS StartShmAddress, using the address noted in step 2 as follows:

   a. Navigate to the following regedit key:

      `HKEY_Local_Machine\Software\FileNet\IMS\CurrentVersion\`

   b. Create a new DWORD value named:.

      `StartShmAddress`

   c. Enter or verify the following in the Edit DWORD Value Screen:

      Value name = `StartShmAddress`

      Value data = *address of largest free memory block*

      From the example above the value will be 4B580000.

      Base is hexadecimal.

   d. Click **OK.**

   e. Exit from regedit.

4. Restart the Process Engine software.

5. Verify the setting you just applied for the shared memory address by executing the following at a command prompt:

```
ipc_tool -A
```

The following is an example of the information that is returned.

```
Image Services software shared memory segment limit: 129 segments
Current configured segment size: 0x01000000 bytes (16 MB)
Before allocating shared memory for Image Services, the SysV library
performs a test to determine the system shared memory limit.  This test
can be used as a reference for performance tuning.  The test results vary
depending on the amount of memory in use by other processes.  The actual
amount of shared memory available during operation may be less.  The test
results are:

   Successfully attached to 27 segments
   Successfully obtained 432 MB of shared memory

The following table displays the number of shared memory segments currently
in use by Image Services.  Segment #0 (called the address manager) is small.
The other segment(s) contain the actual Image Services data.  Note that
running ipc_tool will force the creation of segments #0 and #1 even when no
other Image Services process is up.
Shared Memory Address Manager Information

    Address     Shm id          Creator

Enter <space> to continue, 'q' to quit:

  0 0x4b580000  FNSHM_464d0000  Shared address manager

  1 0x4c580000  FNSHM_464a0000  FileNet server software

Total Image Services shared memory allocated: 16 MB

(This does not include segment #0)
```

NOTE The First shared memory address above 0x4B580000 is the value you would check for this example.

6. Exit ipc_tool. If the shared memory address is correct, proceed to the next installation task. If the value is not correct, verify Step 1 through Step 4 above before proceeding.

To update all isolated regions

Use the following procedure to update all isolated regions, namely to perform a transfer of the `upgrade.cdl` file.

1. Restart the Process Engine software by typing the following at a command line:

   ```
   initfnsw -y restart
   ```

2. Change directories to the following location of the upgrade.cdl file, depending on your operating system:

   Windows

   \fnsw_loc\sd

   UNIX

   /fnsw/local/sd

3. Initiate a transfer on all working isolated regions by entering the following command:

   `vwtfer -o upgrade.cdl -r` All -Y

   where -r All indicates all regions

   At the prompt, log on as a user who is a member of the PEAdministrators group.

Upgrade and configure Application Engine

The tasks for upgrading Application Engine vary depending on whether you start from version 3.5.x or version 4.0.x. You must perform some tasks regardless of your upgrade starting point, as indicated in the following steps.

For upgrades from 3.5.x

1. Prepare for the Application Engine upgrade. Perform Task 1 on page 332.
2. Upgrade Application Engine. Perform Task 2a on page 336.
3. Install Application Engine software updates. Perform Task 3 on page 345.
4. Install the Content Engine Client updates. Perform Task 4 on page 346.
5. Install the Process Engine Client updates. Perform Task 5 on page 350.
6. Configure Application Engine according to your application server type:
   - Configure for Websphere: Task 6a on page 353
   - Configure for WebLogic:Task 6b on page 363
   - Configure for JBoss: Task 6c on page 369
7. Manually copy custom Application Engine data. Perform Task 7 on page 372.
8. Deploy Application Engine according to your application server type:
   - Deploy on Websphere: Task 8a on page 373
   - Deploy on WebLogic: Task 8b on page 375
   - Deploy on JBoss: Task 8c on page 377
9. Complete post-upgrade Application Engine configuration. Perform Task 9 on page 379.

For upgrades from 4.0x

1. Prepare for the Application Engine upgrade. Perform Task 1 on page 332.
2. Upgrade Application Engine. Perform Task 2b on page 341.
3. Install Application Engine software updates. Perform Task 3 on page 345.
4. Install the Content Engine Client updates. Perform Task 4 on page 346.
5. Install the Process Engine Client updates. Perform Task 5 on page 350.
6. Manually copy custom Application Engine data. Perform Task 7 on page 372.
7. Deploy Application Engine according to your application server type:
   - Deploy on Websphere: Task 8a on page 373
   - Deploy on WebLogic: Task 8b on page 375
   - Deploy on JBoss: Task 8c on page 377
8. Complete post-upgrade Application Engine configuration. Perform Task 9 on page 379.

Task 1: Prepare for Application Engine upgrade

This task includes Application Engine upgrade preparation instructions for WebSphere, WebLogic, and JBoss (UNIX and Windows). It applies to Application Engine upgrades from both 3.5.*x* and 4.0.*x* versions.

Before you upgrade Application Engine

- (Upgrades from 3.5.x only) Verify that Content Engine and Process Engine have been upgraded.
- Review the Application Engine details of the guide *Plan and Prepare Your Environment for IBM FileNet P8 Platform*.
- Review the steps needed to retain the AE 3.5.x or 4.0.1 configuration. You will need this information if, for any reason, you want to back out of this installation.
- Verify that you have recorded *all necessary settings*.

  (Best practices for backup) If you want to retain your existing Application Engine settings you must record all necessary settings before you start the upgrade installation.

  – As part of the upgrade you will create backup copies of all important 3.5.x or 4.0.1 configuration files in "Backup, undeploy, and remove the Workplace web application from the J2EE application server." on page 333.

  – During the upgrade, a number of existing configuration files will be moved to the newly created *AE_install_path*\Config\AE directory. See Table 1, "Configuration files that will be moved," on page 335.

  – As part of the upgrade, the installation program automatically creates a backup of your existing AE 3.5.x or 4.0.1 configuration files, appending the suffix .old to the file names. See Table 2, "Configuration files that will be backed up," on page 335.

  – The existing version 3.5.x or 4.0.1 Actions.xml and web.xml files will be merged with the 4.0.2 versions during the upgrade installation.

  – In addition, all comments added to the Actions.xml file will be lost during the merge.
- If you have the IBM FileNet P8 eForms expansion product installed, uninstall it.

  See the "Removing Software" topic in the *IBM FileNet P8 eForms Installation Guide*. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

To prepare for the Application Engine upgrade

1. Log on to the Application Engine server:

   UNIX

   Log on as root.

Windows

Log on as a member of the local Administrators group or a user with equivalent permissions.

2. (Upgrades from 3.5.x only) Verify that the Process Router is stopped.

   NOTE Although the Application Engine router has been deleted as part of the Process Engine upgrade, a local, disconnected instance might be running on your server.

   a. Launch the Process Task Manager from *AE_install_path*/Router:

      UNIX

      ```
      ./routercmd.sh
      ```

      Windows

      ```
      routercmd.bat
      ```

   b. Stop the router, if running.

      NOTE This is the router configured and started as part of the FileNet P8 Platform 3.5.x installation. For more information, see the *FileNet P8 Platform 3.5.x Installation and Upgrade Guide* task "Start the Process Router." To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

   c. Exit the Process Task Manager.

3. (Windows only) Using the Windows Task Manager, verify that no javaw.exe processes are running. If the applications in the preceding steps stopped correctly, no Application Engine related javaw.exe processes should be running.

   a. Use Windows Services to stop the Process Application Engine Services Manager Service. Select **Start > Programs > Administrative Tools > Services**. If the Process Application Engine Services Manager status is started, then right-click the item and select **Stop**.

   b. If there are still javaw.exe processes running after stopping the Process Application Engine Services Manager, use the Windows Task Manager to stop any Application Engine related javaw.exe processes.

4. Backup, undeploy, and remove the Workplace web application from the J2EE application server.

   During this step you will create a backup copy of the deployed web application containing all customized files and all configuration files stored in the WEB-INF directory.

   NOTE Even though the installer automatically creates a backup of your existing AE 3.5.x or 4.0.1 configuration files IBM recommends backing up the directories below to get a complete backup of your deployed system. You might need to use these files in later steps, such as when you copy modified files to the installed Workplace directory and perform the upgrade and configuration of Application Engine, to retain your 3.5.x or 4.0.1 settings.

   WebSphere

   a. Make a backup copy of the deployed Workplace directory (*deploy_backup*):

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

      NOTE Your path may vary if you created a new profile or renamed your WAR file.

b. Stop the Workplace application from the WebSphere administrative console.

c. Uninstall the Workplace application from **Enterprise Applications**.

d. Save the changes and stop the WebSphere server.

e. Delete the temp Workplace directory (default: app_engine_war) from:

WAS_HOME/profiles/default/temp/*node_name/instance*/app_engine.war

WebLogic

a. From the WebLogic Administration Console, stop the Workplace web application module.

b. Make a backup copy of the Workplace folder (*deploy_backup*):

AE_install_path/Workplace

c. Delete the Workplace Web Application Module.

d. Stop the WebLogic server.

JBoss

a. Stop the JBoss server.

b. Make a backup copy of the deployed Workplace directory (*deploy_backup*):

JBoss_HOME/server/default/deploy/Workplace.war

c. Delete the deployed Workplace directory:

JBoss_HOME/server/default/deploy/Workplace.war

d. Delete the temporary Workplace directory for JBoss:

JBoss_HOME\server\default\work\jboss.web\localhost\Workplace

5. Copy modified files to the installed Workplace directory.

During the upgrade, the configuration files listed in Table 1 on page 335 will be moved from the following directory:

AE_install_path/Workplace/WEB-INF

to

AE_install_path/Config/AE

NOTE If you have made modifications to any of these files directly in the deployed Workplace directory (app_engine.war for WebSphere or Workplace.war for JBoss) in your Filenet P8 3.5.x or 4.0.1 environment you must copy modified version of these files from the *deploy_backup*

directory to the installed directory*AE_install_path*/Workplace/WEB-INF before you run the installer. :

Table 1: Configuration files that will be moved

- actions.xml
- bootstrap.properties
- ClassFilter.xml
- ConfigurableLabels.xml
- containericons.properties
- content_redir.properties
- customobjecticons.properties
- download_redir.properties
- fnsoap.xml
- icons.properties
- InfoPages.xml
- PagingConfiguration.xml
- PolicyProcessors.xml
- PrimaryViews.xml
- PropertiesPages.xml
- SimpleSearch.xml
- SystemsPropertiesView.xml

NOTE As part of the upgrade the installation program automatically creates a backup of your existing AE 3.5.x or 4.0.1 configuration files, appending the suffix .old to the file names.

Table 2: Configuration files that will be backed up

- actions.xml.old
- ClassFilter.xml.old
- ConfigurableLabels.xml.old
- containericons.properties.old
- content_redir.properties.old
- icons.properties.old
- SystemsPropertiesView.xml.old
- UpdateActions.xml.old
- UpdateClassFilter.xml.old
- UpdateConfigurableLabels.xml.old
- UpdateProps.xml.old
- UpdateSystemPropertiesView.xml.old
- UpdateWeb.xml.old
- web.xml.old

The backed up files are located in:

AE_install_path/backup-4_0_2

6. Upgrade and configure the Application Engine software. See "Upgrade Application Engine from 3.5.x" on page 336.

Task 2a: Upgrade Application Engine from 3.5.*x*

This topic includes Application Engine installation instructions for all supported application servers, for UNIX and Windows platforms.

NOTES

- If you plan to install and use the IBM FileNet Workplace XT product, installing Application Engine is not required.
- Before installing Application Engine, check the latest version of the IBM *FileNet P8 4.5 Release Notes* for known issues that might impact this software installation. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- Make sure your installation location meets the requirements specific for Application Engine outlined 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.
- If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch.

  Otherwise, if you deploy Application Engine initially as a WAR file and later decide to redeploy as an EAR file, you will have to uninstall Application Engine and then reinstall the application, selecting EAR file deployment, to add the required files to your setup.
- (Highly Available installations) To install Application Engine in a web farm or clustered environment, follow the instructions in the*IBM FileNet P8 Platform High Availability Technical Notice.* The document outlines the required HA install procedure and references this guide for detailed installation and deployment instructions. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- To ensure proper functionality and performance only install one instance of Application Engine per application server (or virtual machine or WebSphere LPAR). You can, however, deploy multiple instances of a single Application Engine version per application server. For details, see "Deploy multiple Application Engine instances" on page 181.
- Before logging on to Workplace for the first time, at least one object store must exist on the Content Engine to hold the site preferences. See "Create the initial object store" on page 60 for more information.
- If you run the installer to upgrade Application Engine, the installer verifies that the currently installed version of Application Engine can be upgraded. See "Prepare for Application Engine upgrade" on page 332 for more information.

To install the Application Engine software

1. Log on to the application server, as appropriate for your operating system:

   UNIX

   Log on as a user with write access to the /bin directory and read, write, and execute access to the directory where you plan to install Application Engine.

   NOTE Although the installing user must have write access to the /bin directory, the Application Engine installer does not write to that directory.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

2. Start the installation process.

   - To install Application Engine interactively:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software.

     ii. Launch the appropriate Setup program (P8AE-4.0.2.0-*operating_system*.bin/.exe) and continue with Step 3 on page 338 below.

   - To install Application Engine silently:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software package, and copy the appropriate AE_silent_input.txt file to a local directory.

     ii. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your installation.

     CAUTION If you are modifying the silent input file to perform an upgrade from AE 3.5.x to AE 4.0.2 you must modify all instances of *AE_install_path* in the script as follows:

     UNIX

     Change `../FileNet/AE` to `../FileNet`

     Windows

     Change `..\FileNet\AE` to `..\FileNet`

     Change `..\\Filenet\\AE` to `..\\FileNet`

     iii. From a command prompt, navigate to, and execute the installer:

     UNIX

     ```
     ./P8AE-4.0.2.0-operating system.bin -options path_to_edited_input_file/
     AE_silent_input.txt -silent
     ```

     Windows

     ```
     P8AE-4.0.2.0-Win.exe -options path_to_edited_input_file\AE_silent_input.txt -
     silent
     ```

NOTE When installing Application Engine remotely on UNIX, run the installer with the standard input stream redirected from `/dev/console`, for example:

```
rsh remote-machine -n P8AE-4.0.2.0-AIX.bin -options AESilentScriptUNIX.txt -
silent < /dev/console
```

If you do not add the redirect, the silent intaller will fail with a "process not attached to terminal" message and the usage message for the "who" command.

iv. Continue with Step 4.

3. Complete the Application Engine Installer screens. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

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 the Application Engine installation program:

- Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select AE installer.
- Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

In addition to the prompts for system-specific information, you will need to consider the following decisions:

- If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the /FileNet/AE part of the path. This affects the following additional paths:
  - Configuration path: Retain the /FileNet/Config/AE part of the path.

    The configuration files for an EAR file deployment, a web farm, or a clustered environment must be located in a shared folder that is accessible by all copies of the Workplace application. For more information, see the *IBM FileNet P8 Platform High Availability Technical Notice*. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
  - Log files: Retain the /FileNet/Logs part of the path.
  - User tokens: Retain the /FileNet/Authentication part of the path.
- You can choose to deploy Application as a WAR file or an EAR file. If you decide to deploy Application Engine as a WAR file and later decide to redeploy as an EAR file you must uninstall Application Engine and then reinstall the application, selecting EAR file deployment. For information, see "If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch." on page 336.
- You can choose whether to use container-managed or application-managed authentication. Choose container-managed if you plan to use SSO.

When you select Container-Managed Authentication, Setup 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 Setup is finished.

- Additional installation notes:
  - To change the Content Engine name later, or to connect to a different Content Engine, edit the WcmApiConfig.properties file. For information, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs.
  - For information on how to reconfigure the Documentation URL after installation is completed, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > Bootstrap properties.
  - The UTCryptoKeyFile.properties file contains the user token cryptography key used by IBM FileNet P8 applications to launch into each other without the need for additional login.

    CAUTION For multiple applications to pass user tokens to each other, each participating application **must** use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all servers that are hosting a token-sharing application.

    For information, see the IBM FileNet P8 Developer Help topic Developer Help > Workplace Integration and Customization Introduction > User Tokens > Configuring Applications to Use Tokens.

4. View the app_engine_install_log_4_0_2.txt file located, located in *AE_install_path*/Logs.

   Verify that no errors or failures were logged. Correct any errors before you proceed.

5. (Solaris only) Set the anon ports.

   To use the IBM FileNet ports listed below for Component Manager on Solaris-based systems, you must first enable the ports by setting the smallest anon port to 32778. When you do this, the ports used by Solaris communication deamons will be 32778 or greater, leaving port 32777 available for IBM FileNet use.

   When Solaris first starts up, it takes the first several ports, called anonports, for its communication daemons. By default, the maximum tcp_smallest_anon_port is 32768. IBM FileNet uses several ports higher than 32768. See the appendix “IBM FileNet P8 ports” in the *Plan and Prepare Your Environment for IBM FileNet P8* for details on which ports IBM FileNet uses.

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

   a. Determine the current tcp_smallest_anon_port setting.

      From a command prompt, enter the following:

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

   b. Enable port 32777.

      If the port returned in the step above is less than 32778, you must enable port 32777.

Solaris 9

Edit the /etc/rc2.d/S69inet file.

Add the following line before the exit 0 entry at the bottom of the file:

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

Solaris 10

Edit the /lib/svc/method/net-init file.

Add the following line before the exit 0 entry at the bottom of the file:

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

c. Reboot the Application Engine server.

You must reboot the Application Engine server to force the release of ports required by the Application Engine that may be in use by the operating system.

CAUTION Failure to reboot after these changes are made can result in the port being unavailable, generating OpenSocket errors.

6. (Optional, UNIX only) Set the the P8TASKMAN_HOME system environment variable.

The Process Engine Client installer will set this variable for you. If you want to specify a different location for the client files, you can do so by setting the P8TASKMAN_HOME system variable. The default value set by the PE Client installer is:

P8TASKMAN_HOME=*AE_install_path*/CommonFiles

7. Continue with "Install Application Engine software updates" on page 345.

Task 2b: Upgrade Application Engine from 4.0.*x*

This topic includes Application Engine installation instructions for all supported application servers, for UNIX and Windows platforms.

NOTES

- If you plan to install and use the IBM FileNet Workplace XT product, installing Application Engine is not required.
- Before installing Application Engine, check the latest version of the IBM *FileNet P8 4.5 Release Notes* for known issues that might impact this software installation. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- Make sure your installation location meets the requirements specific for Application Engine outlined 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.
- The installer creates the folder structure and files needed for Application Engine.
- If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch.
- Otherwise, if you deploy Application Engine initially as a WAR file and later decide to redeploy as an EAR file, you will have to uninstall Application Engine and then reinstall the application, selecting EAR file deployment, to add the required files to your setup.
- (Highly Available installations) To install Application Engine in a web farm or clustered environment, follow the instructions in the*IBM FileNet P8 Platform High Availability Technical Notice.* The document outlines the required HA install procedure and references this guide for detailed installation and deployment instructions. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.
- To ensure proper functionality and performance, only install one instance of Application Engine per application server (or virtual machine or WebSphere LPAR). You can, however, deploy multiple instances of a single Application Engine version per application server, see "Deploy multiple Application Engine instances" on page 181.
- Before logging on to Workplace for the first time, at least one object store must exist on the Content Engine to hold the site preferences. See "Create the initial object store" on page 60 for more information.
- If you run the installer to upgrade Application Engine, the installer verifies that the currently installed version of Application Engine can be upgraded. See "Prepare for Application Engine upgrade" on page 332 for more information.

To install the Application Engine software

1. Log on to the application server:

   UNIX

   Log on as a user with write access to the /bin directory and read, write, and execute access to the directory where you plan to install Application Engine.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

   NOTE Although the installing user must have write access to the /bin directory, the Application Engine installer does not write to that directory.

2. Start the installation process.

   - To install Application Engine interactively:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software.

     ii. Launch the appropriate Setup program (P8AE-4.0.2.0-*operating_system*.bin/.exe) and continue with Step 3 on page 343 below.

   - To install Application Engine silently:

     i. Access the IBM FileNet Application Engine 4.0.2 installation software package, and copy the AE_silent_input.txt file to a local directory.

     ii. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your installation.

     CAUTION If you are modifying the silent input file to perform an upgrade from AE 3.5.x to AE 4.0.2 you must modify all instances of *AE_install_path* in the script as follows:

     UNIX

     Change `../FileNet/AE` to `../FileNet`

     Windows

     Change `..\FileNet\AE` to `..\FileNet`

     Change `..\\Filenet\\AE` to `..\\Filenet`

     iii. From a command prompt, navigate to, and execute the installer, then continue with Step 4 “View the app_engine_install_log_4_0_2.txt file located, located in AE_install_path/Logs.” on page 343.

     UNIX

     ```
     ./P8AE-4.0.2.0-<operating system>.bin -options <path_to_edited_input_file>/
     AE_silent_input.txt -silent
     ```

     Windows

     ```
     P8AE-4.0.2.0-Win.exe -options <path_to_edited_input_file>\AE_silent_input.txt
     -silent
     ```

3. Complete the Application Engine Installer screens. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   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 the Application Engine installation program:

   - Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select AE installer.
   - Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

   In addition to the prompts for system-specific information, you will need to consider the following decisions:

   - You can choose to deploy Application as a WAR file or an EAR file. If you decide to deploy Application Engine as a WAR file and later decide to redeploy as an EAR file you must uninstall Application Engine and then reinstall the application, selecting EAR file deployment. For information, see "If a change of deployment type seems likely for your setup in the future, especially if you think you might switch from a WAR file to an EAR file deployment, it is easier if you always install Application Engine as if you intend to deploy an EAR file. This approach works best because it always creates both a WAR and an EAR file, meaning you have a ready WAR file for a switch." on page 341.
     - Additional installation notes:
     - To change the Content Engine name later, or to connect to a different Content Engine, edit the WcmApiConfig.properties file. For information, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs.
     - For information on how to reconfigure the Documentation URL after installation is completed, see the IBM FileNet P8 help topic FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > Bootstrap properties.
     - The UTCryptoKeyFile.properties file contains the user token cryptography key used by IBM FileNet P8 applications to launch into each other without the need for additional login.

       CAUTION For multiple applications to pass user tokens to each other, each participating application **must** use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all servers that are hosting a token-sharing application.

       For information, see the IBM FileNet P8 Developer Help topic Developer Help > Workplace Integration and Customization Introduction > User Tokens > Configuring Applications to Use Tokens.

4. View the app_engine_install_log_4_0_2.txt file located, located in *AE_install_path*/Logs.

   Verify that no errors or failures were logged. Correct any errors before you proceed.

5. (Optional; UNIX only) Set the the P8TASKMAN_HOME system environment variable.

   The Process Engine Client installer will set this variable for you. If you want to specify a different location for the client files, you can do so by setting the P8TASKMAN_HOME system variable. The default value set by the PE Client installer is:

   `P8TASKMAN_HOME=`*`AE_install_path`*`/CommonFiles`

6. Continue with "Install Application Engine software updates" on page 345.

Task 3: Install Application Engine software updates

Install any fix packs and interim fixes required for Application Engine.

To install the Application Engine software updates

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your support representative.
2. Open the readmes for any subsequent fix packs or interim fixes (typically optional) and perform the installation procedures provided.
3. Install the latest updates for the Content Engine Client and Process Engine Client files using the subsequent tasks.

Task 4: Install the latest Content Engine Client files on Application Engine servers

NOTE If you are performing a staged upgrade, and are upgrading Application Engine before you upgrade Content Engine, you do not need to do this task at this time. When you upgrade Content Engine to 4.5, you must then use this task to run the 4.5 Content Engine Client installer on your Application Engine, and redeploy the Workplace application.

To uninstall CE Client 4.0.x instances

Perform the following procedure to identify and uninstall all Content Engine Client version 4.0.x files on an Application Engine Server machine. Repeat the procedure for all other Application Engine Server machines.

Depending on the CE Client fixpack versions that you have installed, you may have more than one instance of the CE Client 4.0.x installed. Additionally, the version 4.0.x CE Client installer supported installing multiple instance of the CE Client on the same server. You need to uninstall all instances before upgrading to the 4.5 version of CE Client.

1. Locate and uninstall all instances of the CE Client on the Application Engine server.

   Windows

   a. Select **Start > Settings > Control Panel > Add/Remove Programs**.

   b. Select **Content Engine Client Updater** or **IBM FileNet Content Engine Client Installer**, depending on the fixpack version of Content Engine Client that is installed.

      Prior to fix pack P8CE-4.0.1-006: The Content Engine Client is named Content Engine Client Updater.

      As of fix pack P8CE-4.0.1-006: The Content Engine Client is named IBM FileNet Content Engine Client Installer.

   c. Click **Change/Remove** to start the uninstaller.

   d. In the Uninstall IBM FileNet Client Installer screen, click **Next**.

   e. In the Select Components screen, select Complete Uninstall. Click **Next**.

   f. In the Uninstall Complete screen, click **Done**.

   g. Repeat Step b through Step f for each additional instance of **Content Engine Client Updater** or **IBM FileNet Content Engine Client Installer** listed in Add/Remove Programs.

   UNIX

   a. Run one of the following commands, depending on the fixpack version of Content Engine Client that is installed.

      `/ce_client_install_path/ClientUpdater/_uninst/uninstaller.bin`

      `/ce_client_install_path/CEClient/_CEClientuninst/uninstaller`

      Prior to fix pack P8CE-4.0.1-006: The Content Engine Client files are in a directory called ClientUpdater, and the uninstallation program is .../ClientUpdater/_uninst/uninstaller.bin.

As of fix pack P8CE-4.0.1-006: The Content Engine Client files are in a directory called CEClient and the uninstallation program is .../CEClient/_CEClientuninst/uninstaller.

b. In the Uninstall IBM FileNet Client Installer screen, click **Next**.

c. In the Select Components screen, select Complete Uninstall. Click **Next**.

d. In the Uninstall Complete screen, click **Done**.

e. Repeat Step a through Step d for each installed instance of CE Client on this server.

To install the 4.5 release or fix pack version Content Engine Client files, perform the following procedures on all Application Engine servers. Except where noted, these steps apply to WebSphere, WebLogic, and JBoss.

To install the Content Engine Client files

Refer to the appropriate information from your installation worksheet while performing the following steps.

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. On the machine where Application Engine is installed, log on as any user who has the following permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Application Engine is installed
   - Execute permission on the Content Engine Client install software
3. Verify that there is a current backup of the Application Engine.
4. Copy the Content Engine Client install software from the Content Engine installation software to the temporary directory. The version of the install software must match the version of Content Engine.
5. Expand the (TAR or ZIP) Content Engine Client install software within the temporary directory.
6. Close all instances of Workplace and any other Application Engine client applications. From the application server administrative console, stop and un-deploy Application Engine.

   WebSphere

   Stop and un-deploy the FileNet Application Engine application.

   WebLogic

   Stop and un-deploy the FileNet Application Engine application.

   JBoss

   Execute the shutdown command.

7. (WebLogic only) Manually delete the following application server cache directories, substituting your domain name in place of *mydomain*:

   WebLogic UNIX

   /opt/bea/user_projects/domains/*mydomain*/servers/AdminServer/tmp/_WL_user/app_engine

   /opt/ bea/user_projects/domains/*mydomain*/servers/.wlnotdelete

   /opt/bea/user_projects/domains/*mydomain*/servers/AdminServer/cache/EJBCompilerCache

   WebLogic Windows

   C:\bea\user_projects\domains*mydomain*\servers\AdminServer\tmp_WL_user\app_engine

   C:\bea\user_projects\domains*mydomain*\servers\.wlnotdelete

   C:\bea\user_projects\domains*mydomain*\servers\AdminServer\cache\EJBCompilerCache

8. Start the Content Engine client install process.

   - To install the Content Engine client interactively:

     i. Access the IBM FileNet Content Engine client update software.

     ii. Run one of the commands in the table below, *CE_version* is the version of Content Engine you intend to run, for example, 4.5.0..

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN`

     iii. Complete the installation program wizard.

   - To install the Content Engine client files silently:

     i. Make a back up copy of the CEClient_silent_install.txt input file.

     ii. Open the silent input file in a text editor. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your update.

iii. Navigate to the path containing the Content Engine Client installation program, and run one of the commands in the following table to perform the silent install, where:

CE_version is the version of Content Engine you intend to run, for example, 4.5.0.

path is the path that contains the installation program.

Operating System	Install Program
AIX	`P8CE-CLIENT-`*`CE_version`*`-AIX.BIN -f `*`path`*`/CECLIENT.AIX/CEClient_silent_install.txt -i silent`
HPUX	`P8CE-CLIENT-`*`CE_version`*`-HPUX.BIN -f `*`path`*`/CEClient.HPUX/CEClient_silent_install.txt -i silent`
HPUXi	`P8CE-CLIENT-`*`CE_version`*`-HPUXI.BIN -f `*`path`*`/CEClient.HPUXI/CEClient_silent_install.txt -i silent`
Linux	`P8CE-CLIENT-`*`CE_version`*`-LINUX.BIN -f `*`path`*`/CEClient.Linux/CEClient_silent_install.txt -i silent`
Solaris	`P8CE-CLIENT-`*`CE_version`*`-SOL.BIN -f `*`path`*`/CEClient.Solaris/CEClient_silent_install.txt -i silent`
Windows	`P8CE-CLIENT-`*`CE_version`*`-WIN.EXE -f `*`path`*`\CEClient.Windows\CEClient_silent_install.txt -i silent`
zLinux	`P8CE-CLIENT-`*`CE_version`*`-ZLINUX.BIN -f `*`path`*`/CEClient.zLinux/CEClient_silent_install.txt -i silent`

9. Install Process Engine client updates on the Application Engine. Continue at "Install the Latest Process Engine Client files on Application Engine servers" on page 350.

NOTE If you performed this CE Client install task during a staged upgrade of Content Engine, and have previously upgrade Application Engine to 4.5, your next step is to re-deploy the Workplace application. See one of the following topics:

- "Deploy upgraded Application Engine instances (WebSphere)" on page 373
- "Deploy upgraded Application Engine instances (WebLogic)" on page 375
- "Deploy upgraded Application Engine instances (JBoss)" on page 377

Task 5: Install the Latest Process Engine Client files on Application Engine servers

NOTE If you are performing a staged upgrade from 4.0.1, and are upgrading Application Engine before you upgrade Process Engine, you do not need to do this task at this time. When you upgrade Process Engine to 4.5, you must then use this task to run the 4.5 Process Engine Client installer on your Application Engine, and redeploy the Workplace application.

To install the Process Engine Client files, perform the following steps on all Application Engine servers.

To install the Process Engine client files

1. On the machine where Application Engine is to be deployed, log on as any user who has the following permissions:
   - Read and write permission to a temporary directory, such as temp (Windows) or tmp (UNIX), on the machine where Application Engine is installed
   - Execute permission on the Process Engine Client install software
2. Copy the Process Engine Client install software from the Process Engine installation software to the temporary directory. The version of the install software must match the version of Process Engine.
3. Expand the (TAR or ZIP) Process Engine Client install software within the temporary directory.
4. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Application Engine is deployed. Run the program either interactively (using the install wizard) or silently.

   To run the program interactively, run one of the commands in the table below, follow the wizard instructions, and then continue at step 5:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN`

To run the program silently, perform the following steps:

a. In the expanded install software, open the file PEClient_silent_install.txt and edit it as follows:

   i. Change the `Variable_CheckboxAE` line to the following:

   ```
   -V Variable_CheckboxAE="true"
   ```

   ii. Run one of the commands in the following table to perform the silent install:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.0-AIX.BIN -silent -options "PEClient_silent_install.txt"`
HPUX	`P8PE-CLIENT-4.5.0-HPUX.BIN -silent -options "PEClient_silent_install.txt"`
HPUXi	`P8PE-CLIENT-4.5.0-HPUXI.BIN -silent -options "PEClient_silent_install.txt"`
Linux	`P8PE-CLIENT-4.5.0-LINUX.BIN -silent -options "PEClient_silent_install.txt"`
Solaris	`P8PE-CLIENT-4.5.0-SOL.BIN -silent -options "PEClient_silent_install.txt"`
Windows	`P8PE-CLIENT-4.5.0-WIN.EXE -silent -options "PEClient_silent_install.txt"`
zLinux	`P8PE-CLIENT-4.5.0-ZLINUX.BIN -silent -options "PEClient_silent_install.txt"`

5. (For Process Engine upgrades from PE 4.0.x to 4.5 only) Add required JAR files to the Component Manager library.

   a. Launch the Process Task Manager by running one of the following command files from the *AE_install_path*/FileNet/AE/Router directory, depending on your operating system:

      UNIX

      routercmd.sh

      Windows

      routercmd.bat

   b. In the Process Task Manager console, under Application Engine, Component Manager node, select your connection point.

   c. Stop the connection point instance.

   d. Click the Required libraries tab.

   e. Click Add, and browse to /Router/lib.

   f. Select the following files to add:

ContentExtendedOps.jar

smtp.jar

6. (For upgrades from 3.5.x only) Configure Application Engine.

   Follow the instructions for your application server to configure Application Engine:

   - "Configure Application Engine upgrades from 3.5.x (WebSphere)" on page 353
   - "Configure Application Engine upgrades from 3.5.x (WebLogic)" on page 363
   - "Configure Application Engine upgrades from 3.5.x (JBoss)" on page 369

7. (For upgrades from 4.0.x only) Continue with "Manually copy custom data" on page 372.

NOTE If you performed this PE Client install task during a staged upgrade of Process Engine, and have previously upgrade Application Engine to 4.5, your next step is to re-deploy the Workplace application. See one of the following topics:

- "Deploy upgraded Application Engine instances (WebSphere)" on page 373
- "Deploy upgraded Application Engine instances (WebLogic)" on page 375
- "Deploy upgraded Application Engine instances (JBoss)" on page 377

Task 6a: Configure Application Engine upgrades from 3.5.*x* (WebSphere)

This topic covers the configuration of your Application Engine web application (Workplace) on WebSphere. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using WebSphere with container-managed authentication or SSO, edit web.xml. See "To edit web.xml for container-managed authentication or SSO" on page 353.
2. If you are using SSO, edit web.xml. See "(SSO Only) To edit web.xml for SSO (optional)" on page 356.
3. Configure the Application Engine. See "To configure Application Engine (WebSphere)" on page 357.
4. Configure the server ports. See "To configure the server ports" on page 362.

To edit web.xml for container-managed authentication or SSO

NOTE Perform this procedure only if your site uses WebSphere with container-managed authentication or Single Sign-On (SSO). If you are using SSO, you must perform additional configuration steps as directed at the end of this procedure.

1. Make a back-up copy of web.xml.

   AE_install_path/Workplace/WEB-INF/web.xml

2. Edit web.xml.

   NOTE Text in **bold** in the examples below indicates changes made to the original web.xml file.

   a. Search for the parameter challengeProxyEnabled and set it to false.

   ```
   <param-name>challengeProxyEnabled</param-name>
   <param-value>false</param-value>
   ```

   b. Search for the first instance of <web-resource-collection>, and uncomment the url-pattern as noted in the file comments below.

   ```
   <web-resource-collection>
      <web-resource-name>action</web-resource-name>
      <description>Define the container secured resource</description>
      <url-pattern>/containerSecured/*</url-pattern>

               <!--

                   Uncomment this section if all resources that require credentials
         must be secured in order to obtain a secured Thread.  If using WebSphere,
         this section must be
                    uncommented.
         --> Move this commenting tag here from just before the </web-resource-
         collection> closing tag below.

         <url-pattern>/containerSecured/*</url-pattern>
         <url-pattern>/</url-pattern>
         <url-pattern>/author/*</url-pattern>
         <url-pattern>/Browse.jsp</url-pattern>
         <url-pattern>/eprocess/*</url-pattern>
   ```

```
<url-pattern>/Favorites.jsp</url-pattern>
<url-pattern>/GetPortalSitePreferences.jsp</url-pattern>
<url-pattern>/GetTokenSignIn.jsp</url-pattern>
<url-pattern>/GetUserInformation.jsp</url-pattern>
<url-pattern>/GetUserToken.jsp</url-pattern>
<url-pattern>/HomePage.jsp</url-pattern>
<url-pattern>/IntegrationWebBasedHelp.jsp</url-pattern>
<url-pattern>/is/*</url-pattern>
<url-pattern>/operations/*</url-pattern>
<url-pattern>/portlets/Author/edit.jsp</url-pattern>
<url-pattern>/portlets/Author/portlet.jsp</url-pattern>
<url-pattern>/portlets/Browse/edit.jsp</url-pattern>
<url-pattern>/portlets/Browse/portlet.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/edit.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/portlet.jsp</url-pattern>
<url-pattern>/portlets/GroupPageDesign.jsp</url-pattern>
<url-pattern>/portlets/GroupPageSettings.jsp</url-pattern>
<url-pattern>/portlets/Inbox/edit.jsp</url-pattern>
<url-pattern>/portlets/Inbox/portlet.jsp</url-pattern>
<url-pattern>/portlets/MultiPagesDesign.jsp</url-pattern>
<url-pattern>/portlets/OrganizePages.jsp</url-pattern>
<url-pattern>/portlets/PortalPageDesign.jsp</url-pattern>
<url-pattern>/portlets/PortalPageInfo.jsp</url-pattern>
<url-pattern>/portlets/PortletAlias.jsp</url-pattern>
<url-pattern>/portlets/PortletSettings.jsp</url-pattern>
<url-pattern>/portlets/PreviewAndSetup.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/edit.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/portlet.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/edit.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/portlet.jsp</url-pattern>
<url-pattern>/portlets/Workflows/edit.jsp</url-pattern>
<url-pattern>/portlets/Workflows/portlet.jsp</url-pattern>
<url-pattern>/properties/*</url-pattern>
<url-pattern>/redirect/*</url-pattern>
<url-pattern>/regions/*</url-pattern>
<url-pattern>/Search.jsp</url-pattern>
<url-pattern>/select/*</url-pattern>
<url-pattern>/SelectReturn.jsp</url-pattern>
<url-pattern>/Tasks.jsp</url-pattern>
<url-pattern>/UI-INF/*</url-pattern>
<url-pattern>/utils/*</url-pattern>
<url-pattern>/WcmAdmin.jsp</url-pattern>
<url-pattern>/WcmAuthor.jsp</url-pattern>
<url-pattern>/WcmBootstrap.jsp</url-pattern>
<url-pattern>/WcmCloseWindow.jsp</url-pattern>
<url-pattern>/WcmDefault.jsp</url-pattern>
<url-pattern>/WcmError.jsp</url-pattern>
<url-pattern>/WcmJavaViewer.jsp</url-pattern>
<url-pattern>/WcmObjectBookmark.jsp</url-pattern>
<url-pattern>/WcmPortletHelp.jsp</url-pattern>
<url-pattern>/WcmPortletSearch.jsp</url-pattern>
<url-pattern>/WcmQueueBookmark.jsp</url-pattern>
<url-pattern>/WcmSignIn.jsp</url-pattern>
<url-pattern>/WcmSitePreferences.jsp</url-pattern>
<url-pattern>/WcmUserPreferences.jsp</url-pattern>
<url-pattern>/WcmWorkflowsBookmark.jsp</url-pattern>
<url-pattern>/wizards/*</url-pattern>
<url-pattern>/Author/*</url-pattern>
<url-pattern>/axis/*.jws</url-pattern>
<url-pattern>/Browse/*</url-pattern>
<url-pattern>/ceTunnel</url-pattern>
<url-pattern>/CheckoutList/*</url-pattern>
<url-pattern>/downloadMultiTransferElement/*</url-pattern>
<url-pattern>/ExternalUrl/*</url-pattern>
<url-pattern>/findRecordTarget</url-pattern>
<url-pattern>/formCallback/*</url-pattern>
```

```
      <url-pattern>/getAnnotSecurity/*</url-pattern>
      <url-pattern>/getCEAnnotations/*</url-pattern>
      <url-pattern>/getContent/*</url-pattern>
      <url-pattern>/getForm/*</url-pattern>
      <url-pattern>/getISAnnotations/*</url-pattern>
      <url-pattern>/getISAnnotSecurity/*</url-pattern>
      <url-pattern>/getISContent/*</url-pattern>
      <url-pattern>/getMultiContent/*</url-pattern>
      <url-pattern>/getPreview</url-pattern>
      <url-pattern>/getProcessor/*</url-pattern>
      <url-pattern>/getRealms/*</url-pattern>
      <url-pattern>/getUsersGroups/*</url-pattern>
      <url-pattern>/Inbox/*</url-pattern>
      <url-pattern>/integrationCommandProxy</url-pattern>
      <url-pattern>/integrationResponse</url-pattern>
      <url-pattern>/integrationResponseProxy</url-pattern>
      <url-pattern>/integrationWebBasedCommand</url-pattern>
      <url-pattern>/keepAlive</url-pattern>
      <url-pattern>/launch/*</url-pattern>
      <url-pattern>/PublicQueue/*</url-pattern>
      <url-pattern>/putContent/*</url-pattern>
      <url-pattern>/QuickSearch/*</url-pattern>
      <url-pattern>/signingServlet/*</url-pattern>
      <url-pattern>/transport/*</url-pattern>
      <url-pattern>/upload/*</url-pattern>
      <url-pattern>/vwsimsoapservlet</url-pattern>
      <url-pattern>/vwsoaprouter</url-pattern>
      <url-pattern>/Workflows/*</url-pattern>  Move the closing comment tag from
      here to the location indicated at the beginning of this example.
   </web-resource-collection>
```

c. Locate the section <auth-constraint>, comment the wild-card (*) role-name as noted in the file comments below.

```
<auth-constraint>
   <!-- <role-name>*</role-name> -->
   <!-- For WebSphere 6, use the role-name line below instead of the wildcard role
   above.-->

   <role-name>All Authenticated</role-name>

   <!-- For WebSphere 6, add this security-role element below the login-config
   element (below).
      <security-role>
         <description>All Authenticated</description>
         <role-name>All Authenticated</role-name>
      </security-role>
            -->
</auth-constraint>
```

d. Locate the end of the </login-config> element, and add the All Authenticated users role-element after the closing tag.

```
<security-role>
   <description>All Authenticated</description>
   <role-name>All Authenticated</role-name>
</security-role>
```

e. Search for the first instance of a <security-constraint> tag, and add the following <security-constraint> tag before that tag.

CAUTION Enter the information below as single lines without line breaks.

```
<security-constraint>
   <web-resource-collection>
      <web-resource-name>action</web-resource-name>
      <description>Define the non-secured resource</description>
      <url-pattern>/P8BPMWSBroker/*</url-pattern>
   </web-resource-collection>
</security-constraint>
```

3. Save your changes to web.xml and close the file.

4. If your site uses SSO, Continue on with "(SSO Only) To edit web.xml for SSO (optional)" on page 356, otherwise continue with "To configure Application Engine (WebSphere)" on page 357.

(SSO Only) To edit web.xml for SSO (optional)

NOTE Perform this procedure only if your site uses SSO with a proxy server. You must use this procedure to modify web.xml to enable SSO.

1. Open web.xml for editing.

2. At the end of web.xml, comment out the <login-config> element, as follows:

```
<!--
<login-config>
   <auth-method>FORM</auth-method>
   <realm-name>AE Workplace</realm-name>
   <form-login-config>
      <form-login-page>/ContainerLogin.jsp</form-login-page>
      <form-error-page>/ContainerError.jsp</form-error-page>
   </form-login-config>
</login-config>
-->
```

3. As needed, set the ssoProxyContextPath, ssoProxyHost, ssoProxyPort, and ssoProxySSLPort.

   These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, you must set the corresponding sso<*proxy host element*> parameter for that element in the URL to the value for the SSO proxy host.

   The default settings are (in **bold** below):

```
<init-param>
   <param-name>ssoProxyContextPath</param-name>
   <param-value></param-value>
</init-param>
<init-param>
   <param-name>ssoProxyHost</param-name>
   <param-value></param-value>
</init-param>
<init-param>
   <param-name>ssoProxyPort</param-name>
   <param-value></param-value>
</init-param>
<init-param>
```

```
    <param-name>ssoProxySSLPort</param-name>
    <param-value></param-value>
</init-param>
```

In general, the init parameters above must be configured as follows:

- **ssoProxyContextPath**: Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

  For example, if the Workplace deploy host URL is `http://deploy_server:2809/Workplace` and the SSO proxy host URL is `http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

  ```
  <param-name>ssoProxyContextPath</param-name>
  <param-value>/fn/Workplace</param-value>
  ```

- **ssoProxyHost**: Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

  For example, if the host URL where Workplace is deployed is `http://deploy_server/Workplace` and the corresponding SSO proxy host URL is `http://sso_proxy_server/Workplace`, then use the following:

  ```
  <param-name>ssoProxyHost</param-name>
  <param-value>sso_proxy_server</param-value>
  ```

- **ssoProxyPort**: Set the value to the http port on the SSO proxy host.

  For example:

  ```
  <param-name>ssoProxyPort</param-name>
  <param-value>80</param-value>
  ```

- **ssoProxySSLPort**: Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

  For example:

  ```
  <param-name>ssoProxySSLPort</param-name>
  <param-value>443</param-value>
  ```

4. Save your changes to web.xml and close the file.

To configure Application Engine (WebSphere)

1. Open the WebSphere administrative console.
2. Set JVM settings for JAAS login configuration and memory settings.
   a. In the Java Virtual Machine settings, set the JAAS login entry in the Generic JVM argument field to one of the following (do not enter the linebreaks):

   ```
   -Djava.security.auth.login.config=ae_install_path\CE_API\config
   \jaas.conf.WebSphere
   ```

   NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

CAUTION Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the field.

CAUTION (Windows only) On WebSphere for Windows the path cannot contain a space. You must use 8.3-notation for the install path information.

If *AE_install_path* is:

```
C:\Program Files\FileNet\AE
```

use:

```
C:\Progra~1\FileNet\AE
```

b. Set the Initial and Maximum Heap Size.

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

c. Save your changes.

3. (For installations with Content Engine and Application Engine collocated on the WebSphere server, but in different WebSphere profiles) Create an additional JVM property for different WebSphere profiles.

Perform the following steps on *both* the Content Engine profile and the Application Engine profile:

a. In the Java Virtual Machine settings, create a new Custom Property:

com.ibm.websphere.orb.uniqueServerName

b. Set the Value to true.

c. Save your changes.

d. Restart WebSphere.

4. Verify that Application Integration installation is enabled.

a. In the Environment area, in Virtual Hosts, navigate to **default_host** (or the host your application is deployed under).

b. Verify that MIME type is set to application/octet-stream, or set it if it is not. Include .exe in the Extension(s) field.

c. Save your changes.

5. Configure Light weight Third Party Authentication (LTPA).

NOTE Skip this step if your Application Engine and Content Engine are located on the same WebSphere application server profile.

a. On the Content Engine server, do the following:

NOTE If you are already using LTPA with your Content Engine application server, you only need to export the existing keys per Step iii thru Step iv below.

i. Log in to the WebSphere adminstrative console.

ii. If LTPA is not already configured for Content Engine, create an LTPA password, specify the location for the key file on the Content Engine server, and generate the keys.

NOTE For password restrictions, see the WebSphere documentation. If you have already configured Content Engine for LTPA, use the existing password in the Application Engine configuration below.

iii. Export the keys using the WebSphere administrative console.

iv. Copy the key file from CE server location you specified above to a directory on the Application Engine server. On Windows, for example, C:\LTPA*ltpa_key_name*.

v. Save your changes, and stop and restart WebSphere.

b. On the Application Engine server, do the following:

i. Log in to the WebSphere adminstrative console.

ii. Create an LTPA password, using the same password as the Content Engine.

iii. Set the timeout value for LTPA authentication to a value that is larger than the default.

For example, if the timeout value is set to 2 hours, the LTPA key expires and end users will not be able to log in to Workplace after having been logged in for 2 hours.

iv. Specify the path for the key file that you copied to the Application Engine server. For example, C:\LTPA*ltpa_key_name*.

v. Import the keys from the Key file.

vi. Set the following Security settings:

– Turn on (check) **Enable Administrative Security flag**.

– Turn on (check) **Enable application security flag**.

– Turn off (uncheck) **Java 2 Security**.

NOTE The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.

– Set the Active Authentication Mechanism to LTPA (Light weight Third Party Authentication).

6. Configure Lightweight Directory Access Protocol (LDAP).

a. Set "Standalone LDAP registry" to current for the available realm definition.

b. Configure the LDAP provider to exactly match the corresponding settings on the Content Engine application server.

– Primary administrative user name

– Select “Automatically generated server identity."

– Type

– Host

- Port
- Base distinguished name (DN)
- Bind distinguished name (DN)
- Bind password

c. Configure the Advanced Lightweight Directory Access Protocol (LDAP) user registry settings to exactly match the corresponding settings from the Content Engine application server.

- User filter
- Group Filter
- User ID map
- Group member ID map
- Certificate map mode
- Certificate filter

d. Save these settings.

NOTE If you are using Tivoli as your directory server, perform Step e. For all other directory servers, continue with Step f, below.

e. (Tivoli only) Configure WebSphere security for Tivoli.

i. Shut down Websphere.

ii. On the Application Engine WebSphere server, back up the security.xml file. The file is in the following location:

install_path/WebSphere/AppServer/profiles/*appserver_name*/config/cells/*ServerName*Node01Cell/

iii. On the Content Engine WebSphere server, open the security.xml file, located in a path similar to the one in Step iii. Copy the “userRegistries”, “searchFilter”, and “host” entries, as shown in the example below (in bold). Do not use the example settings, as your settings will be different.

```
  <userRegistries xmi:type="security:LocalOSUserRegistry"
xmi:id="LocalOSUserRegistry" serverId="" serverPassword="{xor}" realm=""
useRegistryServerId="true" primaryAdminId=""/>

  <userRegistries xmi:type="security:CustomUserRegistry"
xmi:id="CustomUserRegistry_1" useRegistryServerId="true" primaryAdminId=""
customRegistryClassName="com.ibm.websphere.security.FileRegistrySample"/>

  <userRegistries xmi:type="security:LDAPUserRegistry"
xmi:id="LDAPUserRegistry_1" serverId="CEAdmin"
serverPassword="{xor}Dz4sLCgwLTtu" realm="vauxhall:389" ignoreCase="true"
useRegistryServerId="false" primaryAdminId="CEAdmin"
type="IBM_DIRECTORY_SERVER" sslEnabled="false" sslConfig=""
baseDN="dc=gardens"
bindDN="cn=CEAdmin,ou=Shared,ou=Engineering,ou=FileNet,dc=gardens"
bindPassword="{xor}Dz4sLCgwLTtu" searchTimeout="120" reuseConnection="true">
```

```
    <searchFilter xmi:id="LDAPSearchFilter_1"
userFilter="(&amp;(cn=%v)(objectclass=person))"
groupFilter="(&amp;(cn=%v)(|(objectclass=groupOfNames)(objectclass=groupOfUn
iqueNames)))" userIdMap="cn" groupIdMap="*:cn"
groupMemberIdMap="memberof:member" certificateMapMode="EXACT_DN"
certificateFilter=""/>

    <hosts xmi:id="EndPoint_1" host="vauxhall" port="389"/>

  </userRegistries>

  <userRegistries xmi:type="security:WIMUserRegistry"
xmi:id="WIMUserRegistry_1" serverId="" serverPassword="{xor}"
realm="defaultWIMFileBasedRealm" ignoreCase="true"
useRegistryServerId="false" primaryAdminId="administrator"
registryClassName="com.ibm.ws.wim.registry.WIMUserRegistry"/>

  <authConfig xmi:id="AuthorizationConfig_1" useJACCProvider="false">

    <authorizationProviders xmi:id="AuthorizationProvider_1"
j2eePolicyImplClassName="com.tivoli.pd.as.jacc.TAMPolicy" name="Tivoli Access
Manager"
policyConfigurationFactoryImplClassName="com.tivoli.pd.as.jacc.TAMPolicyConf
igurationFactory"
roleConfigurationFactoryImplClassName="com.tivoli.pd.as.jacc.TAMRoleConfigur
ationFactory"
initializeJACCProviderClassName="com.tivoli.pd.as.jacc.cfg.TAMConfigInitiali
ze" requiresEJBArgumentsPolicyContextHandler="false"
supportsDynamicModuleUpdates="true"/>

  </authConfig>
```

iv. On the Application Engine WebSphere server, open the security.xml file for editing, and paste in the “userRegistries”, “searchFilter”, and “host” entries you copied from the Content Engine file, in the location shown above.

v. Save and close the file.

vi. Restart WebSphere, and log in to the administrative console.

vii. Navigate to Security > Secure administration, applications and infrastructure.

NOTE If you make any changes to your LDAP server settings in the WebSphere Administrative Console after updating the security.xml file, these settings will be overwritten and you will have to add the entries back to the file.

f. Turn on (check) the Enable administrative security flag.

g. Turn off (uncheck) Use Java 2 security to restrict application access to local resources.

h. Save your changes to the master configuration.

i. Test the connection on the Standalone LDAP registry page. If the test fails, correct the error before proceeding. If it passes, click **OK** to return to the previous page.

j. Save your changes.

7. Continue with “To configure the server ports” on page 362.

To configure the server ports

This configuration is not required but is recommended.

1. Stop the WebSphere server.
2. Make a backup copy of serverindex.xml located in:

   WAS_HOME\profiles\default\config\cells*machine_name*Node01Cell\nodes*machine_name*Node01\

3. Edit serverindex.xml.

   Locate the <specialEndpoints> section, and change the port numbers for the three SSL listener addresses to “0” as shown below:

```
<specialEndpoints xmi:id="NamedEndPoint_1155689929072"
endPointName="SAS_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929072" host="host_name" port="0"/>
</specialEndpoints>
<specialEndpoints xmi:id="NamedEndPoint_1155689929073"
endPointName="CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929073" host="host_name" port="0"/>
</specialEndpoints>
<specialEndpoints xmi:id="NamedEndPoint_1155689929074"
endPointName="CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929074" host="host_name" port="0"/>
</specialEndpoints>
```

4. Save and close the serverindex.xml file.
5. Restart WebSphere.
6. Continue with “Manually copy custom data” on page 372.

Task 6b: Configure Application Engine upgrades from 3.5.*x* (WebLogic)

This topic covers the configuration of your Application Engine application (Workplace) on WebLogic. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using SSO, edit web.xml. See "To edit web.xml for SSO (optional)" on page 363.
2. Modify the application server startup script. See "To modify the application server startup script" on page 364.
3. Configure Application Engine. See "To configure Application Engine (WebLogic)" on page 366.
4. Modify config.xml to support passing user credentials to clients such as Application Integration and WebDAV. See "To enable passing user credentials to client applications (WebLogic)" on page 367.

To edit web.xml for SSO (optional)

NOTE Perform this procedure only if your site uses SSO with a proxy server. You must modify web.xml to enable SSO.

1. Make a backup copy of web.xml.

   AE_install_path/Workplace/WEB-INF/web.xml

2. Edit web.xml.

   a. Set the parameter perimeterChallengeMode to true, as in:

   ```
   <init-param>
      <param-name>perimeterChallengeMode</param-name>
      <param-value>true</param-value>
   </init-param>
   ```

   b. As needed, set the ssoProxyContextPath, ssoProxyHost, ssoProxyPort, and ssoProxySSLPort.

   These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, then you must set the corresponding sso* parameter for that element in the URL to the value for the SSO proxy host.

   The default settings are (in **bold**):

   ```
   <init-param>
      <param-name>ssoProxyContextPath</param-name>
      <param-value></param-value>
   </init-param>
   <init-param>
      <param-name>ssoProxyHost</param-name>
      <param-value></param-value>
   </init-param>
   <init-param>
      <param-name>ssoProxyPort</param-name>
   ```

```
    <param-value></param-value>
</init-param>
<init-param>
    <param-name>ssoProxySSLPort</param-name>
    <param-value></param-value>
</init-param>
```

In general, the init parameters above should be configured as follows:

- **ssoProxyContextPath**: Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

  For example, if the Workplace deploy host URL is `http://deploy_server:2809/Workplace` and the SSO proxy host URL is `http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

  ```
  <param-name>ssoProxyContextPath</param-name>
  <param-value>/fn/Workplace</param-value>
  ```

- **ssoProxyHost**: Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

  For example, if the host URL where Workplace is deployed is `http://deploy_server/Workplace` and the corresponding SSO proxy host URL is `http://sso_proxy_server/Workplace`, then use the following:

  ```
  <param-name>ssoProxyHost</param-name>
  <param-value>sso_proxy_server</param-value>
  ```

- **ssoProxyPort**: Set the value to the http port on the SSO proxy host.

  For example:

  ```
  <param-name>ssoProxyPort</param-name>
  <param-value>80</param-value>
  ```

- **ssoProxySSLPort**: Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

  For example:

  ```
  <param-name>ssoProxySSLPort</param-name>
  <param-value>443</param-value>
  ```

3. Save your changes to web.xml and close the file.

To modify the application server startup script

1. Stop the WebLogic application server if running.
2. Make a backup copy of the application server startup script.

   Back up startWebLogic.cmd for Windows or startWebLogic.sh for UNIX.

   NOTE If you are not using a WebLogic domain, backup startWLS.cmd for Windows or startWLS.sh for UNIX.

   a. Edit the MEM_ARGS variable.

Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

NOTE If the MEM_ARGS variable doesn't exist, add it to the startup script.

- For all systems except those using JRockit JAVA.

  Append the following to the MEM_ARGS variable:

  ```
  -XX:MaxPermSize=sizem
  ```

  where *size* is the value, in MB, of the MaxPermSize.

  Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

- For systems using JRockit JAVA.

  Append the following to the MEM_ARGS variable:

  ```
  -Xgc:gencon
  ```

b. Configure JAAS login.

Add one of the following right after the classpath entry for WebLogic.

CAUTION Enter the jaas_login entry (bold below) as a single line without line breaks. Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the script.

NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

Windows

```
@REM Jaas Login configuration setting
set JAAS_LOGIN=%JAAS_LOGIN% -Djava.security.auth.login.config=
"AE_install_path\CE_API\config\jaas.conf.WebLogic"
```

UNIX

```
# Jaas Login configuration setting
JAAS_LOGIN="${JAAS_LOGIN}" -Djava.security.auth.login.config=
"AE_install_path/CE_API/config/jaas.conf.WebLogic"
```

AIX

```
# Jaas Login configuration setting
JAAS_LOGIN="${JAAS_LOGIN}" -Djava.security.auth.login.config=
"AE_install_path/CE_API/config/jaas.conf.WebLogic"
JAAS_LOGIN="${JAAS_LOGIN}" -
Dlogin.configuration.provider=com.ibm.security.auth.login.ConfigFile
```

c. Add %JAAS_LOGIN% section as indicated in the examples below in **bold**.

Windows - in the WLS_REDIRECT_LOG settings

```
If "%WLS_REDIRECT_LOG%"==""(
    echo Starting WLS with line:
```

```
   echo %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN%
-Dweblogic.Name=%SERVER_NAME% -
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS%
    %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN% -
Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
Dweblogic.management.password=%WLS_PW% -
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS%
) else (
echo Redirecting output from WLS window to %WLS_REDIRECT_LOG%
%JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %JAAS_LOGIN% -
Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
Dweblogic.management.password=%WLS_PW% -
Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
%SERVER_CLASS% >”%WLS_REDIRECT_LOG%” 2>&1
)
```

UNIX - in the WLS_REDIRECT_LOG settings

```
${JAVA_HOME}/bin/java ${JAVA_VM} -version

if [ "${WLS_REDIRECT_LOG}" = "" ] ; then
        echo "Starting WLS with line:"
        echo "${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS} -
Dweblogic.Name=${SERVER_
NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_CLASS
}"
        ${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
${JAAS_LOGIN} -Dweblogic.Name=${S
ERVER_NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_
CLASS}
else
        echo "Redirecting output from WLS window to ${WLS_REDIRECT_LOG}"
        ${JAVA_HOME}/bin/java ${JAVA_VM} ${MEM_ARGS} ${JAVA_OPTIONS}
${JAAS_LOGIN} -Dweblogic.Name=${S
ERVER_NAME} -Djava.security.policy=${WL_HOME}/server/lib/weblogic.policy
${PROXY_SETTINGS} ${SERVER_
CLASS}  >"${WLS_REDIRECT_LOG}" 2>&1
fi
```

3. Save and close the server startup script.

To configure Application Engine (WebLogic)

1. (If you selected Container-Managed Authentication during the installation) Enable trust between WebLogic domains for the Content Engine domain and the Application Engine domain.

   Do the following on both the Content Engine application server and the Application Engine application server.

   a. In the WebLogic Administration Console, in the security settings, enter a password for the domain. You must enter the same password for both the Content Engine domain and Application Engine domain.

   b. Save your changes.

   c. Restart the server if needed.

   d. Repeat this procedure in each domain for which you want to enable trust.

2. (If you selected Container-Managed Authentication during the installation) Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
   a. Refer to your Content Engine installation worksheet items and the WebLogic Administration Console settings for Compatibility Security > Realms for Authentication Provider, users, and groups on Content Engine.

      Configure the LDAP provider to exactly match the settings from the Content Engine server.

      - Group Base DN:
      - User Name Attribute:
      - Port:
      - User Base DN:
      - Principal:
      - Credential:
      - Confirm Credential:
      - Host:
      - User From Name Filter:
      - Group From Name Filter:

   b. Restart WebLogic.
3. Set permissions for the user running the application server.

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed AE (*AE_install_path*).
4. Continue with "To enable passing user credentials to client applications (WebLogic)" on page 367.

To enable passing user credentials to client applications (WebLogic)

Perform this procedure to enable passing user credentials between Application Engine and its client applications such as WebDAV and Application Integration.

CAUTION If you do not make this change to config.xml, then end users will be prompted to enter their user name and password to complete any client operations, such as adding a document.

1. Stop the WebLogic server.
2. Make a backup copy of config.xml located in deployment directory.

   For example:

   BEA_home/bea/user_projects/domains/*domain_name*/config/config.xml
3. Edit config.xml.

   CAUTION The enforce-valid-basic-auth-credentials entry should be entered as a single line without line breaks.

a. Locate the <security-configuration> section and add the following line to the end of the section, just before the </security-configuration> tag:

```
<enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-
credentials>
```

b. Save your changes to config.xml and close the file.

4. Restart WebLogic.

5. Continue with "Manually copy custom data" on page 372.

Task 6c: Configure Application Engine upgrades from 3.5.*x* (JBoss)

This topic covers the configuration of your Application Engine application (Workplace) on JBoss. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

To modify the application server startup script

1. Stop JBoss if it is running.
2. Make a backup copy of the application server startup script.

   UNIX

   run.sh

   Windows

   run.bat
3. Edit the application server startup script Java settings.
   a. Add a line to specify the path to the JDK provided by JBoss, as shown in the following example (Windows):

      ```
      set JAVA_HOME=C:\Program Files\Java\jdk1.5.0_06
      ```

      NOTE If your JDK is different from version 1.5.0, substitute your version for the one listed above.

   b. Update the JAVA_OPTS memory settings.

      Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

      In the `JAVA_OPTS` line, change the -Xms and -Xmx values (**bold**) for your configuration.

      Example (Windows):

      ```
      set JAVA_OPTS=%JAVA_OPTS% -Xms128m -Xmx512m
      ```

      Refer to your application server vendor's recommendation for Initial and Maximum heap size values. 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.

   c. Save your edits.
4. Configure JAAS login.

   Add one of the following right after the "$JAVA" $JAVA_OPTS (UNIX) or "%JAVA%" %JAVA_OPTS (Windows) entry in the startup script.

   CAUTION Enter the jaas_login entry (bold below) as a single line without line breaks. Do not copy and paste the text from this guide because hidden formatting can cause problems with the entry. Instead, type the entry into the script.

NOTE Your path may be slightly different depending on the version of your client installations, or whether you have chosen a custom path for installation. Verify the location of the file before you enter the path.

UNIX

```
"$JAVA" $JAVA_OPTS -Djava.security.auth.login.config=”/opt/FileNet/AE/CE_API/
config/jaas.conf.JBoss” "-Djava.endorsed.dirs=$JBOSS_ENDORSED_DIRS" -classpath
"$JBOSS_CLASSPATH" org.jboss.Main $@
```

Windows

```
"%JAVA%" %JAVA_OPTS% "-Djava.security.auth.login.config=C:\Program
Files\FileNet\AE\CE_API\config\jaas.conf.JBoss" "-
Djava.endorsed.dirs=%JBOSS_ENDORSED_DIRS%" -classpath "%JBOSS_CLASSPATH%"
org.jboss.Main %*
```

5. Save and close the server startup script.
6. Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
   a. On the Application Engine server, open login-config.xml, located in *JBoss_home*/server/*server_name*/conf, for editing.
   b. Set the <application-policy name="FileNet"> entry identical to the corresponding entry in the login-config.xml file on the Content Engine server.
   c. Restart JBoss.
7. Set permissions for the user running the application server.

   NOTE On Windows, the following is only required for NTFS formatted partitions

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed Application Engine (*AE_install_path*).
8. (Optional) Disable JBoss logging.

   In development mode, JBoss creates a large number of HTTP Access, “INFO”, “DEBUG” and “TRACE” log messages. This can cause unexpected behavior in the deployed IBM FileNet software. Using the following procedure, you can limit this type of excessive JBoss logging.

   NOTE When logging is disabled, error messages will still be displayed in the JBoss console.

   a. Edit the log4j.xml file (*JBOSS_home*/server/*server_name*/conf/log4j.xml).
      i. Change all threshold values and priority values from "INFO", "DEBUG", or "TRACE" to "ERROR".
      ii. Delete or comment out the "Preserve messages in a local file" to turn off the server log.
   b. To turn off HTTP access logging, open jboss-service.xml with a text editor and delete or comment out the "Access logger" section.

      Location of jboss-service.xml:

      JBoss_Home/server/*server_name*/deploy/jbossweb-tomcat55.sar/META-INF

c. Open web.xml and change the logVerbosityLevel to "FATAL".

   Location of web.xml:

   JBoss_Home/server/*server_name*/deploy/jbossweb-tomcat55.sar/conf

d. Restart the JBoss server.

9. Continue with "Manually copy custom data" on page 372.

Task 7: Manually copy custom data

If your custom data was *not* retained in the upgrade you must manually copy the data from the version 3.5.*x* or 4.0.1 backup files (*deploy_backup*) to the upgraded files.

NOTE You must also copy any data you want to retain if you have made any custom add-ons or modifications to the Application Engine installation being upgraded and chose not to keep these changes after the upgrade. You should manually copy any other custom data you want to retain to the upgraded files listed in "Prepare for Application Engine upgrade" on page 332.

When you have completed the manual copy of your custom data, deploy the Workplace application using one of the following tasks:

- "Deploy upgraded Application Engine instances (WebSphere)" on page 373
- "Deploy upgraded Application Engine instances (WebLogic)" on page 375
- "Deploy upgraded Application Engine instances (JBoss)" on page 377

Task 8a: Deploy upgraded Application Engine instances (WebSphere)

This topic covers the deployment and start of your Application Engine application (Workplace) on WebSphere.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.exe (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.exe (Windows) from the following location:

     AE_install_path/deploy/

To deploy Application Engine

1. Log on to the WebSphere administrative console.
2. Navigate to the dialog for installing a new application.
3. Select the file to deploy.
   - (If the administrative console is running *locally*) Select **Local Path** and enter or browse to the location of the app_engine.war or app_engine.ear file created by the installation program (see below for the default path). Do not enter the machine name.
   - (If the administrative console is *remote*) Select **Server path** and enter the fully-qualified pathname to the app_engine.war or app_engine.ear file. Do not enter the machine name.

     AE_install_path/deploy

4. If you are deploying a WAR file, enter the context root.

   Enter `Workplace` and click **Next** to proceed to deploying a new application.

   NOTE The context root is the name of the application you log in to using the web interface, such as:

   `http://`*`ApplicationEngineServerName`*`:`*`port`*`#/`*`Context_Root`*.

5. Complete the dialogs for installing a new application, using the following settings:

   Application name: `Workplace`, or the name you chose to call the application.

   WebServer: The server you are planning to use. Verify that your application name is selected and associated with the correct WebServer.

   virtual host: Choose the default_host.

6. Save your configuration.

7. Configure the Classloader settings and change the polling interval to a number appropriate for your environment.

   Change Classloader order to have the classes loaded with parent classloader last.

   NOTE Do this only for the specific web application. Do not change the similar settings for the entire application server.

8. In the Manage Modules area, configure the Web Module Classloader setting.

   Change Classloader order to have the classes loaded with parent classloader last.

   NOTE Do this only for the specific web application. Do not change the similar settings for the entire application server.

9. (WebSphere 6.1) If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Map security roles to users/groups, and verify that the **All Authenticated** column is checked for the "All Authenticated" role.

10. (WebSphere 7.0) If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Security roles to user/group mapping. Select the “All authenticated” role and map it to “All Authenticated in Applications realm”.

11. Set permissions for the user running the application server.

    The user that will be running the application server must have read/write permissions on the following (default) folders:

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

    AE_install_path

12. Save all your changes.

13. Stop and restart WebSphere.

14. Continue with Task 9 “Complete post-upgrade Application Engine configuration” on page 379.

Task 8b: Deploy upgraded Application Engine instances (WebLogic)

This topic covers the deployment of your Application Engine application (Workplace) on WebLogic.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.bat (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.bat (Windows) from the following location:

     AE_install_path/deploy/

To deploy as “Workplace” or custom name using a WAR file

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users enter to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the app_engine.war file to reflect the name you want to use using the format *Application Name*.war.

Example:

The default app_engine.war will generate the following context root:

```
http://server_name:port#/app_engine
```

Renaming the WAR file Workplace.war will generate the following context root:

```
http://server_name:port#/Workplace
```

CAUTION You must rename the WAR file every time you regenerate it. The create_app_engine_war.sh/.bat script will by default create a file with the name app_engine.war.

To deploy Application Engine

1. From the WebLogic Administration Console, navigate to the domain you initially created for the Application Engine.
2. Prepare the WebLogic Administration Console to deploy the application.
3. Choose whether to deploy from an exploded folder (*AE_install_path*) or from the WAR or EAR file (default: app_engine.war or app_engine.ear in *AE_install_path*/deploy).
4. Accept the defaults for the deployment, except for the name for the deployment. Use “Workplace” instead of “appengine”.
5. Finish the deployment, and save and activate your changes.

   NOTE To verify that the deployment was successful, expand **Web Applications**. The web application Workplace will be listed.

6. Continue with Task 9 “Complete post-upgrade Application Engine configuration” on page 379.

Task 8c: Deploy upgraded Application Engine instances (JBoss)

This topic covers the deployment and start of your Application Engine application (Workplace) on JBoss.

To recreate the WAR or EAR file

Any time that you make changes to files in the /Workplace directory, such as changes to web.xml, you must recreate the WAR or EAR file and redeploy your changes.

NOTE Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  a. Verify that all modified /Workplace directory files have been saved.
  b. Recreate the app_engine.war file by running create_app_engine_war.sh (UNIX) or create_app_engine_war.exe (Windows) from the following location:

     AE_install_path/deploy/
- If you will be deploying from an EAR file.
  a. Verify that a newly recreated app_engine.war file exists.
  b. Recreate the app_engine.ear file by running create_app_engine_ear.sh (UNIX) or create_app_engine_ear.exe (Windows) from the following location:

     AE_install_path/deploy/

To deploy as “Workplace” or custom name using a WAR file

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users type to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the app_engine.war file to reflect the name you want to use using the format *Application Name*.war.

Example:

The default app_engine.war will generate the following context root:

```
http://server_name:port#/app_engine
```

Renaming the WAR file Workplace.war will generate the following context root:

```
http://server_name:port#/Workplace
```

CAUTION You must rename the WAR file every time you regenerate it. The create_app_engine_war.sh/.exe script will by default create a file with the name app_engine.war.

To deploy and start Application Engine

1. Deploy the Workplace application:

   To deploy from exploded directory

   a. On the JBoss server, copy the /Workplace folder from:

      AE_install_path

      to:

      JBOSS_home/server/default/deploy/

   b. Append the extension .war to the Workplace folder:

      JBOSS_home/server/default/deploy/Workplace.war

   To deploy from a WAR file

   On the JBoss server, copy the app_engine.war file from:

   AE_install_path/deploy

   to:

   JBOSS_home/server/default/deploy/

   To deploy from an EAR file

   On the JBoss server, copy the app_engine.ear file from:

   AE_install_path/deploy

   to:

   JBOSS_home/server/default/deploy/

2. Set permissions for the user running the application server.

   If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the following folders:

   NOTE For Windows this is only required for NTFS formatted partitions:

   JBOSS_home/server/default/deploy/app_engine.war/.ear

   AE_install_path

3. Start or restart the JBoss application server.

4. Verify that the application deployed successfully.

   Verify that the server.log file located in *JBOSS_home*/server/default/log lists deployment of the WAR or EAR file you used.

5. Continue with Task 9 “Complete post-upgrade Application Engine configuration” on page 379.

Task 9: Complete post-upgrade Application Engine configuration

Use the procedure in this topic to complete the configuration of upgraded Application Engine instances. Actions include modifying a configuration file, verifying Workplace sign-in, setting bootstrap preferences, establishing the Process Engine connection point, and verifying the version of an associated software integration.

To complete post-upgrade Application Engine configuration

1. (WebLogic only) If you upgraded from WebLogic 8.1.x to WebLogic 9.x or 10.x, you must set the case-sensitivity for the server.
   a. From the WebLogic Administration Console, click ***my_domain***.
   b. Set **Web App Files Case Insensitive** to false.
   c. Save and activate your changes.
2. If it is not already running, start the Workplace application on your application server.
3. Sign in to Workplace to test your connection.
   a. On any computer, open a browser and type:

      `http://ApplicationEngineServerName:port#/Workplace`

   b. Enter a user name and password, and click **Sign in**.
   c. If the Bootstrap Preferences page is displayed, you must set the preferences.

      Follow the instructions in "Set Application Engine bootstrap preferences" on page 148 to reset your bootstrap properties. Use the notes you made of your bootstrap settings during your upgrade preparation to complete this step.
4. Set the Process Engine Connection Point.

   The Process Router from version 3.5.x has been replaced with the Process Engine Connection Point.

   a. In Workplace click **Admin**.
   b. Click **Site Preferences**.
   c. Under General Settings > Tasks, select a **Process Engine Connection Point** from the drop down list.
   d. Click **Apply** and then **Exit**.
   e. Confirm that Application Engine is communicating with Process Engine.
      i. In Workplace, click **Tasks.**
      ii. Verify that the Tasks page displays.
   f. Sign out of Workplace.

5. (IBM FileNet P8 systems using Image Services Integration only) Verify that you are running a supported version of ISRA.

   See the *IBM FileNet P8 Hardware and Software Requirements* document for details on supported ISRA versions. To download this guide from the IBM support page, see "Access IBM FileNet documentation, compatibility matrices, and fix packs" on page 19.

   For more information on how to upgrade ISRA and configure Image Services Integration, see your ISRA documentation and "Upgrade ISRA servlet" on page 392.

Upgrade add-on components

You must upgrade all of the IBM FileNet P8 Platform core components before upgrading the add-on components listed in this topic. See "Upgrade and configure IBM FileNet P8 Platform" on page 197 for the tasks required to upgrade the IBM FileNet P8 Platform core components.

You can upgrade the add-on components in the following list, and any IBM FileNet P8 Platform expansion products, in any order, except that you must upgrade Process Analyzer before upgrading Process Simulator.

For details and procedures on updating add-on components, see the following topics:

- "Upgrade Application Integration and File Tracker" on page 382.
- "Upgrade IBM FileNet Publishing Components" on page 385.
- "Upgrade server-side scripts and COM objects" on page 386.
- "Upgrade ISRA servlet" on page 392.

Upgrade Application Integration and File Tracker

Upgrading to Workplace Application Integration 4.0.0 involves installing the new version on top of your existing Workplace Application Integration 3.5.*x* install. The upgrade installation program detects the earlier version and notifies you before proceeding with the upgrade.

During the upgrade, the installer detects the Microsoft applications that were previously integrated with Workplace and retains that configuration. For example, if you had Microsoft Outlook integrated with Workplace in the 3.5.*x* release, Microsoft Outlook is upgraded to 4.0.0 during the upgrade installation.

You cannot change your installed components during the upgrade. You can add or remove a Microsoft application from Application Integration after the upgrade using your Add/Remove Programs application.

Verify your computer meets the platform requirements documented 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.

Upgrade considerations

The procedures for upgrading Application Integration and File Tracker depend on which versions are currently installed on the client machine.

- If you upgrade from Application Integration version 3.5.1-001 or earlier, the upgrade to 4.0.0 will install File Tracker for you.
- If you upgrade from Application Integration version 3.5.1-002 or higher, then these conditions apply:
  - If you only had Application Integration installed (without File Tracker), then the upgrade to 4.0.0 will install File Tracker for you.
  - If you only had File Tracker installed (without Application Integration), then see "To upgrade Workplace File Tracker" on page 383.
  - If you had both Application Integration and File Tracker installed, then the order that you installed these determines how you upgrade.
    - If you installed File Tracker 3.5.1-002 or higher before you installed Application Integration, you must upgrade File Tracker before you upgrade Application Integration. See "To upgrade Workplace File Tracker" on page 383, and then see "To upgrade Workplace Application Integration" on page 383.
    - If you installed Application Integration 3.5.1-002 or higher, and later used File Tracker, then you only need to upgrade Application Integration.See "To upgrade Workplace Application Integration" on page 383

NOTE If you upgrade from Application Integration version 3.5.1-002 or earlier, there is a change in the behavior of entry tempates for adding email. Beginning with version 3.5.1-003, entry templates for adding an email are no longer restricted to using the email class or subclass.

To upgrade Workplace Application Integration

1. Log onto the client machine with Application Integration installed using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Click **Download Application Integration for Microsoft Office**. The File Download dialog box for your system appears. Do one of the following:
   - Click **Open** to run the program from its current location.
   - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the ApplicationIntegration.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the upgrade installer.

   NOTE If you have Workplace Application Integration 3.5.x currently installed, you are prompted about the impending upgrade. Click **Yes** to upgrade to the current version you are installing or click **No** to end the installation.
5. After the install is complete, click **Finish** to complete the upgrade process.

To upgrade Workplace File Tracker

1. Log on to the client machine with File Tracker installed using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download File Tracker** and do one of the following:
   - Click **Open** to run the program from its current location.
   - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the FileTracker.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the upgrade installer.
5. After the install is complete, click **Finish** to complete the upgrade process.

To verify your Workplace Application Integration upgrade

1. Start Microsoft Word.
2. From the File menu, click FileNet P8, point to Open Document, and then select Select Item. The Logon dialog box opens.
3. Log on using any valid domain account.
4. Click **Options** to view the Workplace address. The version number appears below the address.
5. Close all dialog boxes and close Microsoft Word.

To install component software updates

1. To obtain the latest component software updates, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the software updates and perform the installation procedures in the readmes on the component machine.

Upgrade IBM FileNet Publishing Components

If you have IBM FileNet Rendition Engine software currently installed, you can upgrade to the latest software version. For instructions, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide* at FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade.

Upgrade server-side scripts and COM objects

As with Content Engine 3.5.x, you can develop and plug in the following server-side components for the 4.x version of Content Engine: event action handlers, lifecycle action handlers, and document classifiers. Whereas these components had to be implemented as COM objects or scripts for Content Engine 3.5.x, they must be implemented as Java classes in Content Engine 4.x. Therefore, for Content Engine 4.x, you must convert any event action handlers, lifecycle action handlers, and document classifiers that you developed for Content Engine 3.5.x.

This section gives you an idea of the effort involved in converting 3.5.x script-based event action handlers to 4.x Java-implemented handlers. For details on developing and deploying event action handlers, lifecycle action handlers, and document classifiers, see the *Java and .NET Developer's Guide* and the *Java API Reference*.

In the following listings, compare the 3.5.x script-based event action handlers to the corresponding Java-implemented handlers required for Content Engine 4.x.

File Document Handler

This handler files a document to a specified folder, determined by the event fired on the document.

Content Engine 3.5.x JScript Version

```
function OnEvent (Event, Subscription)
{
  var doc = Event.SourceObject;
  if ( Event.IsOfClass("CreationEvent")  )
  {
      FileDocInFolder("/Docs", doc);
  }
  else if (Event.IsOfClass("ChangeClassEvent")  )
  {
      FileDocInFolder("/Archives", doc);
  }
}

function FileDocInFolder(otherFolderName, doc)
{
    var os = doc.ObjectStore;
    var rootFld = os.RootFolder;
    var fldSet = new Enumerator(rootFld.SubFolders);
    var subFld;
    for ( ; !fldSet.atEnd(); fldSet.moveNext() )
    {
       subFld = fldSet.item();
       if (subFld.Name == otherFolderName)
       {
          subFld.File(doc, 0, doc.DocumentTitle);
       }
    }
}
```

Content Engine 4.x Java Version

```
import com.filenet.api.constants.*;
import com.filenet.api.constants.DefineSecurityParentage;
import com.filenet.api.core.*;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class FileDocumentAction implements EventActionHandler
{
  public void onEvent(ObjectChangeEvent event, Id subscriptionId)
        throws EngineRuntimeException
  {
      Document doc = (Document)event.get_SourceObject();
      try
      {
            if (event.getClassName().equalsIgnoreCase("CreationEvent"))
                FileDocInFolder("/docs", doc);
            else if (event.getClassName().equalsIgnoreCase("ChangeClassEvent"))
                FileDocInFolder("/Archives", doc);
      }
      catch (Exception e)
      {
            throw new EngineRuntimeException(ExceptionCode.E_FAILED);
      }
  }

  public void fileDocInFolder(String folderName, Document doc)
  {
      try
      {
            Folder folder = (Folder)doc.getObjectStore().getObject("Folder", folderName);
            ReferentialContainmentRelationship rel = folder.file (doc,
                        AutoUniqueName.AUTO_UNIQUE, doc.get_Name(),
                        DefineSecurityParentage.DO_NOT_DEFINE_SECURITY_PARENTAGE);
            rel.save(RefreshMode.NO_REFRESH);
      }
      catch (Exception e)
     {
            e.printStackTrace();
      }
  }
}
```

Log Event Handler

This handler records events to a log file when documents are created.

Content Engine 3.5.x VBScript Version

`Public Sub OnEvent (EventObject, Subscription)`
`  Dim doc, message)`
`  Set doc = EventObject.SourceObject)`
`  WriteToLogFile (doc.Name & " was created on: " & Date))`
`End Sub`
 	
`Public Sub WriteToLogFile (message))`
`  Dim fso, ts, logFile)`
`  Set fso = CreateObject("Scripting.FileSystemObject"))`
`  Set logFile = fso.CreateTextFile("C:\log.txt"))`
`  logFile = nothing)`
`  Set ts = fso.OpenTextFile("C:\log.txt", 8, True))`
`  ts.Write (message))`
`  ts.WriteBlankLines (2))`
`  ts.Close)`
`  Set fso = Nothing)`
`  Set ts = Nothing)`
`End Sub`

Content Engine 4.x Java Version

```
import java.io.File;
import java.io.FileWriter;
import java.io.IOException;

import com.filenet.api.core.Document;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ErrorRecord;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class LogEventAction implements EventActionHandler
{
   public void onEvent(ObjectChangeEvent event, Id subscriptionId) throws EngineRuntimeException
   {
      try
      {
         Document doc = (Document)event.get_SourceObject();
         WriteToLogFile(doc.get_Name() + " was created on: "
            + new java.util.Date() + "\r\n");
      }
      catch (Exception e) {
         ErrorRecord er[] = {new ErrorRecord (e)};
         throw new EngineRuntimeException(e, ExceptionCode.EVENT_HANDLER_THREW, er);

      }
   }

   public void writeToLogFile(String message)
   {
      try
      {
         File outputFile = new File("C:\\log.txt");
         FileWriter out = new FileWriter(outputFile, true);
         out.write(message);
         out.close();
      }
      catch (IOException e) {
         ErrorRecord er[] = {new ErrorRecord (e)};
         throw new EngineRuntimeException(e, ExceptionCode.EVENT_HANDLER_THREW,er);
      }
   }
}
```

Send eMail Handler

This handler sends an email when a new document has been created.

Content Engine 3.5.x VBScript Version

Public Sub OnEvent (EventObject, Subscription)
  Dim myMail, MessageBody
  Set myMail = CreateObject("CDONTS.NewMail")
  myMail.From = "userI@company.com"
  myMail.To = "sysAdmin@company.com"
  myMail.Subject = "Event Notification--New Document created"
  MessageBody = "A document titled """ & Source.DocumentTitle & """ was Created at " & time & " on " & date & "."
  MessageBody = MessageBody + vbCrLf + Subscription.UserString
  myMail.Body = MessageBody
  myMail.Send
  Set myMail = Nothing
End Sub

Content Engine 4.x Java Version

```
import java.util.Date;
import java.util.Properties;

import javax.mail.Message;
import javax.mail.MessagingException;
import javax.mail.Session;
import javax.mail.Transport;
import javax.mail.internet.InternetAddress;
import javax.mail.internet.MimeMessage;

import com.filenet.api.core.*;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class EMailAction implements EventActionHandler
{
   public void onEvent(ObjectChangeEvent event, Id subscriptionId) throws EngineRuntimeException
   {
      Document doc = (Document)event.get_SourceObject();
      try
      {
         Properties props = new Properties();
         props.put("mail.smtp.host", "smtp.company.net");
         props.put("mail.smtp.port", "25");
         Session session = Session.getInstance(props);
         try {
            Message msg = new MimeMessage(session);
            msg.setFrom(new InternetAddress("userl@company.com" ));
            InternetAddress[] address = {new InternetAddress("sysAdmin@company.com" )};
            msg.setRecipients(Message.RecipientType.TO, address);
            msg.setSubject("Test E-Mail through Java");
            msg.setSentDate(new Date());
            msg.setText("Document " + doc.get_Name() + " created with ID " + doc.get_Id());
            Transport.send(msg);
         }
         catch (MessagingException mex) {
            mex.printStackTrace();
         }
      }
      catch (Exception e)
      {
         throw new EngineRuntimeException(ExceptionCode.E_FAILED);
      }
   }
}
```

Upgrade ISRA servlet

As part of upgrading FileNet Image Services Resource Adapter (ISRA) software to run with IBM FileNet P8, you must do the following:

- Ensure that Application Engine has been upgraded.
- Ensure that existing IBM FileNet ISRA software is installed and configured.

  For information on installing, configuring, and deploying FileNet ISRA, refer to the ISRA documentation on the IBM FileNet ISRA installation package.

  TIP Use the Sample Application shipped with IBM FileNet ISRA to confirm that the ISRA installation was successful.

  WARNING In an ISRA upgrade, take care to use the same library name (JNDI connection factory name) that has been previously set in the ISRA install. Changing this variable can cause conflicts when accessing documents.
- Do the following, as documented later in this task topic:
  - Upgrade the Application Engine ISRA Servlet, taking the following into account:
    - The servlet must be deployed on the same application server as FileNet ISRA.
    - The servlet does not need to be collocated with the Application Engine.
  - Configure Workplace site preferences for ISRA support.

ISRA SSL Support

The following table details the supported for SSL configurations.

SSL Configuration	SSL Support
ISRA Servlet and AE Collocated. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet and AE Collocated. AE Configured for SSL logon redirect to a local host.	Supported
ISRA Servlet and AE Collocated. AE and ISRA Servlet running under SSL.	Not Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a local host.	Supported
ISRA Servlet remote from AE. AE running under SSL, ISRA Servlet not running under SSL.	Supported
ISRA Servlet remote from AE. AE and ISRA Servlet running under SSL.	Not Supported

To upgrade the Application Engine ISRA Servlet

The FileNet P8 Application Engine installation CDs contain the ISRA servlet installation programs for the supported P8 AE platforms.

1. Log on to the application server machine:

   UNIX

   Log on as a user with write access to the /bin directory and read, write, execute access to the directory where you plan to install ISRA Servlet.

   Windows

   Log on as a member of the local Administrators group or as a user with equivalent permissions.

2. (WebSphere and WebLogic only) Back up and undeploy the ISRA Servlet application on the application server.

3. Stop the application server if it is running.

4. (JBoss only) Back up and undeploy the ISRA Servlet application.

5. Access the ISRA installation package, and start the Application Engine ISRA Servlet installation (upgrade) program:

   UNIX

   AE-ISRA-Servlet-4.0.2.0-*operating_system*.bin

   Windows

   AE-ISRA-Servlet-4.0.2.0-WIN.exe

6. Complete the installation (upgrade) program screens. For information on parameter values, see "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8*.

   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 the ISRA servlet installation program:

   – Click the AutoFilter drop-down arrow in the "Installation or Configuration Program" column header and select ISRA installer.

   – Click the AutoFilter drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

7. Check the file AE_ISRA_Servlet_install_log-4_0_2_0.txt, located in the *AE_israservlet_install path*/FileNet directory, to see if any errors occurred during the installation.

8. Install unlimited strength jar files.

   Perform this step only if your site is generating or accepting unlimited strength user tokens. Your system must be configured as follows:

   - The Application Engine ISRA Servlet is deployed on a different application server from the Application Engine server.
   - The **Create unlimited strength key** option was selected in the Application Engine User Token Security step of the Application Engine installation.

   NOTE Failure to perform the step will cause an EncryptionException when you log in to the IS Server.

9. (WebSphere and WebLogic only) Start the application server.
10. Deploy *AE_israservlet_install_path*/FileNet/ApplicationEngineISRAServlet/ae_isra.war in the same way you deployed the app_engine.war file for Workplace.
11. (JBoss only) Start the application server.
12. Verify the Application Engine ISRA Servlet installation. Do the following to use an available diagnostic tool to verify that the ISRA Servlet is installed and deployed correctly.

    a. Launch your browser.

    b. Enter the URL for the Application Engine ISRA Servlet. For example,

    ```
    http://ApplicationEngineISRAServlet_servername:port/
    ApplicationEngineISRAServlet/ISRA
    ```

    NOTE ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

    If the ISRA Servlet is installed and deployed correctly, a Congratulations message displays. For example:

    ```
    Congratulations! ISRA Interface Servlet is configured at this URL.

    WcmApiConfigFile = D:\ISRAInterface\jsp\WEB-INF\WcmApiConfig.properties

    WcmApiConfig file exists


    CryptoKeyFile/UserToken = C:\Program
    Files\FileNet\Authentication\UTCryptoKeyFile.properties

    CryptoKeyFile/UserToken exists


    FileNet ISRA classes are in the classpath

    com.filenet.is.ra.cci.FN_IS_CciConnectionSpec
    ```

To configure the Workplace site preferences for ISRA support

Application Engine Setup installs a pre-configured external service called Image Service, which includes the parameterized values necessary to access FileNet Image Service libraries from Workplace. Enable the service by setting the *Image Service* value in Site Preferences to **Show** (the default is Hide), as described in the following procedure.

1. Sign in to Workplace as a user having the Application Engine Administrators access role.
2. Launch Site Preferences as follows:
   a. Click **Admin**.
   b. Click **Site Preferences**.
3. Click **External Services** from the left options list.
4. Click **Modify** for the Image Service (under External Reference Services).

   The External Reference Service Settings site preference page displays.
5. Under General Information, locate *Show on Select File page* and change the value to Show.
6. Click **Accept**.
7. Click **Apply**.

To set the ISRA Interface Servlet URL

1. Click **Bootstrap**.
2. Under Preferences Settings, set the value of ISRA Interface Servlet URL. For example:

   `http://<servername>:<port>/ApplicationEngineISRAServlet/ISRA`

   NOTE ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.
3. Click **Apply.**
4. Click **Exit** to exit the Site Preferences.

To log on to Image Services via LDAP

To log on to the Image Services library using your LDAP account, configure ISRA and Image Services for LDAP authentication. If the LDAP account with which you accessed Workplace is not valid for the Image Services library, or if LDAP authentication is not configured, you will be prompted to log on to the Image Services library.

For information on configuring LDAP authentication for ISRA, refer to the *ISRA Installation and Deployment Guide*. For information on configuring LDAP authentication for Image Services, refer to the *Image Services System Tools Reference Manual*.

To access IS library documents

For informations about accessing IS library documents, see FileNet P8 Help topic User Help > Actions, preferences and tools > Actions > Documents > Add a document (Workplace).

To install component software updates

1. To obtain the latest component software updates, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the software updates and perform the installation procedures in the readmes on the component machine.

Remove software

This section includes:

For instructions on removing the Rendition Engine software, see the IBM FileNet P8 guide FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade.

Remove the IBM FileNet P8 documentation

The following topic explains how to remove the documentation for the FileNet P8 Platform and its expansion products.

NOTE Because of the number of application server configuration possibilities, these examples should be used only for reference. Your specific installation directories and application names may vary.

To remove the FileNet P8 documentation

NOTE In some Windows installations where NTFS is used for the file system (Not FAT or FAT32), there is a known issue with deleting files (and folders) that are longer than 256 characters. For example, if you use a default WebSphere installation location, you may encounter an error where the FileNet P8 Platform documentation files cannot be properly deleted due to the number of characters in the file/folder names. See the Microsoft Knowledge Base article http://support.microsoft.com/?kbid=320081 for additional information about deleting files (and folders) in this environment.

To remove the FileNet P8 documentation from a WebSphere application server

1. Log on to the WebSphere FileNet P8 documentation server.
   - UNIX - Log on as a user with delete access to where the FileNet P8 Platform documentation files are installed.
   - Windows - Log on with a user account that has local Administrative (or Account Operators and Server Operators) rights.
2. Verify that the WebSphere server is running.
3. From the WebSphere administrative console (for example, http://localhost:9060/ibm/console), **Uninstall** the FileNet P8 documentation site (for example, *ecm_help.war*).
4. Delete the entire FileNet P8 documentation folder (for example, *ecm_help.war*) structure from the installation location.
5. Delete any temp folder(s) or log files for the FileNet P8 documentation.

   WARNING Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

To remove the FileNet P8 documentation from a WebLogic application server

1. Log on to the WebLogic application server.
   - UNIX - Log on as a user with delete access to where the **ecm_help** files are located.
   - Windows - Log on with a user account that has local Administrative rights (or Account Operators and Server Operators).
2. Verify that the WebLogic server is running.
3. From the WebLogic Server Administration Console (for example, http://<machinename>:7001/console), **Delete** the FileNet P8 documentation site (for example, *ecm_help*).

4. Delete all folders and files including any temp folder(s) or log files for the FileNet P8 documentation.

   WARNING Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

To remove the FileNet P8 documentation from a JBoss application server

1. Log on to the JBoss FileNet P8 documentation server.
   - UNIX - Log on as a user with delete access to where the FileNet P8 Platform documentation files are installed.
   - Windows - Log on with a user account that has local Administrative (or Account Operators and Server Operators) rights.
2. Stop the JBoss application server.
3. Delete all folders and files including any temp folder(s) or log files for the FileNet P8 documentation (for example, ecm_help).

   WARNING Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

Remove Content Engine

You can uninstall an entire Content Engine installation or selected Content Engine components.

NOTE Uninstalling Content Engine does *not* undeploy it. You must use the application server console or commands to remove the Content Engine EAR file from the application server.

Use one of the following procedures to uninstall part or all of Content Engine.

To remove part or all of a Content Engine installation using the Windows Control Panel

1. Choose **Control Panel > Add/Remove Programs**.
2. Highlight P8 Content Engine in the list of currently installed programs and click **Change/Remove** to launch the Uninstall P8 Content Engine program.
3. At the Uninstall Options screen, choose to uninstall some or all Content Engine components. If you choose to uninstall all components, continue at Step 5; otherwise continue at Step 4.
4. In the Choose Components screen, choose which components to uninstall.
5. In the Uninstall Complete screen, note the directories and files that cannot be removed by the uninstall program and choose whether to have the program restart the machine, or restart it manually.
6. Remove the remaining directories and files, as noted in the Uninstall Complete screen.
7. If you want to completely remove all traces of the Content Engine installation, delete the C:\Program Files\FileNet directory.

To remove an entire Content Engine installation (UNIX)

1. Navigate to the directory *install_path*/FileNet/ContentEngine/_ceuninst, created by the Content Engine installer.
2. To uninstall Content Engine interactively, run the following command:

   ```
   ce_uninstaller
   ```

To remove Content Engine silently

In silent mode, the unisntaller removes all Content Engine Components.

- To uninstall Content Engine silently, run one of the following commands:

  UNIX

  ```
  ce_uninstaller -i
  ```

  Windows

  ```
  ce_uninstaller.exe -i
  ```

To remove data associated with Content Engine

After uninstalling Content Engine, you can remove its associated data, as follows:

1. Use the application server console or command lines to undeploy Content Engine.
2. Use your database tools to drop any databases or tablespaces for the object stores and the GCD.
3. Use your LDAP tools to delete users and groups you created in *Plan and Prepare Your Environment for IBM FileNet P8*.
4. Use your operating system commands to delete any directories, users and groups used for installing and administering Content Engine; and delete file-storage-area directories containing content (for example, documents) and index-area directories (K2 collections).

Remove Content Search Engine

To completely remove 4.5.x Content Search Engine (Autonomy K2) and collections from your IBM FileNet P8 platform installation, you must disable full-text indexing and remove the Autonomy K2 installations from all servers associated with the Content Engine.

CAUTION If you remove or disable Autonomy K2 before you disable CBR and full-text indexing in Enterprise Manager, your system will be rendered unusable and require considerable reconstruction.

NOTES

- This procedure presumes you have a running installation of 4.5.x Content Search Engine and that you have existing collections.
- If you intend to remove Content Engine, skip to step 4.
- The paths listed in this procedure assume you have used the suggested install path. If you have installed to another location, substitute the path as appropriate.

To remove Autonomy K2 and CBR

1. Launch Enterprise Manager.
2. Disable CBR for any classes that have been enabled for CBR.
   a. Right-click the class you want to configure and click **Properties**.
   b. Click the **General** tab.
   c. Clear the **CBR Enabled** checkbox and click **OK**.
   d. A dialog will ask if you wish to propagate this change to the subclasses of this class, click **Yes**.
   e. Repeat this procedure to disable CBR for all classes.
3. Run an index job and re-index any of the following that were previously enabled:
   - Document
   - Annotation
   - Custom Object
   - Folder

   This will disable all full-text indexing and content-based retrieval settings and will delete any associated collections. Once the indexing job is complete, proceed to step 4.

4. As *k2_os_user*, log on to each machine in your IBM FileNet Content Search Engine configuration, starting with the Master Administration Server, and remove the Autonomy K2 software installation as follows:

   Windows

   a. Access the Add or Remove Programs control panel and select **IBM FileNet Content Search Engine**.

   b. Click **Change/Remove**.

   UNIX

   a. Access opt/verity/appser/bin

   b. Stop the Autonomy K2 Administration Server service and the Tomcat application server. Use the following command, according to your environment:

      HP-UX

      ```
      /verity_install_directory/k2/_hpux/bin/k2adminstop
      ```

      AIX

      ```
      /verity_install_directory/k2/_hpux/bin/_rs6k43/bin/k2adminstop
      ```

      Solaris

      ```
      /verity_install_directory/k2/_hpux/bin/_ssol26/bin/k2adminstop
      ```

      Linux

      ```
      /verity_install_directory/k2/_hpux/bin/_ilnx21/bin/k2adminstop
      ```

   c. Change directory to:

      /opt/verity/_cseuninst

   d. Execute the following command:

      ```
      ./CSE_uninstaller
      ```

5. Delete the install directory.

Remove Process Engine (Windows)

This task includes Process Engine removal instructions for Windows platforms.

CAUTION You must remove the software in the order listed below. If you remove Image Services (Step 9) before you remove Process Engine (Step 3) the Process Engine software will be left in a state that will not allow removal with this procedure.

To remove the Process Engine software

1. Stop all of the following components that are running. For procedures and further details, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup.

Component	Server
Process Simulator	Process Simulator
Process Analyzer	Process Analyzer
Component Manager	Application Engine
Application Engine	Application Engine
Content Engine	Content Engine
Process Service	Process Engine
Process Task Manager	Process Engine or Application Engine

2. Navigate to Control Panel > Add/Remove Programs.
3. Click **Remove** for the Process Engine application.
4. Click **Next** at the Welcome screen for Process Engine uninstallation.
5. Click **Next** to stop FileNet BPM software components.
6. Click **Uninstall** to confirm you want to remove Process Engine for Windows 4.5.0 installation.
7. Indicate whether you want to reboot now or later.
8. Click **Finish** after you've read the summary information.
9. Click **Remove** for the FileNet Image Services 4.1.2 software.
10. Enter **Yes** to confirm you want to remove the Image Services software.
11. Press **Enter** to continue with the uninstall.
12. Close the Add/Remove snap-in.
13. Close the Control Panel.

CAUTION (SQL Server only) If you plan to reinstall Process Engine and will configure Process Engine to use a different SQL database, you must remove the database that was configured for the Process Engine installation. In addition, you must remove the following FileNet user IDs from the SQL Server Security folder before you reinstall Process Engine software:

- f_sw
- f_maint

Remove Process Engine (UNIX)

This task includes Process Engine removal instructions for UNIX platforms.

To remove the Process Engine software

1. Stop all of the following components that are running. For procedures and further details, see the IBM FileNet P8 help topic FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup.

Component	Server
Process Simulator	Process Simulator
Process Analyzer	Process Analyzer
Component Manager	Application Engine
Application Engine	Application Engine
Content Engine	Content Engine
Process Service	Process Engine
Process Task Manager	Process Engine or Application Engine

2. Enter the following after the FileNet software is shut down:

   ```
   killfnsw -A -D -y
   ```

3. Log on to the Process Engine server as the root user.
4. On AIX, execute the following:

   ```
   slibclean
   ```

5. If the /fnsw/_pws_uninst2 directory exists, run the following command:

   ```
   /fnsw/_pws_uninst2/pws_uninstall.bin
   ```

   Otherwise run:

   ```
   /fnsw/_pws_uninst/pws_uninstall.bin
   ```

6. Follow the prompts on the screen to remove the Process Engine software.
7. At a command prompt, run the following script:

   ```
   /fnsw/etc/uninstall
   ```

The uninstall script shuts down the Process Engine software, removes the software from the /fnsw directory, and removes the directory structure under /fnsw. The uninstall script removes the Process Engine entries in the /etc/services, /etc/system, /etc/inittab and /etc/devlink.tab files. On an AIX-based Process Engine, the script also removes smit entries.

Remove Application Engine (WebSphere)

This task includes Application Engine removal instructions for WebSphere on UNIX and Windows platforms.

To remove the Application Engine software

1. Log on to the application server.

   UNIX

   Log on as a user with write access to the /bin directory.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

2. Log in to the WebSphere administrative console.
3. Uninstall the Workplace application.
   a. Stop the Workplace process in the admin console.
   b. Uninstall the Workplace application from **Enterprise Applications**.
4. Navigate to the /_uninst folder under the Application Engine installation location.
5. Run the uninstall program:

   UNIX

   uninstaller.bin

   Windows

   uninstaller.exe

   NOTE For Windows, you can also use Add/Remove Programs from the Control Panel to remove the IBM FileNet Application Engine.

   On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

6. Delete the Workplace folder:

   WAS_HOME/temp/*node_name*/*application_server_name*/Workplace

7. Delete the *AE_install_path* directory.
8. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the vpd.properties file. If it exists, delete it.

   WARNING In the following step, do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

9. (UNIX) Remove the P8TASKMAN_HOME system environment variable.

   If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN_HOME system environment variable to complete the uninstallation.

Remove Application Engine (WebLogic)

This task includes Application Engine removal instructions for WebLogic on UNIX and Windows platforms.

To remove the Application Engine software

1. Log on to the application server.

   UNIX

   Log on as a user with write access to the /bin directory.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

2. Undeploy the Workplace application.

   a. Stop the Workplace Web Application Module.

   b. Delete the Workplace Web Application Module.

3. Navigate to the /_uninst folder under the Application Engine installation location.

4. Run the uninstall program:

   UNIX

   uninstaller.bin

   Windows

   uninstaller.exe

   NOTE For Windows, you can also use Add/Remove Programs from the Control Panel to remove Application Engine.

   On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

5. Delete the *AE_install_path* directory.

6. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the vpd.properties file. If it exists, delete it.

7. (UNIX) Remove the P8TASKMAN_HOME system environment variable.

   If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN_HOME system environment variable to complete the uninstallation.

   WARNING Do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

Remove Application Engine (JBoss)

This task includes Application Engine removal instructions for JBoss on Linux and Windows platforms.

To remove the Application Engine software

1. Log on to the application server:

   UNIX

   Log on as a user with write access to the /bin directory.

   Windows

   Log on as a member of the local Administrators group or a user with equivalent permissions.

2. Shut down JBoss.
3. Navigate to the /_uninst folder under the Application Engine installation location.
4. Run the uninstall program:

   UNIX

   uninstaller.bin

   Windows

   uninstaller.exe

   NOTE For Windows, you can also use Add/Remove Programs from the Control Panel to remove Application Engine.

   On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

5. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the vpd.properties file. If it exists, delete it.
6. (UNIX and LINUX) Remove the P8TASKMAN_HOME system environment variable.

   If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN_HOME system environment variable to complete the uninstallation.

   WARNING Do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

7. Remove temporary files and directories.
   a. Remove the Workplace.war folder from *JBoss_HOME*\server\default\deploy\.
   b. Remove the Temp working directory for Workplace from *JBoss_HOME*\server*server_name*\work\MainEngine\localhost\
8. Delete the *AE_install_path* directory.

Remove the Application Engine ISRA servlet

This task includes Application Engine ISRA Servlet removal instructions for Windows and UNIX environments.

NOTE Since the installed names for the ISRA Servlet are configurable on the supported application servers, the information below may not be the same as your environment. Make the appropriate name changes as required for your environment.

To remove the Application Engine Servlet software

1. Log on to the application server.

   UNIX

   Log on as a user with write access to the /bin directory and the directory where ISRA Servlet is installed.

   Windows

   Log on as a user with Administrative rights.

2. Undeploy the ApplicationEngineISRAServlet application. This step is similar to that required to undeploy the Workplace application.

   WebSphere

   a. Stop the ApplicationEngineISRAServlet process in the Admin console.

   b. Uninstall the ApplicationEngineISRAServlet application from Enterprise Applications.

   WebLogic

   a. Stop the ApplicationEngineISRAServlet Web Application Module.

   b. Undeploy or delete the ApplicationEngineISRAServlet Web Application Module.

   JBoss

   a. Shut down JBoss.

3. Navigate to the /_uninstISRA directory under the ISRA Servlet installation location.

4. Run the uninstall program:

   UNIX

   uninstall.bin

   Windows

   uninstall.exe

   NOTE For Windows, you can also use Add/Remove Programs from the Control Panel to remove the FileNet Application Engine ISRA Servlet.

5. Navigate to the /FileNet directory. If there is no other FileNet software installed under this directory, delete the /FileNet directory. If there is some other FileNet software installed under this directory, delete only the /ApplicationEngineISRAServlet subdirectory.

6. (WebSphere only) Delete the following temporary working folders for the Application Engine ISRA Servlet:

   `WAS_Home\WebSphere\AppServer\profiles\default\installedApps\servername\ApplicationEngineISRAServlet.ear\`

Appendixes

This appendix section contains the following major topics:

Configuration Manager reference

Configuration Manager is a tool for configuring and deploying new or upgraded instances of the Content Engine application on an application server. Configuration Manager generates a unique server profile that contains the properties and values needed for the associated configuration and deployment tasks. You must edit the property values to provide information that is specific to each configuration at your site, such as the application server name.

Configuration Manager has a graphical user interface (GUI) and a command-line interface (CLI). Both versions of the tool create the configuration XML files with the property values specific to your site, run tasks to apply your settings, display task status results, and deploy the Content Engine application. The GUI version of the tool displays the properties and default values that you need to set. When you save your changes in the GUI tool, the configuration XML files are updated for you. If you use the command-line version of the tool, you must first generate the configuration XML files with the tool, and then manually edit the default values in the files using a text editor. After you edit the files, you use Configuration Manager to execute the tasks to apply the saved settings. Once you have set the required values, you use either version of Configuration Manager to deploy the Content Engine application. Refer to "Installation and upgrade worksheet" on page 216 in *Plan and Prepare Your Environment for IBM FileNet P8* for information on gathering the values to supply to Configuration Manager.

> **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 you must specify for Configuration Manager:
>
> - Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select Configuration Manager.
> - Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."

The configuration XML files that you create with either version of the Configuration Manager tool can be used with the other version. For example, you can create the files with the command line version, and then use the GUI version to open the profile, edit the values, and run the configuration tasks. If needed, you can create multiple profiles, each of which supports a unique Content Engine instance. These instances can be located on the same server or on different servers, depending on your deployment preferences, the managed or non-managed nature of your application servers, and your clustering or high-availability requirements.

Use the Configuration Manager tool to perform these tasks:

- Set the application server properties. Content Engine will be deployed as an application on the application server. You need to specify the application server type, the software version number, the server name, the administrative user name and password, and other settings. The application server type determines some of the properties and their default values.
- Configure the Java Database Connectivity (JDBC) data sources. The JDBC data source information is used by Content Engine to connect to Global Configuration Data (GCD) and object store databases. The application server uses the JDBC data source information to connect Content Engine to the database. You need to specify the JDBC provider type, the database name, the database user name and password, and other settings. The JDBC provider type determines some of the properties and their default values.

- Configure the directory service (LDAP) provider. Content Engine connects to the directory service provider to authenticate users. The application server uses the directory service information to connect the Content Engine to the directory service provider. You need to specify the directory service provider type, the user and group naming conventions for your provider, the directory service user name for the Content Engine to use for authentication, and other settings. The LDAP provider type determines some of the properties and their default values.
- Configure the Content Engine bootstrap settings. The bootstrap information is needed for creating the GCD and for starting Content Engine.
- Deploy the Content Engine application. This action deploys the Content Engine EAR file with the JDBC, LDAP, and bootstrap settings on the application server.
- Check the status of a particular configuration task.

NOTE Once the Content Engine application is deployed, you use Enterprise Manager to customize Content Engine for your site's requirements.

- For information on how to work with passwords in Configuration Manager, see "Handling passwords in Configuration Manager" on page 417.
- For more information on using the GUI to configure and deploy Content Engine, see "Configuration Manager user interface reference" on page 418.
- For more information on the command line syntax and parameters, see "Configuration Manager command-line reference" on page 429.

Handling passwords in Configuration Manager

To provide the highest possible security, Configuration Manager's default settings do not save passwords from the GUI application. The password save setting is controlled by the cm.allowPasswordSave key in the application's config.ini file. While the default setting provides greater password security, it does require you to reenter all necessary passwords each time you start the GUI. When you close Configuration Manager, the passwords are removed from memory. When you open Configuration Manager again to run a saved task, you will need to reenter each of the following passwords to run the tasks:

- The application server administrator password. Select **File > Update Application Server Properties** to enter the password.
- The database administrator password. Edit the Configure JDBC Data Sources task.
- The directory service bind user password. Edit the Configure LDAP task.
- The bootstrap user password and the master key. Edit the Configure Bootstrap task.

The Configuration Manager command line passes the passwords from an XML configuration file to the required application when you execute a task, but does not save or change password entries in the XML configuration file. When you edit the XML configuration file, you can enter either an encrypted password or a plain text password. Plain text passwords are not recommended. For the procedure to encrypt passwords for the XML file, see "Encrypt passwords" on page 453.

If you later use the Configuration Manager GUI to open a profile with an XML configuration file that you manually edited, the GUI version reads the passwords in the XML file, and will overwrite the existing passwords when you save the file. If the GUI is not configured to save passwords (default setting), the passwords in the XML file will be overwritten with a blank entry. If the GUI is configured to save passwords, the original values or any changed values are encrypted and saved to the XML file.

To change the password save option for the user interface

1. If Configuration Manager is running, close the application.
2. Use a text editor to open *ce_install_path*/tools/configure/CMUI/configuration/config.ini.
3. Change the setting for cm.allowPasswordSave as follows:
   - To save passwords entered in the GUI, use `cm.allowPasswordSave=true`.
   - To not save passwords (default), use `cm.allowPasswordSave=false`.
4. Save the file.
5. Restart the Configuration Manager GUI.

Configuration Manager user interface reference

The Configuration Manager graphical user interface (GUI) lets you create, view and edit your Content Engine configuration profile. You can also make a copy of an existing profile, run configuration tasks, view the session logs, and check the status of a particular task.

NOTE If you need an accessible version of Configuration Manager, use the command line interface instead of the GUI. See "Configuration Manager command-line reference" on page 429.

Starting and stopping Configuration Manager

Configuration Manager is installed by the Content Engine installation program.

To start Configuration Manager

Complete the following procedure:

1. Navigate to the *ce_install_path*/tools/configure directory, where *ce_install_path* is the location where Content Engine server software is installed.
2. Execute the command for your operating system:

   UNIX

   cmui

   Windows

   cmui.exe

To stop Configuration Manager

If you have not saved your changes, you will be prompted to save them when you exit Configuration Manager.

- Select **File** > **Exit**.

Configuration Manager window

The initial Configuration Manager window consists of the Content Engine Task View pane on the left side, the Task Editor pane on the upper right side, and the Console pane in the lower right side. The Content Engine Task View pane displays a profile and the tasks for that profile. Only one profile can be open at a time. The Task Editor pane is empty until a specific task is selected from the Content Engine Task View pane. The Task Editor pane displays the properties and values for a selected task. Each open task is displayed in a separate tab in the Task Editor pane. More than one Task tab may be displayed at a time. The Console pane displays task execution messages, results from the Check Status command, or the session log.

If you close a pane, you can restore it by using the **Window > Show View** command.

Main toolbar

The Main toolbar is located just below the menu bar. The Main toolbar contains the following icons for working with your profiles:

Icon	Command name and Description
	Create a Configuration Profile Click this icon to create a new configuration profile. The current configuration profile will be closed, and the New Configuration Profile Wizard starts. If the existing open profile has been changed, you will be prompted to save your changes.
	Create Upgrade Configuration Profile Click this icon to create a configuration profile for upgrading an existing Content Engine server. The current configuration profile will be closed, and if there have been changes, you will be prompted to save your changes. Then New Configuration Profile Wizard starts.
	Open a Configuration Profile Click this icon to open an existing profile. The current configuration profile will be closed, and if there have been changes, you will be prompted to save your changes. Then New Configuration Profile Wizard starts.
	Update Application Server Properties Click this icon to view or edit the application server properties for the current profile.
	Run All Tasks Click this icon to run all the enabled tasks for the current profile. Click the down arrow to the right of this icon to select a single task to run.
	Save Click this icon to save the current configuration profile settings.
	View Configuration Manager Log File Click this icon to view the session log for the selected task. The session log is cleared when you open Configuration Manager.

Profile toolbar

The Profile toolbar is located at the upper right of the Content Engine Task View pane. The Profile toolbar contains the following icons for working with a selected profile:

Icon	Command name and Description
	Edit the Selected Task Click this icon to open the task in the Task Editor pane to view or edit the property values for that task.
	Run a Single Task Click this icon to run the selected task.
	Check the Status of the Selected Task Click this icon to display the task status in the Console pane. The Check Status results are the same as shown from the command line.

Console toolbar

The Console toolbar is located at the upper right of the Console pane. The Console toolbar contains the following icons for working with the Console pane:

Icon	Command name and Description
	Clear Console Click this icon to clear the display for the currently active tab in the Console pane. Clearing display does not affect any log contents.
	Scroll Lock Click this icon to enable or disable the scroll bars for the currently active tab in the Console pane. When the scroll bars are locked, information in the console might scroll out of view.
	Pin Console Click this icon to lock or unlock the current console location. When pinned (or locked), you cannot move the Console pane to a new location or resize the Console pane.
	Display Selected Console Click this icon to select the console tab to display. Select the desired tab from the list of recently viewed consoles.
	Open Console Click this icon to new tab with the current Console view. For example, you can open a second tab for execution messages for the Deploy Application task.

Menus and commands

Menu and command names	Description
File	Provides commands for creating, saving, or opening a configuration profile.
New Configuration Profile	Creates a new configuration profile.
Create Upgrade Configuration Profile	Creates a profile for upgrading Content Engine.
Open Configuration Profile	Opens an existing profile for viewing or editing. Keyboard shortcut: Ctrl+O
Close Configuration Profile	Closes the current configuration profile.
Save	Saves your changes to the active task. Keyboard shortcut: Ctrl+S
Save All	Saves your changes to all open tasks. Keyboard shortcut: Ctrl+Shift+S
Save Copy of Profile As	Saves the current profile with a new name or path.
Update Application Server Properties	Displays the application server settings for the current profile for you to edit. Keyboard shortcut: Ctrl+U
Run All	Runs all the tasks in the profile to apply your settings. If the Task Enabled setting is turned off for a particular task, that task is skipped when you select Run All Tasks. Keyboard shortcut: Ctrl+R
Exit	Closes Configuration Manager.
Window	Provides commands that give you access to logs and various views of information.
View Log File	Displays the session log in the Console pane. The session log lists results from tasks that you have run since you started Configuration Manager.
Show View	Provides a list of views available.Use to restore a view that you closed previously.
Help	Provides information about the Configuration Manager application.
About Configuration Manager	Opens a page that provides the copyright and related information about Configuration Manager.

Working with Configuration Manager

The following procedures explain how to use the Configuration Manager commands, icons, and panes to work with your profiles. For details on how to complete specific Content Engine configuration tasks, see "Configure Content Engine instances" on page 32 or "Configure Content Engine instances upgraded from 4.0.x" on page 233,

To create a new configuration profile

Complete the following procedure:

1. Create a new profile using one of these methods:
   - Click **Create a Configuration Profile** in the toolbar.
   - Select **File** > **New Configuration Profile**.
2. If a profile is already open, the Action Required message box opens. Respond to the messages as follows:
   a. Click **Yes** to continue creating a new profile, or click **No** to cancel.
   b. If you selected **Yes** and your current profile has any unsaved changes, the Save Resource message box opens. Click **Yes** to save your changes, click **No** to continue without saving your changes, or click **Cancel**.

   The Create New Configuration Profile Wizard opens.
3. Enter a name for the profile. The name must be a valid directory name for your operating system.
4. Specify the path for the profile. Either type in the full path to the profile directory or click **Browse** to locate the directory. The default path is *ce_install_path*/tools/configure/profiles, where *ce_install_path* is the location where Content Engine is installed.
5. Click **Next**.
6. Select the application server type for this profile.
7. Click **Next**.
8. Enter the property values for the application server. The specific properties displayed depend on the server type that you selected in Step 6.
9. Click **Next**.
10. Select the tasks that you want to include in this profile. For a new Content Engine installation, you need to complete all four configuration tasks: **Configure JDBC Data Sources**, **Configure LDAP**, **Configure Bootstrap**, and **Deploy Application**. To create data sources for an object store, you need to select **Configure JDBC Data Sources** only.
11. Click **Finish** to create the profile.

The profile you created will be displayed as an icon in the left-hand pane, along with three icons for the configure tasks, and one icon for the deploy task:

- Configure JDBC Data Sources
- Configure LDAP
- Configure Bootstrap Properties
- Deploy Application

To create a new profile for an upgrade

Use the following procedure to create a profile for upgrading Content Engine. You use the upgrade configuration profile to update the existing Content Engine bootstrap properties and deploy the updated EAR file. In order to create an upgrade profile, the Content Engine installation program must have detected an existing Content Engine installation.

1. Start the Create a Profile Wizard using one of these methods:
   - Click **Create a Configuration Profile for Upgrade** in the toolbar.
   - Select **File** > **Create Upgrade Configuration Profile**.
2. If a profile is already open, the Action Required message box opens. Respond to the messages as follows:
   a. Click **Yes** to continue creating a new profile, or click **No** to cancel.
   b. If you selected **Yes** and your current profile has any unsaved changes, the Save Resource message box opens. Click **Yes** to save your changes, click **No** to continue without saving your changes, or click **Cancel**.
3. In the Upgrade Profile Path screen, specify a name for the profile. Configuration Manager will create a directory with the name you specify.
4. Click **Next**.
5. Review the confirmation information, and click **Next**.
6. Enter the property values for the application server. The specific properties displayed depend on the server type that was selected when the existing Content Engine was installed.
7. Click **Finish** to create the profile.

   The profile you created is displayed as an icon in the left-hand pane with the name presented in the confirmation screen in Step 5, along with icons for **Upgrade Bootstrap** and **Deploy Application**.

To open an existing profile

You can have one profile open at a time in Configuration Manager.

1. Start the Open Configuration Profile Wizard using one of these methods:
   - Click **Open a Configuration Profile** in the toolbar.
   - Select **File** > **Open Configuration Profile**.
2. If a profile is already open, the Action Required message box opens. Respond to the messages as follows:
   a. Click **Yes** to continue creating a new profile, or click **No** to cancel.
   b. If you selected **Yes** and your profile has any unsaved changes, the Save Resource message box opens. Click **Yes** to save your changes, click **No** to continue without saving your changes, or click **Cancel**.
3. Either type in the full path to the profile directory, or click **Browse** to select the Profile directory. Select the desired profile from the Browse for Directory dialog box, and click **OK**.
4. Click **Finish**.

To open a task

More than one task tab can be open at a time in the Task Editor pane.

1. If the profile is collapsed in the Content Engine Task View pane, click **+** next to the profile name to expand it.
2. Use one of the following methods to open the desired task:
   - Click the **task name** in the Content Engine Task View pane, and then click **Edit Selected Task** in the Profile toolbar.
   - Double-click the **task name** in the Content Engine Task View pane. A new tab opens in the Task Editor pane for the selected task. You can then view and edit the values displayed in the Task Editor pane.

To switch between open tasks in the Task Editor pane

Complete the following procedure:

1. Click the **tab name** in the Task Editor pane for the desired task.

To run all the tasks in the profile

You can run each task individually or run all tasks. Use this procedure to run all tasks that have been enabled at once. The Task enabled check box must be selected for a specific task to run.

1. If the profile is collapsed in the Content Engine Task View pane, click **+** next to the profile name to expand it.

2. To run all the tasks in a profile at once, use one of the following methods to run the Run All Tasks command:
   - Click **Run All Tasks** in the Content Engine Task View pane toolbar.
   - Select **File > Run All**.

   The Console pane displays the task execution messages.

   NOTE Task execution will fail if you have not entered passwords correctly during the current Configuration Manager session. See "Handling passwords in Configuration Manager" on page 417 for more information.

To run a single task

You can run each task individually. Use this procedure to run a select task that has been enabled. The Task enabled check box must be selected for a specific task to run.

1. If the profile is collapsed in the Content Engine Task View pane, click **+** next to the profile name to expand it.
2. Use one of the following methods to run the Run a Task command:
   - Click the down arrow just to the right of the **Run All Tasks** icon in the Main toolbar, and then select the task that you want to run.
   - Click the **task name** in the Content Engine Task View pane to select the task. Click **Run a Task** in the Profile toolbar.
   - Right-click the **task name** in the Content Engine Task View pane, and select **Run Task** from the context menu.

   The Console pane displays the task execution messages.

   NOTE Task execution will fail if you have not entered passwords correctly during the current Configuration Manager session. See "Handling passwords in Configuration Manager" on page 417 for more information.

To check the status of a task

Complete the following procedure:

1. If the profile is collapsed in the Content Engine Task View pane, click **+** next to the profile name to expand it.
2. Use one of the following methods to run the Check Status command:
   - Click **Check Status** in the Content Engine Task View pane toolbar.
   - Right-click the **task name** in the Content Engine Task View pane, and select **Check Task Status** from the context menu.

3. The console pane opens with the status listed for the selected task.

   The following table lists the status results and their descriptions.

Status Result	Description
COMPLETED	The task ran successfully.
INCOMPLETE	The task is incomplete.

To view the session log

The session log contains information about the tasks that were run in the current session of Configuration Manager. As you run additional tasks for the same or a different profile, new messages are added to the log. When you exit Configuration Manager, the session log is cleared.

1. Run at least one task.
2. Use one of the following methods to run the View Log command:
   - Click **View Configuration Manager Log File** in the Main toolbar.
   - Select **Window > View Log**.

   The current session log is displayed in the Console pane.
3. Running a task will replace the displayed session log with the current task status. To redisplay the session log, repeat Step 2 as needed.

To save your changes to the task settings

Save changes to the currently open task using one of these methods:

- Click **Save** in the toolbar.
- Select **File** > **Save**.

To close a profile

You will be prompted to save any changes when you close a profile.

- Select **File** > **Close Configuration Profile**.

To restore a closed pane

You can restore a pane that has been closed.

1. Select **Window** > **Show View**.
2. Select one of the following commands:
   - Select **Console** to display the Console pane.
   - Select **Content Engine Task View** to display the Content Engine Task View pane.

cmui.ini parameters

When you install Configuration Manager, the path to the directory that contains the Java binary to be used to launch the graphical user interface is added to the *ce_install_path*/tools/configure/CMUI/cmui.ini file.

Windows example

```
-vm
C:\Program Files\FileNet\ContentEngine\_cejvm/bin
```

UNIX example

```
-vm
/opt/FileNet/ContentEngine/_cejvm/bin
```

Configuration Manager command-line reference

Configuration Manager can be run from a command line or from a graphical user interface. This section covers the command-line version of Configuration Manager.

Running Configuration Manager

How you invoke Configuration Manager from the command line depends on the operating system of the host machine where you run it:

Operating System	File Name
UNIX	`configmgr.sh`
Windows	`configmgr.bat`

NOTE In the following subtopics, substitute `configmgr.sh` or `configmgr.bat` in place of `configmgr`.

Configuration Manager syntax

Configuration Manager uses the following syntax patterns:

- To generate a configuration XML file:

  `configmgr generateConfig -appserver` *`app_server_type`* `-db` *`db_type`* `-ldap` *`ldap_type`* `-deploy` *`deploy_type`* `-task` *`task_name`* `[-path` *`mypath`*`] [-silent]`

- To execute a configuration task:

  `configmgr execute -task` *`task_name`* `[-path` *`mypath`*`] [-silent]`

- To check the status of a configuration task that you have executed:

  `configmgr checkStatus [ -task` *`taskname`* `] [-path` *`mypath`*`]`

- To display a description of the settings for a configuration task:

  `configmgr describe -task` *`taskname`*

The command names (generateConfig, execute, checkStatus, describe) are case sensitive; the parameter values (task_name, app_server_type, etc.) are not.

Command descriptions

The Configuration Manager commands are as follows:

checkStatus

The checkStatus command checks whether the results of executing a configuration XML file are correct.

execute

The execute command executes a configuration XML file. Typically, you will edit a configuration XML file (to insert your site-specific properties and values) before executing it.

describe

The describe command displays a description of what each task does.

generateConfig

The generateConfig command generates a configuration XML file, which depends on the task you specify in the command line.

Parameter descriptions

The Configuration Manager parameters are as follows:

-appserver *app_server_type*

The *appserver_type* value specifies the type of application server and must be one of the following: WebSphere, WebLogic, or JBoss.

-db *database_type*

The *database_type* value specifies the type of database to be used by Content Engine and must be one of the following: mssql, oracle, oracle_rac, db2, or db2zos. This parameter applies only to the generateConfig command.

-deploy *deploy_type*

The *deploy_type* value specifies the type of Content Engine deployment and applies only to the generateConfig command. The value must be one of the following: standard, cluster, or netdeploy (network deployment).

Specify standard if you are deploying Content Engine to a standalone (that is, a server that is neither managed nor clustered) WebSphere, WebLogic, or JBoss application server

Specify cluster if you are deploying Content Engine to a WebSphere, WebLogic, or JBoss application server cluster

Specify netdeploy if you are deploying Content Engine to a managed WebSphere application server instance (using Network Deployment for managing individual servers that are not necessarily in a cluster).

-ldap *ldap_type*

The *ldap_type* value specifies the type of directory service repository Content Engine uses for authenticating users and must be one of the following: activedirectory, adam, edirectory, sunjava, or tivoli. The adam option applies to both Microsoft ADAM and AD LDS. This parameter applies only to the generateConfig command.

-task *task_name*

task_name specifies which task to perform. Each task has one or more associated configuration XML files. The following table lists the tasks you can perform, the affected configuration XML files, and the contents of the files:

Task Name	Configuration XML File	Task Description or Contents of Configuration XML File
ConfigureApplicationServer	configureapplicationserver.xml	Parent configuration task that generates a complete set of configuration files. The parent XML file identifies the location of the following files: • configurebootstrap.xml • configurejdbc.xml • configureldap.xml • deployapplication.xml
ConfigureBootstrap	configurebootstrap.xml	Settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data (GCD) and for starting Content Engine.
ConfigureJDBC	configurejdbc.xml	Settings for configuring JDBC data source connections to the Global Configuration Data (GCD) and object store databases that Content Engine uses.
ConfigureLDAP	configureldap.xml	Settings for connecting to and searching within a directory server for authenticating Content Engine users.
ConfigureDotNetAPI	none	Configures the Content Engine .Net API for client use.
ConfigureDotNetClient	none	Configures the Content Engine .NET API client and FileNet Enterprise Manager.

Task Name	Configuration XML File	Task Description or Contents of Configuration XML File
DeployApplication	deployapplication.xml	Settings for deploying the Content Engine instance on the application server.
RedeployApplication	redeployapplication.xml	Settings for deploying an updated Content Engine EAR file
RegisterEM	none	Settings for registering or unregistering FileNet Enterprise Manager from the Windows Registry.
Upgrade	upgrade.xml	Parent configuration that identifies the location of the RedeployApplication, UpgradeApplicationServer, UpgradeBootstrap configuration XML files
UpgradeApplicationServer	upgradeapplicationserver.xml	Settings for upgrading Global Configuration Data (GCD) and starting Content Engine.
UpgradeBootstrap	upgradebootstrap.xml	Settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data (GCD) and for starting Content Engine.

NOTE The *task_name* parameter is not case sensitive. For example, you can enter ConfigureApplicationServer, Configureapplicationserver, or ConFiGureApplicationServer for the *task_name*.

-path *mypath*

The -path *mypath* entry is optional and specifies the path to the configuration XML files that the Configuration Manager generates, executes, or verifies, as follows:

- If *mypath* is an absolute path name (that is, a path that begins at the root directory), Configuration Manager uses *mypath*.
- If *mypath* is a relative path (that is, a path that does not begin at the root directory) or just a single directory name, Configuration Manager uses the following path:

  CE_install_path/tools/configure/profiles/*mypath*

- If you omit -path *mypath*, Configuration Manager uses the following path:

  CE_install_path/tools/configure/profiles

NOTE If you are deploying multiple Content Engine instances on the same machine, you will need to generate, edit, and deploy a complete set of configuration files for each instance. Store the configuration files for each instance in a separate directory.

-silent

The -silent parameter is optional. When -silent is specified, then no prompts or informational messages are displayed. Failure messages are displayed as needed. If the configuration XML file to be generated already exists, then the task fails if the silent switch is selected.

New Content Engine classes and properties

This appendix lists the new classes and properties that have been added to Content Engine in the 4.0.0 and 4.5.0 release. If you have previously defined any classes or properties in a 3.5.*x* or 4.0.x object store with names that conflict with any of the class or property names listed here, it will not be possible to upgrade the object store to 4.5.0. The upgrade tool will check for naming conflicts and require you to change the names. To complete the upgrade, you must change the name of your user-defined class or property.

Consider the following upgrade scenarios:

- If you are upgrading from 3.5.x to 4.5.x, you must resolve conflicts between your 3.5.x names and the 4.0.0 names, as well as any conflicts between 3.5.x and 4.5.0 names. Consult the lists in both of the following sections:
  - "New in Content Engine 4.0.0 (for upgrades from 3.5.x only)" on page 435
  - "New in Content Engine 4.5.0 (for upgrades from 3.5.x and from 4.0.x)" on page 448
- If you are upgrading from 4.0.x to 4.5.0, you must resolve conflicts between your 4.0.0 names and the 4.5.0 names. See the following section.
  - "New in Content Engine 4.5.0 (for upgrades from 3.5.x and from 4.0.x)" on page 448

New in Content Engine 4.0.0 (for upgrades from 3.5.x only)

New Content Engine classes in 4.0.0

The following table displays the new class names in Content Engine 4.0.0. To ensure successful upgrade, resolve any conflicts between your 3.x object store class names and the new names.

Class Symbolic Name	Class ID
AddOnInstallationRecord	Add On Installation Record
AsyncProcessingConfiguration	Async Processing Configuration
CenteraSiteSettings	Centera Site Settings
CodeModule	Code Module
ComponentRelationship	Component Relationship
ContentCacheArea	Content Cache Area
ContentCacheConfiguration	Content Cache Configuration
ContentConfiguration	Content Configuration
DatabaseStorageArea	Database Storage Area
DirectoryConfiguration	Directory Configuration
DirectoryConfigurationAD	Directory Configuration AD
DirectoryConfigurationADAM	Directory Configuration ADAM
DirectoryConfigurationIBM	Directory Configuration IBM
DirectoryConfigurationNovell	Directory Configuration Novell
DirectoryConfigurationSunOne	Directory Configuration SunOne
Domain	Domain
EntireNetwork	EntireNetwork
FileStorageArea	File Storage Area
FixedStorageArea	Fixed Storage Area
Group	Group
ImageServicesClassDescription	Image Services Class Description
ImageServicesImportAgentConfiguration	Image Services Import Agent Configuration

Class Symbolic Name	Class ID
ImageServicesPropertyDescription	Image Services Property Description
ImageServicesSiteSettings	Image Services Site Settings
IndexArea	Index Area
IndexJob	Index Job
IndexJobClassItem	Index Job Class Item
IndexJobCollectionItem	Index Job Collection Item
IndexJobItem	Index Job Item
IndexJobSingleItem	Index Job Single Item
IsolatedRegion	Isolated Region
PEConnectionPoint	PE Connection Point
PublishCompleteEvent	Publish Complete Event
PublishRequest	Publish Request
PublishRequestEvent	Publish Request Event
PublishingConfiguration	Publishing Configuration
Realm	Realm
RenditionEngineConnection	Rendition Engine Connection
SecurityPrincipal	Security Principal
ServerCacheConfiguration	Server Cache Configuration
ServerInstance	Server Instance
Site	Site
SiteSettings	Site Settings
StorageArea	Storage Area
StoragePolicy	Storage Policy
SubsystemConfiguration	Subsystem Configuration
TakeFederatedOwnershipEvent	Take Federated Ownership Event
TraceLoggingConfiguration	Trace Logging Configuration
UpgradeAddOn	Upgrade Add On
User	User

Class Symbolic Name	Class ID
VerityCollection	Verity Collection
VerityDomainConfiguration	Verity Domain Configuration
VerityIndexArea	Verity Index Area
VerityServerConfiguration	Verity Server Configuration
VirtualServer	Virtual Server

New Content Engine properties in 4.0.0

The following table displays the new property names in Content Engine 4.0.0. To ensure successful upgrade, resolve any conflicts between your 3.x object store property names and the new names.

Property Symbolic Name	Property Name
AbandonedContentCleanupInterval	Abandoned Content Cleanup Interval
AbandonedDBContentCleanupInterval	Abandoned Database Content Cleanup Interval
ActiveDirectorySiteDNS	Active Directory Site DNS
AddOn	Add On
AddOnInstallationRecords	Add On Installation Records
AddOns	Add Ons
AddOnName	Add On Name
AllRealms	All Realms
AllowFederatedDeletes	Allow Federated Deletes
AllowsContentToBeCached	Allows Content To Be Cached
AllowsDelete	Allows Delete
APITraceFlags	API Trace Flags
AppenderNames	Appender Names
ApplyDefinition	Apply Definition
AsynchronousProcessingTraceFlags	Asynchronous Processing Trace Flags
AuditedDeletePrefix	Audited Delete Prefix
BaseClassIDs	Base Class IDs
BatchDelay	Batch Delay

Property Symbolic Name	Property Name
BlobReadAheadSize	Blob Read Ahead Size
BlobWriteCollisionsAvoidanceFlag	Blob Write Collisions Avoidance Flag
BrokerPort	Broker Port
FPPoolBufferSize	Buffer Size
CacheStatus	Cache Status
CanAcceptForwardedRequests	Can Accept Forwarded Requests
CanForwardRequests	Can Forward Requests
CBRLocale	CBR Locale
CBRTraceFlags	CBR Trace Flags
CFSDaemonTraceFlags	CFS Daemon Trace Flags
ChildComponent	Child Component
ChildDocuments	Child Documents
ChildRelationships	Child Relationships
ChildVersionSeries	Child Version Series
ClassDefinition	Class Definition
ClosureDate	Closure Date
FPPoolClusterNonAvailTime	Cluster Nonavailable Time
CodeModule	Code Module
CodeModuleCacheEntryTTL	Code Module Cache Entry TTL
CodeModuleCacheMaxFileSpace	Code Module Cache Maximum File Space
CodeModuleCacheMaxMemory	Code Module Cache Maximum Memory
CodeModuleTraceFlags	Code Module Trace Flags
CollectionName	CollectionName
ComponentCascadeDelete	Component Cascade Delete
ComponentPreventDelete	Component Prevent Delete
ComponentRelationshipType	Component Relationship Type
ComponentSortOrder	Component Sort Order
CompoundDocumentState	Compound Document State

Property Symbolic Name	Property Name
ConcurrentReaders	Concurrent Readers
ConnectionState	Connection State
ConnectionTimeout	Connection Timeout
ContentCacheArea	Content Cache Area
ContentCacheAreas	Content Cache Areas
ContentCacheTraceFlags	Content Cache Trace Flags
ContentElementCount	Content Element Count
ContentElementKBytes	Content Element Kbytes
ContentElementsCreated	Content Elements Created
ContentElementsDeleted	Content Elements Deleted
ContentQueueMaxWorkerThreads	Content Queue Max Worker Threads
ContentStorageTraceFlags	Content Storage Trace Flags
ContentTempDirectoryRoot	Content Temp Directory Root
CSMCache	CSM Cache
CurrentUser	Current User
DatabaseContentUploadBufferSize	Database Content Upload Buffer Size
DatabaseServerPort	Database Server Port
DatabaseTraceFlags	Database Trace Flags
DatabaseType	Database Type
FPPoolDefaultCollisionAvoidance	Default Collision Avoidance
DefaultISDocumentClass	Default IS Document Class
DefaultRetentionDays	Default Retention Days
DefaultRetentionPassThrough	Default Retention Pass Through
DefaultSite	Default Site
DeleteMethod	Delete Method
DeviceAddress	Device Address
DeviceRootDirectory	Device Root Directory
DirectoryConfigurations	Directory Configurations

Property Symbolic Name	Property Name
DirectoryServerHost	Directory Server Host
DirectoryServerPassword	Directory Server Password
DirectoryServerPort	Directory Server Port
DirectoryServerProviderClass	Directory Server Provider Class
DirectoryServerType	Directory Server Type
DirectoryServerUserName	Directory Server User Name
DirectoryStructure	Directory Structure
DispatcherEnabled	Dispatcher Enabled
DispatcherWaitInterval	Dispatcher Wait Interval
DistinguishedName	Distinguished Name
DNSName	DNS Name
DocumentsPerBatch	Documents Per Batch
Domain	Domain
DomainID	Domain ID
EJBForwardingEndpoint	EJB Forwarding Endpoint
EJBTraceFlags	EJB Trace Flags
EmailAddress	Email Address
EmbeddedDataThreshold	Embedded Data Threshold
FPPoolEnableMulticlusterFailover	Enable Multicluster Failover
EncryptionAlgorithm	Encryption Algorithm
EngineTraceFlags	Engine Trace Flags
ErrorCode	Error Code
ErrorDescription	Error Description
EventsTraceFlags	Events Trace Flags
ExpiredBatchSelectionSize	Expired Batch Selection Size
FixedContentDevice	Fixed Content Device
FixedContentDevices	Fixed Content Devices
FixedContentProviderTraceFlags	Fixed Content Provider Trace Flags

Property Symbolic Name	Property Name
FolderCacheMaxAgeDelta	Folder Cache Maximum Age Delta
FolderCacheMaxEntries	Folder Cache Maximum Entries
FolderCacheReapFrequency	Folder Cache Reap Frequency
FromVersions	From Versions
FullTextRowDefault	Full Text Row Default
FullTextRowMax	Full Text Row Max
GCDCacheTTL	GCD Cache TTL
GCDTraceFlags	GCD Trace Flags
GroupBaseDN	Group Base DN
GroupDisplayNameAttribute	Group Display Name Attribute
GroupMembershipSearchFilter	Group Membership Search Filter
GroupNameAttribute	Group Name Attribute
GroupSearchFilter	Group Search Filter
Groups	Groups
ImageServicesClassDescriptions	Image Services Class Descriptions
ImageServicesClassName	Image Services Class Name
ImageServicesDataType	Image Services Data Type
ImageServicesPropertyDescriptions	Image Services Property Descriptions
ImageServicesPropertyName	Image Services Property Name
ImplementationClass	Implementation Class
ImportAgentRetryCount	Import Agent Retry Count
InboundFileNameCacheMaxEntries	Inbound File Name Cache Maximum Entries
IndexArea	Index Area
IndexAreas	Index Areas
IndexItems	Index Items
IndexationID	Indexation ID
InitiatingUser	Initiating User
InlineContentRetrievalLimit	Inline Content Retrieval Limit

Property Symbolic Name	Property Name
InputDocument	Input Document
InstallationDate	Installation Date
InstallationReport	Installation Report
Installer	Installer
IsCBREnabled	Is CBR Enabled
ISDomainName	IS Domain Name
IsFederatedOnly	Is Federated Only
ISOrganization	IS Organization
ISPassword	IS Password
IsSSLEnabled	Is SSL Enabled
ISTempDir	IS Temporary Directory
ISUserName	IS User Name
IsolatedRegion	Isolated Region
IsolatedRegions	Isolated Regions
JNDIDataSource	JNDI DataSource
JNDIXADataSource	JNDI XA DataSource
JobAbortRequested	Job AbortRequested
JobStatus	Job Status
LabelBindValue	Label Bind Value
LeaseDuration	Lease Duration
LocalDomain	Local Domain
MarkingSetCacheEntryTTL	Marking Set Cache Entry TTL
MarkingSetCacheMaxEntries	Marking Set Cache Maximum Entries
MarkingSets	Marking Sets
MarkingUseGranted	Marking Use Granted
MaxBatchSize	Max Batch Size
MaxCollections	Max Collections
MaxInMemoryElementState	Max In Memory Element State

Property Symbolic Name	Property Name
MaxInMemoryQueueItems	Max In Memory Queue Items
MaxObjectsPerCollection	Max Objects Per Collection
MaxReaderSemaphoreWaitTime	Max Reader Semaphore Wait Time
FPPoolMaxConnections	Maximum Connections
MaximumContentElements	Maximum Content Elements
MaxResolutionBatchSize	Maximum Resolution Batch Size
MaximumRetentionDays	Maximum Retention Days
MaximumSizeKBytes	Maximum Size Kbytes
MaximumTimeToLive	Maximum Time To Live
MemberOfGroups	Member Of Groups
MetadataTraceFlags	Metadata Trace Flags
MinimumRetentionDays	Minimum Retention Days
FPPoolMulticlusterDeleteStrategy	Multicluster Delete Strategy
FPPoolMulticlusterExistsStrategy	Multicluster Exists Strategy
FPPoolMulticlusterQueryStrategy	Multicluster Query Strategy
FPPoolMulticlusterReadStrategy	Multicluster Read Strategy
FPPoolMulticlusterWriteStrategy	Multicluster Write Strategy
MyRealm	My Realm
NeverDeleteClipsOrContent	Never Delete Clips Or Content
NoWorkDelay	No Work Delay
NumImportAgents	Num Import Agents
ObjectSecurityCacheEntryTTL	Object Security Cache Entry TTL
ObjectSecurityCacheMaxEntries	Object Security Cache Maximum Entries
ObjectStores	Object Stores
FPOpenStrategy	Open Strategy
OptimizationInterval	Optimization Interval
OriginalObject	Original Object
OriginalOrdinal	Original Ordinal

Property Symbolic Name	Property Name
OutputFolder	Output Folder
OutputLocation	Output Location
ParentComponent	Parent Component
ParentDocuments	Parent Documents
ParentRelationships	Parent Relationships
PartialResolutionChunkSize	PartialResolutionChunkSize
PEConnectionPoints	PE Connection Points
PersistenceType	Persistence Type
PoolAddress	Pool Address
FPPoolPrefetchSize	Prefetch Size
PreloadOnCreate	Preload On Create
Prerequisites	Prerequisites
PrivilegedSettability	Privileged Settability
PruneAmount	Prune Amount
PruneThresholdContentElements	Prune Threshold Content Elements
PruneThresholdSizeKBytes	Prune Threshold Size KBytes
PublicKey	Public Key
PublicationDocument	Publication Document
PublishRequestType	Publish Request Type
PublishStyleTemplate	Publish Style Template
PublishTemplate	Publish Template
PublishTraceFlags	Publish Trace Flags
PublishingStatus	Publishing Status
QueueItemDatabaseTimeout	Queue Item Database Timeout
QueueItemMaxDispatchers	Queue Item Max Dispatchers
QueueItemRetryCount	Queue Item Retry Count
ReferencingActions	Referencing Actions
RegionPassword	Region Password

Property Symbolic Name	Property Name
RenameFileRetryAttempts	Rename File Retry Attempts
RenditionEngineConnection	Rendition Engine Connection
RenditionEngineConnections	Rendition Engine Connections
RenditionEnginePassword	Rendition Engine Password
ResourceStatus	Resource Status
ResourceString	Resource String
RetentionPeriod	Retention Period
RetrievalRetryAttempts	Retrieval Retry Attempts
FPPoolRetryCount	Retry Count
FPPoolRetrySleep	Retry Sleep
ReturnNameAsDN	Return Name As DN
RollFwdBatchRetryAttempts	Roll Forward Batch Retry Attempts
RootDirectoryPath	Root Directory Path
SearchCrossForestGroupMembership	Search Cross Forest Group Membership
SearchServersToAttach	Search Servers To Attach
SearchTraceFlags	Search Trace Flags
SecurityDescCacheMaxEntries	Security Descriptor Cache Maximum Entries
SecurityTraceFlags	Security Trace Flags
ServerInstances	Server Instances
ShortName	Short Name
SingleItem	Single Item
Site	Site
SiteSettings	Site Settings
Sites	Sites
SnapLockPassword	SnapLock Password
SnapLockUserName	SnapLock User Name
SSITraceFlags	SSI Trace Flags
StagingAreaPath	Staging Area Path

Property Symbolic Name	Property Name
Status	Status
StatusDescription	Status Description
StorageArea	Storage Area
StorageAreas	Storage Areas
SubjectCacheEntryTTL	Subject Cache Entry TTL
SubjectCacheMaxEntries	Subject Cache Maximum Entries
SubmittedCount	Submitted Count
SubsystemConfigurations	Subsystem Configurations
TemplateType	Template Type
TempDBContentLifetime	Temporary Database Content Lifetime
TempDirectoryPath	Temp Directory Path
TempFileLifetime	Temporary File Lifetime
ThreadCount	Thread Count
TimeAllSubmitted	Time All Submitted
TimeLastProcessed	Time Last Processed
FPPoolTimeout	Timeout
TimeOutSeconds	Timeout Seconds
ToVersions	To Versions
TraceLoggingEnabled	Trace Logging Enabled
URIValue	URI Value
UserBaseDN	User Base DN
UserDisplayNameAttribute	User Display Name Attribute
UserDomain	User Domain
UserGroup	User Group
UserName	User Name
UserNameAttribute	User Name Attribute
UserPassword	User Password
UserSearchFilter	User Search Filter

Property Symbolic Name	Property Name
UserTokenCacheEntryTTL	User Token Cache Entry TTL
UserTokenCacheMaxEntries	User Token Cache Maximum Entries
Users	Users
VerityAdminServerHostname	Verity Admin Server Hostname
VerityAdminServerPort	Verity Admin Server Port
VerityBrokerName	Verity Broker Name
VerityBrokers	Verity Brokers
VerityCollections	Verity Collections
VerityDomainConfiguration	Verity Domain Configuration
VerityIndexServers	Verity Index Servers
VerityMasterAdminServerHostname	Verity Master Admin Server Hostname
VerityMasterAdminServerPort	Verity Master Admin Server Port
VeritySearchServers	Verity Search Servers
VersionBindType	Version Bind Type
VirtualServer	Virtual Server
VirtualServers	Virtual Servers
WSITraceFlags	WSI Trace Flags

New in Content Engine 4.5.0 (for upgrades from 3.5.x and from 4.0.x)

New Content Engine classes in 4.5.0

The following table displays the new class names in Content Engine 4.5.0. To ensure successful upgrade, resolve any conflicts between your 3.5.x or 4.0.x object store class names and the new names.

Class Symbolic Name	Class ID
ReplicationGroup	Replication Group
CMODApplicationGroup	CMOD Application Group
CMODFixedContentDevice	CMOD Fixed Content Device
CMODRepository	CMOD Repository
ContentFederatingRepository	Content Federating Repository
DirectoryConfigurationCA	Directory Configuration CA
DITARenditionEngineConnection	DITA Rendition Engine Connection
ExternalAlias	External Alias
ExternalClassAlias	External Class Alias
ExternalClassDescription	External Class Description
ExternalIdentity	External Identity
ExternalParticipant	External Participant
ExternalPropertyAlias	External Property Alias
ExternalPropertyDescription	External Property Description
ExternalRepository	External Repository
ImageServicesRepository	Image Services Repository
ObjectStoreParticipant	Object Store Participant
Replicable	Replicable

Class Symbolic Name	Class ID
ReplicableClassDefinition	Replicable Class Definition
ReplicationConfiguration	Replication Configuration
ReplicationJournalEntry	Replication Journal Entry
ReplicationParticipant	Replication Participant
Repository	Repository

New Content Engine properties in 4.5.0

The following table displays the new property names in Content Engine 4.5.0. To ensure successful upgrade, resolve any conflicts between your 3.5.x or 4.0.x object store property names and the new names.

Property Symbolic Name	Property Name
AliasDirection	Alias Direction
ApplicationGroupNumber	Application Group Number
ClassFamily	Class Family
ClassIdentity	Class Identity
CMODApplicationGroups	CMOD Application Groups
CMODLanguage	CMOD Language
CMODPassword	CMOD Password
CMODPort	CMOD Port
CMODServerName	CMOD Server Name
CMODTraceLevel	CMOD Trace Level
CMODUserName	CMOD User Name
DatabaseIndexStorageLocation	Database Index Storage Location
DatabaseLOBStorageLocation	Database LOB Storage Location
DatabaseTableStorageLocation	Database Table Storage Location
DateCheckedIn	Date Checked In
DestinationRepository	Destination Repository
DITADatabaseTimeout	DITA Database Timeout
DITAHome	DITA Home

Property Symbolic Name	Property Name
DITAJVMArguments	DITA JVM Arguments
DITARenditionEngineConnection	DITA Rendition Engine Connection
DITARenditionEngineConnections	DITA Rendition Engine Connections
DITAWorkingDirectory	DITA Working Directory
DuplicateSuppressionEnabled	Duplicate Suppression Enabled
DynamicGroupMemberAttribute	Dynamic Group Member Attribute
DynamicGroupObjectClass	Dynamic Group Object Class
DynamicGroupQueryAttribute	Dynamic Group Query Attribute
ExternalAliases	External Aliases
ExternalClassDescriptions	External Class Descriptions
ExternalObjectIdentity	External Object Identity
ExternalPropertyDescriptions	External Property Descriptions
ExternalReplicaIdentities	External Replica Identities
ExternalRepositories	External Repositories
ExternalRepository	External Repository
FcpPoolIdleTimeoutSeconds	FCP Pool Idle Timout Seconds
FcpPoolMaxInUse	FCP Pool Max In Use
FcpPoolMaxWaitSeconds	FCP Pool Max Wait Seconds
FcpPoolPreferredSize	FCP Pool Preferred Size
FederateInboundContent	Federate Inbound Content
GroupUniqueIDAttribute	Group Unique ID Attribute
GroupUniqueIDAttributeIsBinary	Group Unique ID Attribute Is Binary
HeldUntilDate	Held Until Date
IICEConfigurationServerURL	Content Integrator URL
InboundBatchSize	Inbound Batch Size
InboundBusyWaitInterval	Inbound Busy Wait tInterval
InboundDispatcherEnabled	Inbound Dispatcher Enabled
InboundIdleWaitInterval	Inbound Idle Wait tInterval

Property Symbolic Name	Property Name
LastFailureReason	Last Failure Reason
MappableDirection	Mappable Direction
MaxInboundWorkers	Max Inbound Workers
MaxOutboundWorkers	Max Outbound Workers
ObjectStateRecordingLevel	Object State Recording Level
ObjectStoreSchemaDB2	Object Store Schema DB2
ObjectStoreSchemaMSSQL	Object Store Schema MSSQL
ObjectStoreSchemaOracle	Object Store Schema Oracle
OutboundBatchSize	Outbound Batch Size
OutboundBusyWaitInterval	Outbound Busy Wait tInterval
OutboundDispatcherEnabled	Outbound Dispatcher Enabled
OutboundIdleWaitInterval	Outbound Idle Wait tInterval
PartnerObjectStore	Partner Object Store
PropertyIdentity	Property Identity
RegionKey	Region Key
ReplicableDirection	Replicable Direction
ReplicationData	Replication Data
ReplicationGroup	Replication Group
ReplicationGroups	Replication Groups
ReplicationMode	Replication Mode
ReplicationOperation	Replication Operation
ReplicationParticipants	Replication Participants
ReplicationStatus	Replication Status
ReplicationTraceFlags	Replication Trace Flags
RestrictMembershipToConfiguredRealms	Restrict Membership To Configured Realms
SearchDynamicGroup	Search Dynamic Group

Property Symbolic Name	Property Name
SuperclassIdentity	Superclass Identity
SupersededAddOnIds	Superseded AddOn Ids
TargetObjectStore	Target Object Store
UserUniqueIDAttribute	User Unique ID Attribute
UserUniqueIDAttributeIsBinary	User Unique ID Attribute Is Binary

Encrypt passwords

Several passwords are required to configure IBM FileNet P8 components. To accommodate your security requirements, you can encrypt these passwords, as follows, before you enter them into the resource files.

To encrypt a password for Configuration Manager

Configuration Manager and the encryption tool run only on Windows.

1. Navigate to *ce_install_path*/tools/configure, where *ce_install_path* is where Content Engine is installed.
2. Run one of the following commands to generate the encrypted password:

   Windows

   ```
   encryptpassword.bat passwordtoencrypt
   ```

   UNIX

   ```
   encryptpassword.sh passwordtoencrypt
   ```

   where `passwordtoencrypt` is the password you want to encrypt.
3. An encrypted password is generated and displayed. Copy the encrypted password and paste into the appropriate property in the XML configuration file.

To encrypt a password for Process Engine

1. The encryption tool is located on the installation media for Process Engine software, in the Tools directory. Copy the following files from the Tools directory to a local drive:
   - fnencryptutils - an executable .jar file
   - RunEncryptApp - a batch file
2. Run one of the following executable files to invoke the application:

   Windows

   RunEncryptApp.bat

   UNIX

   RunEncryptApp.sh

   NOTE Before running the file, be sure Java is installed and its location is in your PATH environment variable.
3. Enter the appropriate values for the user name and password.
4. Reenter the password to confirm it.
5. Click **Generate**.

6. An encrypted password will be generated and displayed in the encrypted password field.
7. Copy and paste this password into the appropriate configuration file.

Index

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

GC19-2553-00