IBM FileNet P8 Platform
Version 4.5.1

Upgrade Guide
for installation on AIX with Oracle, IBM
WebSphere Application Server, and
Windows Active Directory

IBM ®

IBM FileNet P8 Platform
Version 4.5.1

Upgrade Guide
for installation on AIX with Oracle, IBM
WebSphere Application Server, and
Windows Active Directory

IBM ®

Contents

ibm.com and related resources

Product support and documentation are available from ibm.com.

Support and assistance

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

FileNet Content Manager Support
> http://www.ibm.com/software/data/content-management/filenet-content-manager/support.html

Information center

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

PDF publications

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

See the following PDF publications Web sites:

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

"How to send your comments"
Your feedback is important in helping to provide the most accurate and highest quality information.

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

How to send your comments

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

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

Consumability survey

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

Contacting IBM

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

To learn about available service options, call one of the following numbers:
* In the United States: 1-888-426-4343
* In Canada: 1-800-465-9600

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

Part 1. Upgrading and configuring IBM FileNet P8 Platform

Upgrading the IBM FileNet P8 Platform software can require that you make installation and configuration changes to one or more of the core application components and the common documentation, and the upgrade process can vary depending on the product version you are currently running.

Important: You must upgrade the core components before you upgrade any expansion products or optional add-on components.

"Upgrading IBM FileNet P8 documentation" on page 3
You can point your IBM FileNet P8 components to the information center on ibm.com. You can also choose to upgrade your existing IBM FileNet P8 documentation server by installing the new files and redeploying the documentation on the Web application server.

"Upgrading and configuring Content Engine and Content Search Engine" on page 9
You can upgrade Content Engine and Content Search Engine to version 4.5.1 from version 3.5.2, 4.0.1, or 4.5.0.

"Upgrading and configuring Process Engine" on page 131
Upgrade the Process Engine by performing several pre-upgrade steps, installing the Process Engine software and performing multiple post-upgrade steps.

"Upgrading and configuring Application Engine" on page 161
You can upgrade Application Engine from version 3.5 or version 4.0 to version 4.0.2. The procedures for upgrading Application Engine vary depending on the version that you are upgrading.

Upgrading IBM FileNet P8 documentation

You can point your IBM FileNet P8 components to the information center on ibm.com. You can also choose to upgrade your existing IBM FileNet P8 documentation server by installing the new files and redeploying the documentation on the Web application server.

You can link your FileNet P8 software directly to the IBM FileNet P8 Platform Version 4.5.1 information center on www.ibm.com. The advantages of linking to the ibm.com information center are as follows:

- You do not have to maintain a documentation server or information center.
- The documentation is always up to date.

If you want to use the online information center, you do not need to install or upgrade any documentation in your environment. However, you must configure your IBM FileNet P8 Platform components to use the new P8 documentation URL:

http://publib.boulder.ibm.com/infocenter/p8docs/v4r5m1/topic/com.ibm.p8.doc

To change the documentation URL for your FileNet P8 components, see the following help topics:

- **IBM FileNet P8 documentation** › **System Administration** › **Content Engine Administration** › **FileNet P8 domain** › **How to...** › **Configuring the FileNet P8 Platform documentation location**
- **IBM FileNet P8 documentation** › **System Administration** › **Process Administration** › **System administration tasks** › **Configure optional features** › **Reconfigure the FileNet documentation server URL**

You can also choose to upgrade your existing documentation package installation. You must also rebuild the documentation search index if you upgrade an installation containing more than the core IBM FileNet P8 Platform documentation.

Consider the following details as you get ready to upgrade your documentation:

- As of the IBM FileNet P8 4.0 release, you can deploy the IBM FileNet P8 documentation on a Content Engine server. Like Application Engine, Content Engine 4.0 and later runs on a web application server. If you are an existing 3.5 customer and want to install the 4.5.1 documentation onto a Content Engine server or a new application server, you must do a new installation rather than an upgrade.
- Before you update any IBM FileNet P8 documentation, review the *IBM FileNet P8 Hardware and Software Requirements* for the required software versions and updates for vendor software.
- Before you install 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 allows 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, be aware of the following requirements:
  – You must copy the associated documentation for all expansion products onto the updated IBM FileNet P8 documentation server.

– You must update the index for the documentation Search feature as documented. This action ensures that searches return all expected results. If you add documentation for other expansion products later, you must rerun the procedure for updating the documentation Search index.

– You must also update the documentation Search index if you update the installed IBM FileNet P8 documentation 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 is automatically copied to the `Index/IndexOld` subdirectory. If you need to return to your previous indexed state, you can reapply these backed-up files to the `core` subdirectory. However, you must first remove any newer existing files at that location.

- Depending on your operating system and application server levels, your options might vary slightly from what is documented.

To upgrade the documentation requires that you copy new files to the documentation server and reindex those files as necessary. You can then redeploy, start, and verify the upgraded documentation web application.

"Refreshing or updating IBM FileNet P8 documentation"
You can choose one of two methods to upgrade an IBM FileNet P8 documentation installation. The method you choose depends on whether the existing installation includes expansion product and other documentation in addition to core IBM FileNet P8 Platform documentation.

"Updating the documentation search index" on page 7
You must update the documentation search index only if you add to the core IBM FileNet P8 Platform documentation that is installed on an application server. For example, you must perform this task if you add expansion product documentation or auxiliary documentation such as release notes or technical notices. You must also update the search index any time you install an update of one of these collocated pieces of documentation.

"Starting and verifying the IBM FileNet P8 documentation web application" on page 8
After you install and optionally index the IBM FileNet P8 documentation again on the application server, verify that you can access its URL address and run its search feature. These tests ensure that the installed documentation can be used as a help system in the various IBM FileNet P8 component applications.

Refreshing or updating IBM FileNet P8 documentation

You can choose one of two methods to upgrade an IBM FileNet P8 documentation installation. The method you choose depends on whether the existing installation includes expansion product and other documentation in addition to core IBM FileNet P8 Platform documentation.

"Refreshing IBM FileNet P8 documentation without uninstalling" on page 5
You can refresh an existing IBM FileNet P8 documentation installation by copying the newer documentation directly over the existing files. Use this method if you have only the core IBM FileNet P8 Platform documentation installed and it is version 3.5 or later.

"Updating IBM FileNet P8 documentation by uninstalling and reinstalling" on page 6
If your existing documentation installation includes both core IBM FileNet P8 Platform documentation and other IBM FileNet P8 documentation, you must uninstall it completely before you can upgrade it.

Refreshing IBM FileNet P8 documentation without uninstalling

You can refresh an existing IBM FileNet P8 documentation installation by copying the newer documentation directly over the existing files. Use this method if you have only the core IBM FileNet P8 Platform documentation installed and it is version 3.5 or later.

If you update the core documentation package, you must also reindex the files for search support as a follow-up task.

Important:

- If you have IBM FileNet P8 3.5 or later documentation installed with expansion product documentation included, use the procedure "Updating IBM FileNet P8 documentation by uninstalling and reinstalling" on page 6.
- 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 server types, even those types such as IBM WebSphere® Application Server that required a WAR file for initial deployment.

To refresh the IBM FileNet P8 documentation without first uninstalling it:

1. Stop the IBM FileNet P8 documentation Web application or the application server 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 the search-index files located in:

Option	Description
UNIX®	*deployment_path*/ecm_help/search/index/core

3. Access the new IBM FileNet P8 Platform documentation package.
4. Copy the expanded `ecm_help` directory from the package over the deployed `ecm_help` directory.
5. As a best practice, download the latest web-posted updates of the planning and installation guides. For IBM FileNet P8 Platform, copy these PDF files to the `ecm_help/installation` directory on your IBM FileNet P8 documentation server.
6. (Optional) Download any auxiliary documentation to the `ecm_help/installation/web` directory on your IBM FileNet P8 Platform documentation server.
7. Continue with one of the following tasks:
    - If you added documentation to the core IBM FileNet P8 Platform documentation, then continue with "Updating the documentation search index" on page 7.
    - If you added no other documentation to the core IBM FileNet P8 Platform documentation, then continue with "Starting and verifying the IBM FileNet P8 documentation web application" on page 8.

Related information

➦ Product documentation for IBM FileNet P8 Platform
Download updated PDF's of the IBM FileNet P8 Platform documentation.

Updating IBM FileNet P8 documentation by uninstalling and reinstalling

If your existing documentation installation includes both core IBM FileNet P8 Platform documentation and other IBM FileNet P8 documentation, you must uninstall it completely before you can upgrade it.

Use the procedure provided for your application server type.

"Updating the IBM FileNet P8 documentation on IBM WebSphere Application Server"
To update IBM FileNet P8 documentation running on IBM WebSphere Application Server, you must uninstall the existing web application and then install the new documentation files.

Updating the IBM FileNet P8 documentation on IBM WebSphere Application Server

To update IBM FileNet P8 documentation running on IBM WebSphere Application Server, you must uninstall the existing web application and then install the new documentation files.

To update the documentation:
1. Verify that the IBM WebSphere Application Server machine is running, and then start the WebSphere administrative console.
2. From the WebSphere administrative console, choose **Uninstall** to remove the existing documentation site.
3. From the initial installation location, delete the installed IBM FileNet P8 documentation files but leave the `ecm_help` directory in place.
4. Copy the IBM FileNet P8 `ecm_help.war` file from the IBM FileNet P8 documentation package to the local hard disk drive.
5. From the WebSphere administrative console, deploy the `ecm_help.war` file by 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.

   You can copy more than one expansion product documentation set to the documentation server before continuing, for example, for Content Federation Services and IBM InfoSphere™ Enterprise Records. When you are finished, you have one ecm_help directory structure containing files for the core documentation and one or more sets of expansion product documentation.
7. As a best practice, download the latest Web-posted updates of the planning and installation guides. For IBM FileNet P8 Platform, copy these PDF files to the `ecm_help/installation` directory on your IBM FileNet P8 documentation server. Check the documentation page on the IBM customer support Web site http://www.ibm.com/software/data/support/ for the latest versions of these guides.
8. (Optional) Download any auxiliary documentation to the `ecm_help/installation/web` directory on your IBM FileNet P8 Platform documentation server.
9. From the WebSphere administrative console, restart the IBM FileNet P8 documentation site.

10. Continue with one of the following tasks:
    - If you added expansion products, continue with "Updating the documentation search index."
    - If you have no further documentation to install, continue with "Starting and verifying the IBM FileNet P8 documentation web application" on page 8.

Updating the documentation search index

You must update the documentation search index only if you add to the core IBM FileNet P8 Platform documentation that is installed on an application server. For example, you must perform this task if you add expansion product documentation or auxiliary documentation such as release notes or technical notices. You must also update the search index any time you install an update of one of these collocated pieces of documentation.

If no conditions require you to update the search index, skip to "Starting and verifying the IBM FileNet P8 documentation web application" on page 8.

To update the help Search index:

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

   Important: Ensure that you copied all 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 must repeat this procedure if you install additional documentation later.

2. Open a terminal or command prompt on the application server.

3. From the command line, navigate to the `search` subdirectory under the application root directory, for example, `ecm_help`.

4. Complete the search indexing steps:

Option	Description
UNIX	1. If not already done, apply chmod 755 permissions on the `indexFiles.sh` file, which is preset to read-only permissions in the documentation package.
	2. Using a text editor, open the indexFiles.sh file.
	3. 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. **Important:** The Java JRE installation subdirectory can be user-defined, so substitute your actual location.
	4. Save your changes and close the text editor.
	5. From the `ecm_help/search` directory, run the updated indexFiles.sh search-indexing script file.
	As you run the script file, you might notice periodic Parse Abort errors. You can ignore these error conditions because they are benign and do not affect the overall indexing process.

Starting and verifying the IBM FileNet P8 documentation web application

After you install and optionally index the IBM FileNet P8 documentation again on the application server, verify that you can access its URL address and run its search feature. These tests ensure that the installed documentation can be used as a help system in the various IBM FileNet P8 component applications.

To verify the IBM FileNet P8 documentation installation:

1. Start the IBM FileNet P8 documentation Web application (`ecm_help`) on the application server.
2. From a Web browser, access the documentation URL by using the application server name and port number for your Web environment, for example:

Option	Description
IBM WebSphere Application Server	http://*yourdocserver*:9080/ecm_help/

If the installation was a success, the IBM FileNet P8 documentation opens the same way as for an IBM FileNet P8 component application.

3. Note the URL for your application server because you must enter it as the documentation URL for the IBM FileNet P8 components while running installation programs or setting site preferences later, for example, in Workplace or Workplace XT.
4. Click the **Search** link on the Help Directory toolbar and verify that it opens the documentation Search page.
5. Enter a value for your Search query and verify that the query runs as expected.
6. Select one of the Search query result links and verify that the help page opens.

Upgrading and configuring Content Engine and Content Search Engine

You can upgrade Content Engine and Content Search Engine to version 4.5.1 from version 3.5.2, 4.0.1, or 4.5.0.

Important: In some customized versions of this documentation, Content Search Engine information is removed from this topic and its subtopics.

Perform the following upgrade tasks that apply to your existing version of Content Engine and Content Search Engine.

"Upgrading from version 3.5.2"
Upgrading Content Engine and Content Search Engine from version 3.5.2 is a multi-step process.

"Upgrading from version 4.0.1 or 4.5.0" on page 92
Upgrading Content Engine and Content Search Engine from version 4.0.1 or 4.5.0 is a multi-step process.

Upgrading from version 3.5.2

Upgrading Content Engine and Content Search Engine from version 3.5.2 is a multi-step process.

Important: Before you start to upgrade the Content Engine software, ensure that you are at the correct version and fix pack level, as indicated in *Plan and Prepare Your Environment for IBM FileNet P8*.

1. "Completing pre-upgrade Content Engine configuration" on page 10
   You must complete several pre-upgrade procedures to prepare the Content Engine for an upgrade from version 3.5.2.

2. "Upgrading Content Engine and related client software" on page 16
   If you are upgrading from Content Engine version 3.5.2 on Windows, you can upgrade and continue to use Windows, or you can migrate your Content Engine to UNIX.

3. "Configuring Content Engine" on page 23
   You can configure and deploy all of your Content Engine instances by using Configuration Manager. A single Content Engine application instance equates to one deployed application on your application server.

4. "Deploying upgraded Content Engine instances" on page 44
   You must deploy the upgraded Content Engine on the web application server. Deploying makes the Content Engine application available for use.

5. "Installing storage device source files" on page 47
   If your upgraded IBM FileNet P8 system includes Tivoli Storage Manager or EMC Centera devices, you must install files on the Content Engine server.

6. "Completing Content Engine post-deployment steps" on page 50
   You must complete several Content Engine post-deployment steps before you can put an IBM FileNet P8 system into production.

7. "Installing the latest Content Engine client files on other IBM FileNet P8 servers (for staged upgrades)" on page 53
   When you upgrade Content Engine 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.

8. "Establishing the FileNet P8 domain and Global Configuration Data" on page 53
   You must use IBM FileNet Enterprise Manager to create a FileNet P8 domain and the Global Configuration Data (GCD).

9. "Configuring storage devices for upgrades" on page 56
   When you upgrade from Content Engine version 3.5.2, you must configure your existing file store devices.

10. "Creating the data sources for an object store" on page 59
    Each object store in your site requires its own JDBC data sources. You can configure the data sources by using the Configuration Manager graphical user interface or the command line.

11. "Upgrading Content Search Engine software from version 3.5.2" on page 63
    If you used the full-text search feature (CBR) in version 3.5.2, you upgrade by installing and configuring version 4.5.1 of Content Search Engine (Autonomy K2). After the upgrade, you recreate index areas, called *K2 collections*, to replace the version 3.5.2 content search indexes.

12. "Updating the CFS-IS settings" on page 74
    If Content Federation Services for FileNet Image Services (CFS-IS) is in your version 3.5.2 system, then you must enable the FileNet Image Services dispatcher. You must also modify the default map between CFS-IS and the object stores in the FileNet P8 domain.

13. "Upgrading Content Engine data" on page 75
    You must upgrade your current Content Engine data (GCD, addons, fixed content devices, object stores, and related information).

14. "Updating the database storage area parameters" on page 88
    After you have upgraded your object stores, you need to initialize the parameter values for the associated database storage areas.

15. "Completing Content Search Engine upgrade from 3.5.2" on page 88
    After you upgrade Content Engine data, you must create new content-search index areas in Enterprise Manager. Then you can enable indexing for designated classes and properties and reindex your stored data. The reindexing process creates new Autonomy K2 collections.

16. "Upgrading Enterprise Manager" on page 91
    If you did not upgrade Enterprise Manager as part of a Content Engine upgrade, you can upgrade Enterprise Manager interactively or silently.

Related information

➡ General requirements for all FileNet P8 Platform systems
Check to make sure your software is at the correct fix pack level before you upgrade.

Completing pre-upgrade Content Engine configuration

You must complete several pre-upgrade procedures to prepare the Content Engine for an upgrade from version 3.5.2.

1. "Verifying that all in-progress event actions have finished" on page 11
   You must verify that all in-progress events have completed for each object store before you begin the upgrade from version 3.5.2.

2. "Verifying that publishing requests have been completed" on page 12
   You must verify that all publishing requests have completed for each object store before you begin the upgrade from version 3.5.2.

3. "Deleting custom subclasses of the PublishRequest object" on page 13
   You must delete any custom subclasses of the PublishRequest class for each object store before you begin the upgrade from version 3.5.2.

4. "Verifying that all in-progress Content Engine transactions involving content files have finished" on page 13
   You must verify that all pending content transactions, indexing requests, and fixed-content migration requests have completed before you begin the upgrade from version 3.5.2. Use the Content Resource Manager utility to view the queues.

5. "Renaming storage policies" on page 13
   Check that every storage policy name does not exceed the maximum allowed number of characters. You must shorten the name of any storage policy in version 3.5.2 of IBM FileNet P8 so that it does not exceed 58 characters.

6. "Removing accented characters from file store paths" on page 14
   If you are moving your file stores from Windows to UNIX during your upgrade to version 4.5.1, you must remove any accented characters from the file store path prior to moving the file store.

7. "Determining the CBR-enabled properties in an object store" on page 14
   For each object store to be upgraded, you must determine the associated classes and properties that are currently enabled for content-based retrieval (CBR). You enable these properties after the upgrade is complete.

8. "Preventing name conflicts with reserved words" on page 14
   You must verify that the object store names, custom classes, or custom properties in your current version do not conflict with reserved words before you begin the upgrade from version 3.5.2.

9. "Upgrading IBM InfoSphere Enterprise Records object stores" on page 15
   If you site includes FileNet Records Manager, you must update the FileNet Records Manager data model on each file plan object store (FPOS) before you install Content Engine version 4.5.1.

10. "Running the Content Consistency Checker" on page 15
    Run the Content Consistency Checker to ensure that the documents in the object store database are synchronized with file storage area content.

11. "Moving IBM TotalStorage DR550 fixed content device data" on page 15
    Although IBM TotalStorage DR550 devices have been supported since version 4.5.0 FP2 of Content Engine, there is no upgrade path to move file stores that are on such devices from version 3.5.2 to 4.5.1. If you need to move the content of your DR550 fixed content devices from Content Engine version 3.5.2 to version 4.5.1, contact IBM Support.

12. "Backing up the database" on page 15
    You must back up the database before you begin the upgrade from version 3.5.2.

13. "Stopping and disabling all services related to Content Engine" on page 15
    You must stop and disable all services related to Content Engine on all servers in the FileNet P8 domain before you begin the upgrade from version 3.5.2.

Verifying that all in-progress event actions have finished

You must verify that all in-progress events have completed for each object store before you begin the upgrade from version 3.5.2.

To locate and remove in-progress items from the work queue for the object store:

1. Start Enterprise Manager if it is not already running.
2. In the list view, under the object store icon, right-click the **Search Results** folder and select **New Search**.

3. In the Content Engine Query Builder window, click **Select From Table** →
   QueueItem.
4. Search for unfinished items in the queue or delete unwanted items from the
   queue:

Option	Description
Search for unfinished items in the queue	1. Retain the default settings and click **OK**. 2. Click **Yes** at the prompt for a WHERE clause. Any remaining event items in the queue are displayed in the Query Status window.
Remove unwanted items from the queue	1. Select the **Delete objects** check box in the **Action** tab of the Search dialog box. 2. Click **OK**. Click **OK** again to confirm the deletion.

5. Repeat this procedure for each object store in your environment.

Verifying that publishing requests have been completed

You must verify that all publishing requests have completed for each object store
before you begin the upgrade from version 3.5.2.

To verify that publishing requests complete:
1. Start Enterprise Manager if it is not already running.
2. Search for publishing requests in the queue by expanding the **Publishing** folder
   and clicking **Queue**. Then, click **View** → **All requests**.
3. Verify that the queue (right pane) of publishing requests is empty. If the queue
   is not empty, wait until the publish requests are processed. 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).
4. In the list view, under the object store icon, right-click the **Search Results**
   folder, and choose **New Search**.
5. In the Content Engine Query Builder window, click **Select From Table** →
   PublishRequest.
6. Retain the default settings and click **OK**.
7. Click **Yes** at the prompt for a WHERE clause. Any remaining event items in the
   queue are displayed in the Query Status window.
8. Remove unwanted items from the queue.
   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 window, click **Select From Table** →
      PublishRequest.
   c. Select the **Delete objects** check box in the **Actions** tab of the Search
      window.
   d. Click **OK**. Click **OK** again to confirm the deletion.
9. Repeat this procedure for each object store in your environment.

Deleting custom subclasses of the PublishRequest object

You must delete any custom subclasses of the PublishRequest class for each object store before you begin the upgrade from version 3.5.2.

To delete custom subclasses of the PublishRequest object:

1. Start version 3.5 of IBM FileNet Enterprise Manager.
2. In the left pane of Enterprise Manager, open an object store, and navigate to **Other Classes** → **Custom Object** → **Publish Request**.
3. In the right pane, delete any custom subclasses of Publish Request and any instances of such subclasses.
4. Repeat this procedure for each object store in your environment.
5. Close Enterprise Manager.

Verifying that all in-progress Content Engine transactions involving content files have finished

You must verify that all pending content transactions, indexing requests, and fixed-content migration requests have completed before you begin the upgrade from version 3.5.2. Use the Content Resource Manager utility to view the queues.

Important: All transactions and requests must complete before you upgrade Content Engine. Otherwise, the upgrade will fail.

To verify that all transactions and requests are complete:

1. Start File Store Service if it is not already running on the server where Content Engine is installed.
2. Stop Object Store Service and Content Cache Service.
3. Start the *ce_install_path*\RMU.exe utility, where *ce_install_path* is the location where the Content Engine software is installed.
4. 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 are complete.
   d. Navigate to **Indexing Service** → **Index Control**. If the value of **Current index queue files** is zero, then all pending indexing requests are complete.
   e. Navigate to **Fixed Content** → **Migration Queuing**. If the value of **Current number in queue** is zero, then all fixed-content migration requests are complete.
5. If all the activities checked are complete, then stop the File Store Service on the server.
6. If any of the activities are not complete, then close the Content Resource Manager utility. Wait several minutes, and then return to Step 1.
7. Repeat this procedure on each server in your environment where Content Engine is installed.

Renaming storage policies

Check that every storage policy name does not exceed the maximum allowed number of characters. You must shorten the name of any storage policy in version 3.5.2 of IBM FileNet P8 so that it does not exceed 58 characters.

To check storage policy names:

1. Start version 3.5.2 of IBM FileNet Enterprise Manager and log in to the P8 domain FileNet P8 domain you are upgrading.

2. For each object store in the left panel, select the Storage Policies node and complete the following substeps for each storage policy:

   a. Right-click the storage policy node and choose **Properties** from the context menu that opens.

   b. If the storage policy name has more than 58 characters, change the name so that it does not exceed 58 characters and click **OK**.

Removing accented characters from file store paths

If you are moving your file stores from Windows to UNIX during your upgrade to version 4.5.1, you must remove any accented characters from the file store path prior to moving the file store.

To remove any accented characters from the file store path:

1. Start the version 3.5.2 Enterprise Manager.

2. Use the Move File Store wizard to change the path for the file store. Do not use any accented characters in the path. See the FileNet P8 Help topic **IBM FileNet P8 documentation** → **System Administration** → **Content Engine Administration** → **Content storage** → **Storage areas** → **File storage areas** → **How to** → **Move a file storage area**

Determining the CBR-enabled properties in an object store

For each object store to be upgraded, you must determine the associated classes and properties that are currently enabled for content-based retrieval (CBR). You enable these properties after the upgrade is complete.

To determine the classes and properties that are CBR-enabled:

1. Using an RDBMS tool, access the object store database as a user who has appropriate privileges to query the object store tables.

2. Run the following SQL query command on each object store database to determine which properties are enabled for CBR:

   ```
   select p.dbg_class_name, p.dbg_display_name, c.is_cbr_enabled,
   p.is_cbr_enabled from ClassDefinition c inner join PropertyDefinition p
   on p.parent_id = c.object_id where p.is_cbr_enabled = 1 AND
   c.is_cbr_enabled = 0
   ```

3. Record the results of the query commands. After you upgrade the object stores, you need to run FileNet Enterprise Manager to enable CBR on these properties after you enable CBR on the classes that these properties are associated with.

Related tasks

"Completing Content Search Engine upgrade from 3.5.2" on page 88
After you upgrade Content Engine data, you must create new content-search index areas in Enterprise Manager. Then you can enable indexing for designated classes and properties and reindex your stored data. The reindexing process creates new Autonomy K2 collections.

Preventing name conflicts with reserved words

You must verify that the object store names, custom classes, or custom properties in your current version do not conflict with reserved words before you begin the upgrade from version 3.5.2.

Verify that your custom classes and properties do not conflict with the reserved words listed in "New Content Engine classes and properties" in *Plan and Prepare Your Environment for IBM FileNet P8* that describes how you must change any custom class or custom property names in your existing object stores. If you have previously defined any classes or properties in a version 3.5.2 object store with symbolic names that are the same as any of the reserved class or property symbolic names, the upgrade tool detects the naming conflicts and requires you to change your class or property names before it upgrades the object store to version 4.5.1. Therefore, you must resolve conflicts between your version 3.5.2 symbolic names and the version 4.0 and 4.5 symbolic names

Related concepts

⬚➡ Upgrade tasks
For details on the necessary steps to prepare for an upgrade, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading IBM InfoSphere Enterprise Records object stores

If you site includes FileNet Records Manager, you must update the FileNet Records Manager data model on each file plan object store (FPOS) before you install Content Engine version 4.5.1.

To complete and confirm your IBM FileNet P8 Platform upgrade, you must update at least one object store. If your configuration contains only FileNet Records Manager object stores, you must update at least one of the object stores or create a non-Records Manager object store prior to upgrading and use that object store to confirm the upgrade.

Run the RM37PreUpgrader tool. See the topic "Prepare to Upgrade Object Stores" in *IBM FileNet Records Manager Installation and Upgrade Guide*

Running the Content Consistency Checker

Run the Content Consistency Checker to ensure that the documents in the object store database are synchronized with file storage area content.

For information about the Content Consistency Checker, see the IBM FileNet P8 Help topic **FileNet P8 Administration** ➔ **Content Engine Administration** ➔ **Content storage** ➔ **Content Consistency Checker Overview**.

Moving IBM TotalStorage DR550 fixed content device data

Although IBM TotalStorage DR550 devices have been supported since version 4.5.0 FP2 of Content Engine, there is no upgrade path to move file stores that are on such devices from version 3.5.2 to 4.5.1. If you need to move the content of your DR550 fixed content devices from Content Engine version 3.5.2 to version 4.5.1, contact IBM Support.

Backing up the database

You must back up the database before you begin the upgrade from version 3.5.2.
1. Use your existing database backup solution to back up each of the Content Engine object store databases.
2. Make a backup copy of the version 3.5.2 Global Configuration Data GCD file. The default location is `C:\Program Files\FileNet\Content Engine\sysconfig\sysinit\sysinit.dat`.

Stopping and disabling all services related to Content Engine

You must stop and disable all services related to Content Engine on all servers in the FileNet P8 domain before you begin the upgrade from version 3.5.2.

1. Log on with local administrator permissions to the server where Content Engine version 3.5.2 is installed.
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
   - FileNet P8 CFS Server for FileNet Image Services
3. Right-click the Apache services monitor in the Windows system tray, and click **Exit**.
4. Repeat this procedure on each Content Engine server in your FileNet P8 domain.

Upgrading Content Engine and related client software

If you are upgrading from Content Engine version 3.5.2 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 3.5.2 on a Windows server to Content Engine version 4.5.1 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 IBM FileNet Enterprise Manager and the Content Engine Upgrader tool.

It is a best practice to retain a Windows machine where version 3.5.2 of IBM FileNet Enterprise Manager is installed. If you use the Content Engine Upgrader tool, you need access to this Windows-based tool. If you use Content Engine Web Upgrade Tool, you can browse to it only from a machine where version 8.0 of Microsoft Internet Explorer is installed.

1. "Installing or upgrading Content Engine" on page 17
   You can install or upgrade Content Engine, Configuration Manager, .NET Clients (including IBM FileNet Enterprise Manager), Content Engine Upgrader, and FileNet Deployment Manager either interactively or silently.
2. "Installing Content Engine software updates" on page 19
   You must apply any updates, fix packs, or interim fixes for the Content Engine software that have been released.
3. "Installing the latest Content Search Engine Client files on Content Engine servers" on page 19
   If your environment includes Content Search Engine, you must install the files that Content Engine uses to communicate with Content Search Engine.
4. "Installing the latest Process Engine client files on Content Engine servers" on page 21
   You must install the Process Engine Client files on all application server machines where Content Engine Server is deployed.

Installing or upgrading Content Engine

You can install or upgrade Content Engine, Configuration Manager, .NET Clients (including IBM FileNet Enterprise Manager), Content Engine Upgrader, and FileNet Deployment Manager either interactively or silently.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

If you are installing in a managed environment, then the application server must be the Deployment Manager node for WebSphere Application Server or the Administrator node for Oracle WebLogic Server.

Some Content Engine components must be installed on a Windows® operating system. The following table lists the Content Engine components and the required operating system.

Table 1. Content Engine components

Component	Required Platform
Content Engine Server	UNIX or Windows
Configuration Manager	UNIX or Windows
.NET Clients (including IBM FileNet Enterprise Manager)	Windows only
Content Engine Upgrader	Windows only
FileNet Deployment Manager	Windows only

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

Important: You must keep one running version of Enterprise Manager version 3.5.2 until after you upgrade the Content Engine data. You must use Enterprise Manager version 3.5.2 to prepare your version 3.5.2 object stores for the Content Engine data upgrade.

The Content Engine Upgrader tool or Content Engine Web Upgrade Tool is required to upgrade Content Engine and Content Search Engine data from version 3.5.2.

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.

You can install or upgrade Content Engine interactively or silently.

"Installing Content Engine interactively" on page 18
You can interactively install or upgrade Content Engine, Configuration Manager, .NET Clients (including IBM FileNet Enterprise Manager), Content Engine Upgrader, and IBM FileNet Deployment Manager.

"Installing Content Engine silently" on page 18
You can silently install or upgrade Content Engine, Configuration Manager, .NET Clients (includingIBM FileNet Enterprise Manager), Content Engine Upgrader, and IBM FileNet Deployment Manager.

Related concepts

➥ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuring storage devices for migration to UNIX" on page 56
When you upgrade from Content Engine version 3.5.2 and migrate to a UNIX
server, you must prepare the UNIX server to mount your existing file store devices.

Installing Content Engine interactively:

You can interactively install or upgrade Content Engine, Configuration Manager,
.NET Clients (including IBM FileNet Enterprise Manager), Content Engine
Upgrader, and IBM FileNet Deployment Manager.

Be sure that you have available the Installation and Upgrade Worksheet that was
filled out during your planning activities.

Important: You must keep one running version of Enterprise Manager version 3.5
until after you upgrade the Content Engine data. You must use Enterprise Manager
version 3.5 to prepare your version 3.5 object stores for the Content Engine
Upgrader tool.

To install Content Engine:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command
   is enabled. To view only Content Engine values, filter by **CE Installer** in the
   Installation or Configuration Program column.

2. Log on as *ce_install_user* to the application server machine where you are going
   to install Content Engine software.

3. Navigate to the Content Engine software package in the installation media.

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

Platform	Command
AIX®	P8CE-4.5.1-AIX.BIN

5. When the installation completes, click **Done** and check for errors in the Content
   Engine *ce_install_path*/ce_install_log_4_5_1.txt error log file, where
   ce_install_path is the location where Content Engine is installed.

Installing Content Engine silently:

You can silently install or upgrade Content Engine, Configuration Manager, .NET
Clients (includingIBM FileNet Enterprise Manager), Content Engine Upgrader, and
IBM FileNet Deployment Manager.

Be sure that you have available the Installation and Upgrade Worksheet that was
filled out during your planning activities.

Important: You must keep one running version of Enterprise Manager version 3.5.2 until after you upgrade the Content Engine data. You must use Enterprise Manager version 3.5.2 to prepare your version 3.5.2 object stores for the Content Engine Upgrader tool.

To install Content Engine:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Installer** in the **Installation or Configuration Program** column.

2. Log on as *ce_install_user* to the application server machine where you are going to install Content Engine software.

3. Navigate to the Content Engine software package in the installation media.

4. Edit the `CE_silent_install.txt` file to reflect the installation choices in your worksheet.

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

Platform	Command
AIX	`P8CE-4.5.1-AIX.BIN -f` `CE_silent_install.txt -i silent`

6. When the installation completes, check for errors in the Content Engine error log file *ce_install_path*/`ce_install_log_4_5_1.txt`, where *ce_install_path* is the location where Content Engine is installed.

Installing Content Engine software updates

You must apply any updates, fix packs, or interim fixes for the Content Engine software that have been released.

If no Content Engine software updates are available, skip this procedure.

To install the Content Engine software updates:

Open the readmes for the Content Engine software updates, and perform the installation procedures in the readmes on each Content Engine instance.

Installing the latest Content Search Engine Client files on Content Engine servers

If your environment includes Content Search Engine, you must install the files that Content Engine uses to communicate with Content Search Engine.

Important: With the initial release of Content Engine 4.5.1, the Content Search Engine 4.5.1 client files are included. Therefore, you are not required to manually install these client files on a new Content Engine server. However, if fix packs of these client files become available, install them manually before you use Configuration Manager to configure and deploy Content Engine. For instructions on how to obtain the latest Content Search Engine Client software updates, access the IBM FileNet P8 Platform support site.

You can manually install Content Search Engine Client files either interactively or silently:

"Installing Content Search Engine Client files interactively"
You can install the Content Search Engine Client files interactively by running the installation program on all application server machines where Content Engine Server is installed.

"Installing Content Search Engine Client files silently"
You can install the Content Search Engine Client files silently by editing a silent input file and running the installation program from a command line. You must install Content Search Engine Client files on all application server machines where Content Engine Server is installed.

Installing Content Search Engine Client files interactively:

You can install the Content Search Engine Client files interactively by running the installation program on all application server machines where Content Engine Server is installed.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install Content Search Engine Client files:
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Search Engine Client values, filter by **CSE Client Installer** in the **Installation or Configuration Program** column.
2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.
3. On the machine where Content Engine is to be deployed, log on as *ce_install_user*.
4. Locate the Content Search Engine Client installation software. It must be the same version as your Content Search Engine Server software. You must install Content Search Engine and Content Engine software on the same operating system. For example, if Content Engine runs on Windows, locate the Content Search Engine Client installation software that runs on Windows.
5. Start the installation program.

Operating System	Install Program
AIX	P8CSE-CLIENT-4.5.1-AIX.BIN

6. Complete the Content Search Engine Client installation screens by using the values in your worksheet.
7. Review the following Content Search Engine Client log files:
   - UNIX - *cse_client_install_path*/CSEClient_install_log_4.5.1.0.log

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing Content Search Engine Client files silently:

You can install the Content Search Engine Client files silently by editing a silent input file and running the installation program from a command line. You must install Content Search Engine Client files on all application server machines where Content Engine Server is installed.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install Content Search Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ⟩ **Filter** ⟩ **AutoFilter** command is enabled. To view only Content Search Engine Client values, filter by **CSE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. On the machine where Content Engine is to be deployed, log on as *ce_install_user*.

4. Locate the Content Search Engine Client installation software. It must be the same version as your Content Search Engine Server software. You must install Content Search Engine and Content Engine software on the same operating system. For example, if Content Engine runs on Windows, locate the Content Search Engine Client installation software that runs on Windows.

5. Locate the `CSE_silent_install.txt` file within the directory that contains the Content Search Engine Client installation software, and edit the file to specify the location where you want to install the Content Search Engine Client files.

   The default installation paths are as follows:

   UNIX `/opt/IBM/FileNet/ContentEngine/CSEClient`

6. Start the silent installation of Content Search Engine Client files by running the following command.

Operating System	Install Program
AIX	`P8CSE-CLIENT-4.5.1-AIX.BIN —f CSEClient_silent_install.txt —i silent`

7. Review the following Content Search Engine log files:
   - UNIX - *cse_client_install_path*/ `CSEClient_install_log` *release_number*.log

Related concepts

⏩ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing the latest Process Engine client files on Content Engine servers

You must install the Process Engine Client files on all application server machines where Content Engine Server is deployed.

The upgrade wizard provides an interactive way to upgrade the Process Engine

Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

"Upgrading the latest Process Engine client files on Content Engine servers silently"
You can upgrade the Process Engine Client software on the Content Engine server by using a silent input text file and running the upgrade program from the command line. Use the values in your installation and upgrade worksheet to edit the silent input text file before you run the upgrade.

Upgrading the latest Process Engine client files on Content Engine servers interactively:

The upgrade wizard provides an interactive way to upgrade the Process Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To upgrade the Process Engine Client files:
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.
2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.
3. On the machine where Content Engine is to be deployed, log on as *ce_upgrade_user*.
4. 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.
5. Expand the (TAR or ZIP) Process Engine Client install software.
6. Shut down all instances of Enterprise Manager and any other Content Engine client applications, such as Application Engine.
7. Run the command for the installation program according to your operating system.

Operating System	Install Program
AIX	P8PE-CLIENT-4.5.1-AIX.BIN

8. Complete the Process Engine client install screens.

Related concepts

⏩ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading the latest Process Engine client files on Content Engine servers silently:

You can upgrade the Process Engine Client software on the Content Engine server by using a silent input text file and running the upgrade program from the

command line. Use the values in your installation and upgrade worksheet to edit the silent input text file before you run the upgrade.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install the Process Engine Client files silently:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. On the machine where Content Engine is to be deployed, log on as *ce_upgrade_user*.

4. 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.

5. Expand the (TAR or ZIP) Process Engine Client install software.

6. Shut down all instances of Enterprise Manager and any other Content Engine client applications, such as Application Engine.

7. In the expanded install software, open the file `PEClient_silent_install.txt` and edit it as follows: Change the `Variable_CheckboxCE` line to the following line:

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

8. Save your edits.

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

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.1-AIX.BIN -silent` `-options "PEClient_silent_install.txt"`

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring Content Engine

You can configure and deploy all of your Content Engine instances by using Configuration Manager. A single Content Engine application instance equates to one deployed application on your application server.

Configuration is a multiple step process. You must provide information about your Content Engine application environment, then you apply the settings by running the configuration tasks, and then you deploy the application. You can configure multiple instances before you deploy any of them, or you can configure and deploy one instance at a time.

You use Configuration Manager to define the following information for the Content Engine instance:

- Application server properties
- Java Database Connectivity (JDBC) data source properties for the Global Configuration Data (GCD) database
- Java Database Connectivity (JDBC) data source properties for each object store database
- Directory service (LDAP) provider properties
- Content Engine application login modules
- Content Engine bootstrap properties

Remember the following points when you prepare to configure Content Engine:
- (WebSphere Application Server only) For best results, configure no more than one Content Engine instance in a WebSphere profile.
- If you need an accessible software version of Configuration Manager for people with disabilities, use the command-line version of Configuration Manager.
- If your FileNet P8 domain uses multiple non-managed application servers, then you use a copy of a single EAR file with the bootstrap settings on all the servers. After you complete the tasks on the initial server, you copy the EAR file with the bootstrap settings to the additional servers.

Begin the configuration process by ensuring that you have the necessary operating system permissions to use Configuration Manager.

To grant permissions to the Configuration Manager user:
1. Log on as *ce_install_user* to the server where you installed Content Engine.
2. Grant permissions to the *config_mgr_user* user for the executable file that you intend to use:

Option	Description
UNIX - graphical user interface	Grant execute permission to `configmgr`
UNIX - command line	Grant execute permission to `configmgr_cl`

3. Grant write permission to the directory where you want Configuration Manager to place the configuration XML files that it creates.

    If you do not specify the profile directory when you run Configuration Manager, grant write permission on the default directory `ce_install_path/tools/configure` and all its files and subdirectories.
4. Log off the Content Engine server and log on again as *config_mgr_user*.

1. "Configuring Content Engine instances by using the graphical user interface" You can use the graphical user interface version of Configuration Manager to configure and deploy a Content Engine application instance on an application server.
2. "Configuring instances by using the command line" on page 32 You can configure a Content Engine instance on a web application server by using the command-line version of Configuration Manager.

Configuring Content Engine instances by using the graphical user interface

You can use the graphical user interface version of Configuration Manager to configure and deploy a Content Engine application instance on an application server.

First, you create a configuration profile that defines:

- The JDBC data source settings
- The login modules for the Content Engine application
- The directory service provider (LDAP) settings
- The Content Engine bootstrap settings

You can follow the recommended order to edit and run the individual configuration tasks. However, you can edit and run the Configure GCD JDBC Data Sources, Configure Login Modules, Configure LDAP properties, and Configure Bootstrap tasks in any order. You do not need to complete work on one configuration task before starting another. You can save your edits, switch between tasks, close the tasks, and reopen tasks as needed. However, you must complete all of these tasks before you deploy the application.

1. "Starting Configuration Manager"
   You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.
2. "Changing the password save preference" on page 26
   By default, any passwords that you enter in the graphical user interface in Configuration Manager are not saved to a file. The passwords are stored in memory as long as you have the profile open in Configuration Manager. You can change the password save preference to save the passwords each time that you save your changes to a profile.
3. "Creating a new configuration profile" on page 26
   Configuration Manager stores your settings for deploying a Content Engine application instance in a configuration profile. The profile defines the application server settings, the JDBC data source settings, the login modules settings, the directory service provider (LDAP) settings, the Content Engine bootstrap settings, and the application deployment settings. You must create a new profile for each Content Engine instance that you are configuring.
4. "Configuring the Global Configuration Data JDBC data source settings" on page 27
   You must provide property information about the JDBC data sources for the Global Configuration Data (GCD) database. Content Engine uses the data source information to connect to and update the GCD database.
5. "Configuring the login modules (WebSphere and JBoss only)" on page 28
   You must configure the login modules for the Content Engine application. The login modules provide authentication information for the Content Engine application.
6. "Configuring the directory service provider (LDAP) settings" on page 29
   You must configure the directory service provider (LDAP) settings. The directory service provider (LDAP) settings define the directory service and the users and groups that Content Engine uses for authentication.
7. "Configuring the bootstrap settings" on page 30
   You must configure the bootstrap settings. The bootstrap information is needed for creating the Global Configuration Data (GCD) and for starting Content Engine.

Starting Configuration Manager:

You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.

To start Configuration Manager:

1. 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:

Option	Description
UNIX	Run this command: `ce_install_path/tools/configure/configmgr`

The first time that you start Configuration Manager, the Welcome is displayed.

2. Select one of the links in the Welcome to learn more or to start working in a wizard, or close the Welcome by clicking the **X** in the tab at the upper left. You can reopen the Welcome later, as needed, from the **Help** menu.

Related tasks

"Deploying upgraded instances by using the graphical user interface" on page 44 You can deploy an upgraded Content Engine Server instance on a web application server by using the graphical user interface version of Configuration Manager. Deploying makes the Content Engine application available for use.

Changing the password save preference:

By default, any passwords that you enter in the graphical user interface in Configuration Manager are not saved to a file. The passwords are stored in memory as long as you have the profile open in Configuration Manager. You can change the password save preference to save the passwords each time that you save your changes to a profile.

When you close the profile, the passwords are erased from memory. Each time that you start Configuration Manager or open a saved profile, you must specify any passwords required by the tasks and for the application server properties before you run the tasks; otherwise, the tasks won't run successfully. If your site security requirements permit you to save passwords to a file, you can change the password save preference setting.

To change the password save preference:

1. Click **Window** → **Preferences**.
2. Complete one of the following actions:

Option	Description
To save passwords to file	Select the **Save all passwords to file when saving a task or profile** check box.
To prevent writing passwords to file	Clear the **Save all passwords to file when saving a task or profile** check box.

3. Click **OK**.

Creating a new configuration profile:

Configuration Manager stores your settings for deploying a Content Engine application instance in a configuration profile. The profile defines the application server settings, the JDBC data source settings, the login modules settings, the directory service provider (LDAP) settings, the Content Engine bootstrap settings, and the application deployment settings. You must create a new profile for each Content Engine instance that you are configuring.

The information for a profile is collected in XML files in the form of properties and values that describe the associated configuration and deployment tasks. You must provide values for the profile properties that are specific to each configuration at your site, such as the application server name. The XML files are stored in a directory that is unique to a given profile. Because the profile name is used for the directory name and for the configuration profile file name, you must provide a profile name that is a valid directory name for your operating system. By default, the profiles are stored in the *ce_install_path*/`tools/configure/profiles` directory, where *ce_install_path* is the location where Content Engine is installed. For more information on profiles, see "Configuration profile concepts" on page 216.

Tip: For more information on the properties and values that you set in Configuration Manager, roll your mouse over the property name to view the tool tip help for the property.

To create a configuration profile:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** �assistant **Filter** ⁺ **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Create New Installation Profile** in the **Installation or Configuration Program** column.

2. Start the Create New Installation Profile wizard by selecting **File** ⁺ **New Installation Profile**.

3. Complete the wizard screens by using the values in your worksheet.

4. Optional: WebSphere Application Server and Oracle WebLogic Server only. In the Set Properties for Application Server screen, click **Test Connection** to test the connection between Configuration Manager and the application server by using the information that you provided. The test is optional, and you can proceed in the wizard even if the test fails. If the test fails, make sure that the application server is running and that the application server property values that you entered match the values that are defined in your application server.

The profile you created is displayed as an icon in the profile pane (left pane), along with icons for the tasks you selected. By default, the **Deploy Application** task is disabled. You must enable the task later in the configuration process.

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Configuring the Global Configuration Data JDBC data source settings:

You must provide property information about the JDBC data sources for the Global Configuration Data (GCD) database. Content Engine uses the data source information to connect to and update the GCD database.

Ensure that the application server is running or stopped, depending on its type:

Table 2. Required application server state

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

To configure the JDBC settings:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ▸ **Filter** ▸ **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Configure GCD JDBC Data Sources** in the **Installation or Configuration Program** column.

2. If your configuration profile is not open in Configuration Manager, open the profile.

3. Right-click **Configure GCD JDBC Data Sources** in the profile pane and select **Edit Selected Task**.

4. Enter the property values for your database by using the values in your worksheet.

5. Optional: (WebSphere and WebLogic only) Click **Test Database Connection** to test the connection to the database by using the database user name, database server name, database name, port number, and password that you provided. The test does not create the data sources.

6. Select **File** ▸ **Save** to save your changes.

7. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, select **Configure GCD JDBC Data Sources (disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.

8. Apply the JDBC property settings by right-clicking **Configure GCD JDBC Data Sources** in the profile pane, and selecting **Run Task**. Running the configuration task can take several minutes. The task execution status messages are displayed in the console pane below the JDBC properties.

9. Close the Configure GCD JDBC Data Sources task pane. In this step, you created the GCD data sources. You create the initial object store data sources later in the configuration process.

Related concepts

⬆ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Configuring the login modules (WebSphere and JBoss only):

You must configure the login modules for the Content Engine application. The login modules provide authentication information for the Content Engine application.

Ensure that the application server is running or stopped, depending on its type:

Table 3. Required application server state

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

To configure the login module settings:

1. If your configuration profile is not open in Configuration Manager, open the profile.

2. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, select **Configure Login Modules (disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.

3. Create the login modules by right-clicking **Configure Login Modules** in the profile pane, and selecting **Run Task**. Running the configuration task can take several minutes. The task execution status messages are displayed in the console pane below the login modules properties.

4. Close the Configure Login Modules task pane.

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Configuring the directory service provider (LDAP) settings:

You must configure the directory service provider (LDAP) settings. The directory service provider (LDAP) settings define the directory service and the users and groups that Content Engine uses for authentication.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

If you plan to configure Content Engine to use the directory server's e-mail attribute or, for Active Directory, the userPrincipalName (UPN) to be the user short name used for login, then you must perform additional configuration steps and enter specific values for your LDAP settings. For detailed steps, see the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **How to...** → **Configuration Content Engine to use e-mail or UPN for login**.

To configure the LDAP settings:

1. Open your completed Installation and Upgrade Worksheet file.

Tip: In the worksheet file, verify that the **Data** ⇾ **Filter** ⇾ **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Configure LDAP** in the **Installation or Configuration Program** column.

2. If your configuration profile is not open in Configuration Manager, open the profile.
3. Enter property values for the LDAP provider:
   a. Right-click **Configure LDAP** in the profile pane, and select **Edit Selected Task**.
   b. Enter the property values for your LDAP provider, by referring to the values from your worksheet.
4. Optional: (WebSphere and WebLogic only) Click **Test LDAP Connection** to test the connection to the directory service provider by using the directory service bind user name, host name, port number, and password that you provided.
5. Click **File** ⇾ **Save** to save your changes.
6. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, select **Configure LDAP (disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.
7. Apply the LDAP property settings by right-clicking **Configure LDAP** in the profile pane and selecting **Run Task**. Running the configuration task can take several minutes. The task execution status messages are displayed in the console pane below the LDAP properties.
8. Close the Configure LDAP task pane.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Configuring the bootstrap settings:

You must configure the bootstrap settings. The bootstrap information is needed for creating the Global Configuration Data (GCD) and for starting Content Engine.

Be sure that you have available the Installation and Upgrade Worksheet that was filled out during your planning activities.

If your FileNet P8 domain uses multiple non-managed application servers, then complete this procedure only on the initial server. You use a copy of the same EAR file with the bootstrap settings on all the servers.

To edit the bootstrap settings:
1. Open your completed Installation and Upgrade Worksheet file.

Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Configure bootstrap properties** in the **Installation or Configuration Program** column.

2. If your configuration profile is not open in Configuration Manager, open the profile.

3. Right-click **Configure Bootstrap Properties** in the profile pane and select **Edit Selected Task**.

4. In the **Bootstrap operation** field, select **Configure New**.

5. Enter the property values for your database by using the values in your worksheet.

6. Optional: Click **Verify Bootstrap Information** to display the bootstrap information from the modified EAR file.

7. Click **File → Save** to save your changes.

8. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, select **Configure Bootstrap Properties (disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.

9. Apply the bootstrap property settings by right-clicking **Configure bootstrap properties** in the profile pane, and select **Run Task**. Running the configuration task can take a few minutes. The task execution status messages are displayed in the console pane below the bootstrap properties.

10. Close the Configure Bootstrap Properties task pane.

11. If your FileNet P8 domain uses multiple non-managed application servers, then copy the *ce_install_path*/ContentEngine/tools/configure/profiles/*my_profile*/ear/*Engine-app_server_type.ear* file to each additional Content Engine server.

   install_dir
   > The location where Content Engine is installed.

   my_profile
   > The directory for the Configuration Manager profile that you created.

   Engine-app_server_type.ear
   > The EAR file for your application server type: `Engine-ws.ear`, `Engine-wl.ear`, or `Engine-jb.ear`.

   Important: Only run the **Configure bootstrap properties** task on the initial server.

Related concepts

↪ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Configuring instances by using the command line

You can configure a Content Engine instance on a web application server by using
the command-line version of Configuration Manager.

Begin by generating the configuration XML files that define the application server
settings, the Global Configuration Data (GCD) data source JDBC data source
settings, the login modules for the Content Engine application, the directory
service provider (LDAP) settings, and the Content Engine bootstrap settings. Next,
edit the files to provide values for your environment. Then, apply the settings by
executing the tasks from a command prompt.

You must generate, edit, and run a complete set of configuration XML to configure
a Content Engine application. If you are deploying multiple Content Engine
application instances on the same application server, you must 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 complete the configuration tasks by generating all the configuration XML
files before editing, running, or verifying any of them; or you can generate, edit,
run, and verify one file at a time.

1. "Generating the configuration XML files for a Content Engine instance"
   The configuration XML files contain the properties and values for the various
   configuration tasks. From the command line, you can generate all the XML files
   at the same time, or you can generate a single configuration XML file at a time.
2. "Editing the configuration XML files for a Content Engine instance" on page 37
   You must edit each configuration XML file to provide the property values for
   your environment. You can use any text editor to open and edit the files.
3. "Running the configuration XML files" on page 40
   Running the configuration XML files applies the settings. You use the **execute**
   command to apply your configuration settings from the command line.
4. "Checking the completion status of Content Engine configuration tasks" on
   page 43
   Task execution messages are displayed in the console when you run a task, and
   you can view the status of a specific task at any time by running the
   checkStatus command. From the command line, you can check the status of all
   the configuration tasks or check the status of a single task.

Generating the configuration XML files for a Content Engine instance:

The configuration XML files contain the properties and values for the various
configuration tasks. From the command line, you can generate all the XML files at
the same time, or you can generate a single configuration XML file at a time.

The following table lists the configuration XML files that you need to generate to configure a new Content Engine instance.

Table 4. Configuration XML files

File Name	Description
`applicationserver.xml`	Settings for the application server, including the location of the application server software and the name of the server. This file is generated when any other configuration file is generated (either all at the same time or individually) and is used by all of the configuration tasks. If you generate the files one at a time, this file is created only once.
`configurebootstrap.xml`	Settings for creating the Global Configuration Data (GCD) and starting Content Engine.

`configurejdbcgcd.xml`
 	
`configurejdbcgos.xml`
 	
`configurejdbcos.n.xml`, where *n* is an integer starting with 2
	Settings for configuring JDBC connections to the databases used by Content Engine. You must generate, edit, and run the configure the JDBC task once for the data sources for the Global Configuration Data (GCD) and then again for the data sources for each object store.
	
	If you generate a second JDBC configuration file, then it is named `configurejdbc.2.xml`. The file name increments for each new file that you generate.

`configureldap.xml` `configureldap.n.xml`, where *n* is an integer starting with 2	Settings for connecting to and searching within a directory server. If your site uses multiple LDAP providers, you can generate an additional file for each provider. If you generate a second LDAP configuration file, then it is named `configureldap.2.xml`. The file name increments for each new file that you generate.
`configureloginmodules.xml`	Settings for creating the application server login modules for Content Engine.

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

"Generating all of the configuration XML files at the same time"
From the command line, you can generate all of the required configuration XML files for Content Engine at the same time with a single command.

From the command line, you can generate each of the required configuration XML files for Content Engine one file at a time.

Generating all of the configuration XML files at the same time:

From the command line, you can generate all of the required configuration XML files for Content Engine at the same time with a single command.

When you run the **generateconfig** command, all the required configuration XML files are created.

To generate all the configuration XML files at the same time:

1. Log on to the application server as *config_mgr_user*, the user who runs Configuration Manager.

2. Change the current directory to *ce_install_path*/tools/configure, where *ce_install_path* is the location where the Content Engine software is installed.

3. Run the following command. Do not type any line breaks when you enter the command.

   WebSphere Application Server
   ```
   configmgr_cl generateconfig -appserver app_server_type
     -repositorytype ldap_repository_type
     -db db_type -ldap ldap_type
     -bootstrap bootstrap_operation
     -deploy deploy_type -profile myprofile
   ```

   Where:

   -appserver *appserver_name*
   > The **-appserver** *appserver_type* parameter specifies the type of application server and must be WebSphere, WebLogic, or JBoss.

   -repositorytype *ldap_repository_type*
   > WebSphere Application Server only. The **-repositorytype** *ldap_repository_type* parameter specifies the type of LDAP repository to use and must be standalone or federated.

   -db *database_type*
   > The **-db** *database_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the configurejdbcgcd or configurejdbcos option. This parameter specifies the type of database to be used by Content Engine and must be mssql, oracle, oracle_rac, db2, or db2zos.

   -ldap *ldap_type*
   > The **-ldap** *ldap_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the configureldap option. This parameter specifies the type of directory service repository that Content Engine uses for authenticating users and must be activedirectory, adam, edirectory, oid, sunjava, or tivoli. The adam value applies to both Microsoft® ADAM and AD LDS.

   -bootstrap *bootstrap_operation*
   > The **-bootstrap** *bootstrap_operation* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the configurebootstrap option. This parameter specifies the bootstrap operation for the profile and must be new, modify, or upgrade.

   -deploy *deploy_type*
   > The **-deploy** *deploy_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the deployapplication option. This parameter specifies the type of Content Engine deployment. The value must be standard, cluster, or netdeploy (network deployment).
   >
   > Specify standard if you are deploying Content Engine to a stand-alone (that is, a server that is neither managed nor clustered) WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server.
   >
   > Specify cluster if you are deploying Content Engine to a WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server cluster.

Specify `netdeploy` if you are deploying Content Engine to a managedWebSphere Application Server instance. That is, you are using Network Deployment to manage individual servers that are not necessarily in a cluster.

-profile *myprofile*

The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
- The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
- The fully qualified path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg".

For example, the following command generates all the configuration XML files for a new installation profile for a standard deployment on WebSphere with IBM Tivoli® Directory Server that uses a stand-alone LDAP repository and DB2 in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:

```
configmgr_cl generateconfig -appserver WebSphere
 -repositorytype standalone -db db2 -ldap tivoli
 -bootstrap new -deploy standard -profile wstdb2
```

Generating one configuration XML file at a time (WebSphere):

From the command line, you can generate each of the required configuration XML files for Content Engine one file at a time.

You must generate the configuration XML files for each of these required configuration tasks:

- Create the XML file for the bootstrap properties file by using the `configurebootstrap` option.
- Create the XML file for configuring the JDBC Data Sources for the Global Configuration Data (GCD) database by using the `configurejdbcgcd` option.
- Create the XML file for configuring the JDBC Data Sources for a single object store database by using the `configurejdbcos` option.
- Create the XML file for configuring the LDAP provider by using the `configureldap` option.
- WebSphere Application Server and JBoss Application Server only. Create the XML file for configuring the login modules by using the `configureloginmodules` option.

To generate a single configuration XML file:

1. Log on to the application server as *config_mgr_user*, which is the user who will run Configuration Manager.
2. Change the current directory to *ce_install_path*/tools/configure.
3. Run the appropriate command for the task you need. Do not type any line breaks when you enter the command.

   WebSphere Application Server only. Generate the `configurebootstrap.xml` file for the `configurebootstrap` task:
   ```
   configmgr_cl generateconfig -appserver app_server_type
    -repositorytype  ldap_repository_type
    -task configurebootstrap -bootstrap bootstrap_operation
    -profile myprofile
   ```

   Generate the `configurejdbcgcd.xml` file for the `configurejdbcgcd` task for the GCD data sources:
   ```
   configmgr_cl generateconfig -appserver app_server_type
    -repositorytype ldap_repository_type -db db_type
    -task configurejdbcgcd -profile myprofile
   ```

   Generate the `configureldap.xml` file for the `configureldap` task:
   ```
   configmgr_cl generateconfig -appserver app_server_type
    -repositorytype  ldap_repository_type -ldap ldap_type
    -task configureldap -profile myprofile
   ```

   Generate the `configureloginmodules.xml` file for the `configureloginmodules` task:
   ```
   configmgr_cl generateconfig -appserver app_server_type
    -repositorytype  ldap_repository_type -task configureloginmodules
    -profile myprofile
   ```

   where:

   -appserver *appserver_name*
   > **-appserver** *appserver_type* specifies the type of application server and must be WebSphere.

   -repositorytype *ldap_repository_type*
   > WebSphere Application Server only. The **-repositorytype** *ldap_repository_type* parameter specifies the type of LDAP repository to use and must be `standalone` or `federated`.

   -db *database_type*
   > The **-db** *database_type* parameter is required only when you are generating files by using the `configurejdbcgcd` or `configurejdbcos` option. This parameter specifies the type of database to be used by Content Engine and must be `mssql`, `oracle`, `oracle_rac`, `db2`, or `db2zos`.

   -ldap *ldap_type*
   > The **-ldap** *ldap_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the `configureldap` option. This parameter specifies the type of directory service repository that Content Engine uses for authenticating users and must be `activedirectory`, `adam`, `edirectory`, `oid`, `sunjava`, or `tivoli`. The `adam` value applies to both Microsoft ADAM and AD LDS.

   -bootstrap *bootstrap_operation*
   > The **-bootstrap** *bootstrap_operation* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the `configurebootstrap` option. This parameter specifies the bootstrap operation for the profile and must be `new`, `modify`, or `upgrade`.

- *myprofile*

The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
- The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
- The fully qualified path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg".

4. Repeat as needed to generate all the required XML configuration files.

Editing the configuration XML files for a Content Engine instance:

You must edit each configuration XML file to provide the property values for your environment. You can use any text editor to open and edit the files.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To edit values in the configuration XML files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Configuration Manager values, in the **Installation or Configuration Program** column filter by each of the following items, depending on which XML file you are editing:
   - **CM: Set Application Server properties**
   - **CM: Configure GCD JDBC Data Sources**
   - **CM: Configure LDAP**
   - **CM: Configure Bootstrap Properties**

2. Use a text editor or XML editor to open one of the following configuration XML files that you generated:
   - applicationserver.xml
   - configurejdbcgcd.xml for the Global Configuration Data (GCD)
   - configureldap.xml
   - configurebootstrap.xml

   Do not edit the configureloginmodules.xml, configurejdbcos.xml, or deployapplication.xml files at this time.

3. Make the following changes to each XML configuration file:

a. Replace each occurrence of `****INSERT VALUE****` with a value appropriate for your site. See the descriptions in the file for more information.

   Important: You do not need to supply values for passwords. You can run the **storepasswords** command later to add encrypted passwords to the file.

b. Verify that the default values for the remaining properties are correct for your site.

c. Set the **enabled** attribute value in the <configuration> tag to `true` in any configuration XML file you edit if you want to run the configuration task.

4. (WebSphere Application Server only) If you 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 `configurejdbcgcd.xml` file, set the **enabled** attribute value to `false` in the <configuration> tag 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, that are associated with the GCD database.

5. Save your edits and close the XML file.

6. Optional: Add encrypted passwords to the XML files by running the **storepasswords** command.

   a. Enter the following command on one line:

      ```
      configmgr_cl storepasswords [-task task_type | -taskfile task_file_name]
       -profile myprofile
      ```

      where:

      -task *task_type*
         The **-task** *task_type* parameter is optional and specifies a task for which to encrypt passwords. The *task_type* value is not case sensitive. The following table lists each valid task name, its associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 5. task_type values

Task	Configuration file	Description
omitted	`applicationserver.xml` `configurebootstrap.xml` `configurejdbc.xml` `configureldap.xml` `deployapplication.xml`	When you omit the `-task task_type` parameter, you are prompted to enter the passwords for each configuration XML file in the profile. Each password is encrypted before it is added to the XML file.
configurebootstrap	`configurebootstrap.xml`	Encrypts the password for the `BootstrapPassword` property that is used to create the Global Configuration Data and to start Content Engine.

Table 5. task_type values (continued)

Task	Configuration file	Description
configurejdbcgcd	configurejdbcgcd.xml	Encrypts the password for the `DatabasePassword` property that Content Engine uses to access the GCD database.
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Encrypts the password for the `DatabasePassword` property that Content Engine uses to access the object store database. If you have more than one `configurejdbcos.`*n*`.xml` file, run the command for each task file.

-taskfile *task_file_name*
> The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.
>
> If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.
>
> If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the `ce_install_path`/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2"` or opt/IBM/FileNet/ContentEngine/tools/ configure/profiles/ce_was_tiv_db2.
> - The absolute path to the profile input file, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"` or opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg.
>
> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"`.

7. Repeat this procedure as needed until you edit all the required files.

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Running the configuration XML files:

Running the configuration XML files applies the settings. You use the **execute**
command to apply your configuration settings from the command line.

If you need to run the configuration XML files for a profile that was created or
edited in the Configuration Manager graphical user interface, verify that the XML
files contain values for the required passwords before you attempt to run the files.
See "Handling passwords in Configuration Manager" on page 219 for more
information.

To run the configuration XML files:

1. Start or stop the application server.

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

2. Run the configuration XML files, either all at the same time or one file at a
   time.

   "Running all of the configuration XML files at the same time"
   From the command line, you can run all of the required configuration XML
   files for Content Engine at the same time with a single command. Any
   configuration XML file that has the `<TaskEnbled>` element set to false is
   skipped.

   "Running one configuration XML file at a time" on page 41
   From the command line, you can run each of the required configuration XML
   files one file at a time. You must run each of the required files to complete the
   Content Engine configuration.

Running all of the configuration XML files at the same time:

From the command line, you can run all of the required configuration XML files
for Content Engine at the same time with a single command. Any configuration
XML file that has the `<TaskEnbled>` element set to false is skipped.

If your FileNet P8 domain uses multiple non-managed application servers, then
you use a copy of a single EAR file with the bootstrap settings on all the servers.
After you run all the tasks on the initial server, you copy the EAR file with the
bootstrap settings to the additional servers. Then, when you run the tasks for the
additional servers, do not use this procedure to run all tasks at once on the
additional servers. Instead, only run the tasks for configuring the JDBC data
sources for the GCD, configuring the JDBC data sources for the object stores, and
configuring the LDAP provider.

For all other configurations, you can run all the tasks at once on each Content
Engine server.

To run all the configuration XML files at the same time:

1. Change the current directory to *ce_install_path*/tools/configure, where
   ce_install_path is the location where the Content Engine software is installed.
2. Run the following command.

   configmgr_cl execute -profile *myprofile*

   where the **-profile** *myprofile* parameter specifies the profile to use. The *myprofile*
   value can be one of the following items:

   - The name of the profile, such as ce_was_tiv_db2. The profile must be located
     in the *ce_install_path*/tools/configure/profiles directory, where
     ce_install_path is the location where the Content Engine software is installed.
   - The absolute path to the profile directory, such as "C:\Program
     Files\IBM\FileNet\ContentEngine\tools\configure\profiles\
     ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/
     profiles/ce_was_tiv_db2.
   - The absolute path to the profile input file, such as "C:\Program
     Files\IBM\FileNet\ContentEngine\tools\configure\profiles\
     ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/
     tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

   Remember: If the path includes a directory name with spaces, enclose the
   entire path in quotation marks. For example, enter "C:\Program
   Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\
   ce_was_tiv_db2.cfg".

3. If your FileNet P8 domain uses multiple non-managed application servers, then
   copy the *ce_install_path*/ContentEngine/tools/configure/profiles/
   my_profile/ear/*Engine-app_server_type.ear* file to each additional Content
   Engine server.

   install_dir
   > The location where Content Engine is installed.

   my_profile
   > The directory for the Configuration Manager profile that you created.

   Engine-app_server_type.ear
   > The EAR file for your application server type: Engine-ws.ear,
   > Engine-wl.ear, or Engine-jb.ear.

   Important: For each additional non-managed server, only run the tasks for
   configuring the JDBC data sources for the GCD, configuring the JDBC data
   sources for the object stores, and configuring the LDAP provider.

The task execution status messages are displayed.

Running one configuration XML file at a time:

From the command line, you can run each of the required configuration XML files
one file at a time. You must run each of the required files to complete the Content
Engine configuration.

If you are running tasks for a profile that was created or edited in the
Configuration Manager graphical user interface, verify that the XML files contain
values for the required passwords before you attempt to run the configuration
XML files. See "Handling passwords in Configuration Manager" on page 219 for
more information.

You must run the configuration XML files for each of the required configuration tasks:

- Apply the bootstrap properties file by using the `configurebootstrap` option.

  If your FileNet P8 domain uses multiple non-managed application servers, then you use a copy of a single EAR file with the bootstrap settings on all the servers. After you apply the bootstrap properties file on the initial server, you copy the EAR file with the bootstrap settings to the additional servers. Then, when you run the tasks for the additional servers, do not use the `configurebootstrap` option. Only run the tasks for configuring the JDBC data sources for the GCD, configuring the JDBC data sources for the object stores, and configuring the LDAP provider.

  For all other configurations, you can run all the tasks on each Content Engine server.

- Configure the JDBC Data Sources for the Global Configuration Data (GCD) database by using the `configurejdbcgcd` option.
- Configure the JDBC Data Sources for a single object store database by using the `configurejdbcos` option.
- Configure the LDAP provider by using the `configureldap` option.
- Configure the login modules by using the `configureloginmodules` option.

To run the configuration XML files one file at a time:

1. Change the current directory to *ce_install_path*/tools/configure, where *ce_install_path* is the location where Content Engine is installed.
2. Run the appropriate command for the task you need to complete. Do not type any line breaks when you enter the command.

   Run the `configurejdbcos.xml` file in a profile with one `configurejdbcos` task:
   ```
   configmgr_cl execute -task configurejdbcos -profile myprofile
   ```

   Run the `configurejdbc.2.xml` file in a profile with multiple `configurejdbcos` tasks:
   ```
   configmgr_cl execute -taskfile task_file_name
   configurejdbc.2.xml -profile myprofile
   ```

   Run the `configureldap.xml` file in a profile with one `configureldap` task:
   ```
   configmgr_cl execute -task configureldap -profile myprofile
   ```

   Run the `configureldap.3.xml` file in a profile with multiple `configureldap` tasks:
   ```
   configmgr_cl execute -taskfile configureldap.3.xml -profile myprofile
   ```

   Run the `configureloginmodules.xml` file:
   ```
   configmgr_cl execute -task configureloginmodules -profile myprofile
   ```

   Run the `configurebootstrap.xml` file in a profile with one `configurebootstrap` task:
   ```
   configmgr_cl execute -task configurebootstrap -profile myprofile
   ```

   Where:

   -taskfile *task_file_name*
   > The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.
   >
   > If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional, but do not use the **-task** *task_type* parameter with the **-taskfile** *task_file_name* parameter.

If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*
The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
- The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
- The fully qualified path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg".

Remember: The values that you entered for the configurejdbcgcd task were for the Global Configuration Data (GCD) data sources. You must edit and run the configurejdbcos task once for each object store later in the installation process.

3. If your FileNet P8 domain uses multiple non-managed application servers, then copy the *ce_install_path*/ContentEngine/tools/configure/profiles/ *my_profile*/ear/*Engine-app_server_type.ear* file to each additional Content Engine server.

install_dir
The location where Content Engine is installed.

my_profile
The directory for the Configuration Manager profile that you created.

Engine-app_server_type.ear
The EAR file for your application server type: Engine-ws.ear, Engine-wl.ear, or Engine-jb.ear.

Important: For each additional non-managed server, only run the tasks for configuring the JDBC data sources for the GCD, configuring the JDBC data sources for the object stores, and configuring the LDAP provider. When you run the tasks for the additional servers, do not use the configurebootstrap option on the additional servers.

4. Repeat this procedure as needed to run one of the other configuration XML files until you run the files for each of the four required tasks on this server.

The task execution status messages are displayed.

Checking the completion status of Content Engine configuration tasks:

Task execution messages are displayed in the console when you run a task, and you can view the status of a specific task at any time by running the **checkStatus** command. From the command line, you can check the status of all the configuration tasks or check the status of a single task.

The following table lists the status results and their descriptions.

Table 6. checkstatus command results

Status Result	Description
COMPLETED	The task ran successfully.
INCOMPLETE	The task is incomplete.
NO STATUS AVAILABLE	The task has not been run.
FAILED	The task failed to complete. Additional information about the failure is displayed.

Deploying upgraded Content Engine instances

You must deploy the upgraded Content Engine on the web application server. Deploying makes the Content Engine application available for use.

Restriction: Use the command-line version of Configuration Manager if any of these conditions is true:

- You need an accessible software version of Configuration Manager for people with disabilities to use.

  "Deploying upgraded instances by using the graphical user interface"
  You can deploy an upgraded Content Engine Server instance on a web application server by using the graphical user interface version of Configuration Manager. Deploying makes the Content Engine application available for use.

  "Deploying upgraded instances by using the command line" on page 45
  You can deploy an upgraded Content Engine instance on a web application server by using the command-line version of Configuration Manager. Deploying makes the Content Engine application available for use.

Deploying upgraded instances by using the graphical user interface

You can deploy an upgraded Content Engine Server instance on a web application server by using the graphical user interface version of Configuration Manager. Deploying makes the Content Engine application available for use.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

Make sure that the application server is running or stopped, depending on its type:

Table 7. Required Application Server State

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

For more information on the properties and values you set in the Configuration Manager, roll your mouse over the property name to view the tool tip help for the property.

To deploy Content Engine:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data ⇨ Filter ⇨ AutoFilter** command is enabled. To view only Content Engine values for this task, filter by **CM: Deploy Application** in the **Installation or Configuration Program** column.

2. Log on to the application server machine as *config_mgr_user*, the Configuration Manager user.

3. Start Configuration Manager. See "Starting Configuration Manager" on page 25.

4. Select **File ⇨ Open Configuration Profile**, and use the information in your worksheet to specify the values for the parameters in the wizard screens.

5. In the Open Configuration Profile wizard, enter the path to the profile for the Content Engine instance, and then click **Finish**.

6. Right-click the **Deploy Application** task in the profile pane (left pane), and select **Edit Selected Task**.

7. Provide the property values for your deployment by using the values in your worksheet.

8. Select **File ⇨ Save**.

9. Enable the task. By default, the **Deploy Application** task is disabled. When the task is disabled, the task name includes the text **(Disabled)**. To enable the task, select **Deploy Application (Disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.

10. Right-click the **Deploy Application** task in the left pane, and select **Run Task**.

    Running the deploy task might take a few minutes. The task execution status messages are displayed in the Console pane below the deploy application properties.

11. Return to the parent topic and complete the task.

Related concepts

⇨ Accounts for Content Engine upgrade
For details on accounts and required permissions, see *Plan and Prepare Your Environment for IBM FileNet P8*.

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Starting Configuration Manager" on page 25
You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.

Deploying upgraded instances by using the command line

You can deploy an upgraded Content Engine instance on a web application server by using the command-line version of Configuration Manager. Deploying makes the Content Engine application available for use.

Ensure that the application server is running or stopped:

Table 8. Required Application Server State

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

See "Configuration Manager command-line reference" on page 238 for syntax details.

To deploy a Content Engine instance from the command line:

1. Log on to the application server machine as the *config_mgr_user* user, which is the Configuration Manager user.
2. If you upgraded the version of your application server, or if you moved it to a new directory or device and then redeployed Content Engine before upgrading Content Engine to version 4.5.1 from version 4.0.1 or 4.5.0, 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.
3. Set the current directory to *ce_install_path*`/tools/configure`, where *ce_install_path* is the path where you installed Content Engine.
4. Run the following command to apply the settings in the `deployapplication.xml` file:

   `configmgr execute -task deployapplication -profile` *myprofile*

   where the *myprofile* value can be one of the following items:

   - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*`/tools/configure/profiles` directory, where *ce_install_path* is the location where the Content Engine software is installed.
   - The absolute path to the profile directory, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2"` or `opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2`.
   - The absolute path to the profile input file, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"` or `opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg`.

   Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg"`.
5. Optional: Run the following command to check the completion status of the DeployApplication task.

   `configmgr checkstatus -task deployapplication -profile` *myprofile*
6. Repeat this procedure as needed to deploy any additional Content Engine instances.

Related concepts

➥ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Configuration Manager user interface reference" on page 222
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.

Installing storage device source files

If your upgraded IBM FileNet P8 system includes Tivoli Storage Manager or EMC
Centera devices, you must install files on the Content Engine server.

"Installing Tivoli Storage Manager client and adding native API library paths
(WebSphere)"
The client software and the native API library paths enable you to use Tivoli
Storage Manager to access a fixed content device.

"Installing or updating EMC Centera SDK library files" on page 48
EMC Center SDK library files enable access between the IBM FileNet P8
environment and EMC Centera fixed content devices.

Installing Tivoli Storage Manager client and adding native API library paths (WebSphere)

The client software and the native API library paths enable you to use Tivoli
Storage Manager to access a fixed content device.

1. "Installing Tivoli Storage Manager client"
   Install the Tivoli Storage Manager client software on each application server
   where Content Engine is deployed.

2. "Copying the Tivoli Storage Manager API library files to additional servers"
   If you are running a Content Engine server farm, the Tivoli Storage Manager
   API libraries must be on each server in the farm.

3. "Creating a shared library definition for Tivoli Storage Manager native API
   library files" on page 48
   The shared library definition specifies the location of the Tivoli Storage
   Manager JAR file.

Installing Tivoli Storage Manager client:

Install the Tivoli Storage Manager client software on each application server where
Content Engine is deployed.

1. Download the Tivoli Storage Manager client software from the IBM Support
   site at http://www-01.ibm.com/support/docview.wss?uid=swg24019757.

2. Complete the platform-specific installation instructions included with each
   Tivoli Storage Manager client download package. Record the path where you
   install Tivoli Storage Manager client because you must specify the path as the
   DSMI directory when you create a Tivoli fixed content device in FileNet
   Enterprise Manager. If you are installing Tivoli Storage Manager client on
   multiple UNIX hosts, the installation path must be the same on each host. On
   AIX, install Tivoli Storage Manager client at `/usr/tivoli/tsm/client/ba/bin`.

Copying the Tivoli Storage Manager API library files to additional servers:

If you are running a Content Engine server farm, the Tivoli Storage Manager API libraries must be on each server in the farm.

Copy the entire `tsm100`directory structure from the Content Engine installation directory to each of the servers in the farm. It is a best practice to use the same directory structure on each server in the farm. For example:

Option	Description
UNIX	`/opt/IBM/FileNet/ContentEngine/tsm100`

Creating a shared library definition for Tivoli Storage Manager native API library files:

The shared library definition specifies the location of the Tivoli Storage Manager JAR file.

1. Log on to the WebSphere administrative console.
2. Create a shared library definition according to the type of deployment (stand-alone or clustered), the version of WebSphere, and the operating system on which WebSphere runs, as shown in the following substeps. For the detailed procedure on creating the shared library, visit http://www.ibm.com/software/webservers/appserv/was/library/.
   a. Specify a Node scope for the library.
   b. In a server farm, if you installed the `tsm100` directory in different locations, choose Server scope and add a Shared Library entry for each server in your server farm.
   c. Provide a name for the shared library, for example `TSMAPLIB`.

Installing or updating EMC Centera SDK library files

EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.

Update the EMC Centera SDK library files if Centera fixed content devices are in your existing environment or if Centera fixed content devices will be added to your IBM FileNet P8 environment.

1. "Installing EMC Centera SDK version 3.2 library files"
   EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.
2. "Configuring EMC Centera SDK environment variables for version 3.2" on page 49
   Specify values for the environment variables so that Content Engine can access the EMC Centera SDK library files.

Installing EMC Centera SDK version 3.2 library files:

EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.

1. Log on to the Content Engine Server machine as *ce_install_user*.
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`

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

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.

Option	Description
UNIX	`install.sh`

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.

Configuring EMC Centera SDK environment variables for version 3.2:

Specify values for the environment variables so that Content Engine can access the EMC Centera SDK library files.

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:

Option	Description
UNIX	`setCenteraLibPath.sh`

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`

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`

4. Stop and start the application server.

Completing Content Engine post-deployment steps

You must complete several Content Engine post-deployment steps before you can put an IBM FileNet P8 system into production.

Perform the post-deployment procedures for your application server type, and then perform the Content Engine verification procedures.

1. "Completing post-deployment steps (WebSphere Application Server)"
   If you are using IBM WebSphere Application Server federated user repositories, you must specify the primary administrative user for the web application server.
2. "Verifying the Content Engine deployment" on page 51
   You can verify that the Content Engine deployment was successful by accessing the IBM FileNet P8 System Health page.

Completing post-deployment steps (WebSphere Application Server)

If you are using IBM WebSphere Application Server federated user repositories, you must specify the primary administrative user for the web application server.

To set the administrative user

1. Log in to the WebSphere administrative console.
2. Navigate to **Security** → **Secure administration, applications, and infrastructure**, and click **Configure**.
3. If you have not already done so, specify a unique user (short name) for the **Primary administrative user name**.

Important: This name must exist in one of the realms and must be unique.

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

5. Save your changes to the master configuration.

6. If you are upgrading from version 3.5, enable administrative security and application security. The deployment of Content Engine does not enable or check these settings.

   Important: You must disable Java2 security; otherwise, Content Engine will not be able to start or process requests.

7. Restart the web application server where Content Engine is deployed, as follows:
   - For a stand-alone server, stop and start the application server.
   - For a network deployment on WebSphere Application Server Network Deployment, stop and start the application server whereContent Engine is deployed, including the Deployment Manager and all managed nodes.
   - For a cluster, stop and start the cluster.

Verifying the Content Engine deployment

You can verify that the Content Engine deployment was successful by accessing the IBM FileNet P8 System Health page.

1. Verify the state of Content Engine deployment by browsing to the IBM FileNet P8 System Health page:

   `http://`*`server:port`*`/P8CE/Health`

   where:

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

   port is the WSI port used by the Web application server on the machine where Content Engine is deployed.

   The following table lists an example address for your application server:

Table 9. Example FileNet P8 System Health page address

Application Server Type	Web Page Address
IBM WebSphere Application Server	`http://`*`server:`*`9080/P8CE/Health`

2. Verify that the FileNet P8 System Health page contains the Content Engine instance host name and port number. The FileNet P8 System Health page provides status for the items in the following table.

   Tip: At this point, red icons appear on the page at the left of those entries that do not yet exist. In particular, if you are upgrading from version 3.5.2, a red icon appears before the GCD, directory configurations, object stores, storage areas, and sites entries. Also, if you are upgrading from version 3.5.2, no domain name appears yet in the Domain entry.

Table 10. FileNet P8 System Health page contents

Section	Description
Domain	Displays the FileNet P8 domain name if a domain was found.

Table 10. FileNet P8 System Health page contents (continued)

Section	Description
Global Configuration Database	Verifies that the GCD contains a valid domain object, that the XA and non-XA data sources are defined and have unique names, and that the bootstrap user name and password are defined. If any of these verification tests fail, then the failed icon is displayed.
Directory Configurations	Verifies that at least one directory service is configured and lists the number of configured directory service providers.
PE Connection Points	Displays the number of Process Engine connection points.
PE Isolated Regions	Displays the number of Process Engine isolated regions.
Fixed Content Devices	Lists the number of fixed content devices (FCDs) For each FCD listed, there is at least one associated fixed storage area that is in the open state.
Object Stores	Verifies that at least one object store exists and lists the number of object stores.
Storage Areas	Verifies that at least one storage area is defined, lists the number of storage areas, and the status for each storage area, such as online or offline.
Content Cache Areas	Displays the number of content cache areas.
Sites	Verifies that at least one site is defined and lists the number of sites. For each site listed, at least one virtual server or server instance exists.

3. Optional: Bookmark the FileNet P8 System Health page Web page address in your browser for later use.

4. Browse to the FileNet Engine Web page:

   `http://server:port/FileNet/Engine`

   where:

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

   port is the HTTP port used by the application server where Content Engine is deployed.

   The following table lists an example address for your application server:

Table 11. Example FileNet Engine Web page address

Application Server Type	Web Page Address
IBM WebSphere Application Server	http://myserver:9080/FileNet/Engine

5. Verify that the FileNet Engine Web page contains the following information:
   - Verify the value in the **Startup Message** key. The Content Engine build and version (for example, dap451.097), must match the build and version in the `ce_install_path`/ce_version.txt file, where `ce_install_path` is the location where the Content Engine software is installed.
   - Verify the value in the **Process Engine** key. The Process Engine Client build and version (for example, pui451.075), must match the build and version in the version.txt file in the TAR or ZIP file of the Process Engine Client installation software.
   - Verify that the values for JDBC driver, server instance, operating system, and other properties match the values that you entered when you configured Content Engine.

6. Optional: Bookmark the FileNet Engine Web page in your browser for later use.

Installing the latest Content Engine client files on other IBM FileNet P8 servers (for staged upgrades)

When you upgrade Content Engine 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:

- "Upgrading the Content Engine client files on Process Engine servers" on page 148
- "Installing the latest Content Engine client files on Application Engine servers" on page 169

Important:

- You must install the Content Engine client files for any expansion products that use Content Engine Client files, such as:
  - IBM FileNet Workplace XT
  - IBM InfoSphere Enterprise Records
  - IBM FileNet Business Process Framework
- You must redeploy any components that are web-application-server-based, such as Application Engine, Workplace XT, and InfoSphere Enterprise Records, after you upgrade the Content Engine Client files on the associated machines.

Establishing the FileNet P8 domain and Global Configuration Data

You must use IBM FileNet Enterprise Manager to create a FileNet P8 domain and the Global Configuration Data (GCD).

Complete the procedures for the FileNet P8 domain and the Global Configuration Data (GCD).

1. "Creating the FileNet P8 domain"
   You use IBM FileNet Enterprise Manager to create the FileNet P8 domain.
2. "Configuring directory service authentication" on page 54
   You must configure the directory service authentication for the domain. To set up multi-realm authorization, you must run the Create a Directory Configuration wizard one time for each realm.
3. "Configuring permissions for a FileNet P8 domain" on page 55
   You must configure the permissions for the FileNet P8 domain. To complete this configuration, you add your authorized users and groups to the list of Global Configuration Data (GCD) administrators.
4. "Verifying the trace log file location" on page 55
   You must verify that trace log file location after you deploy Content Engine.

Creating the FileNet P8 domain

You use IBM FileNet Enterprise Manager to create the FileNet P8 domain.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

Important: 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.

To create a FileNet P8 domain:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values for this task, filter by **EM: Create FileNet P8 domain** in the **Installation or Configuration Program** column.

2. Start Enterprise Manager by double-clicking the **Enterprise Manager** icon on the desktop, or by choosing **Start** → **All Programs** → **FileNet P8 Platform** → **Enterprise Manager**.

3. Click **Add** and complete the configuration screens by using the values in your worksheet. The initial values displayed in the screens are default values that you must change to match your site.

Related concepts

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring directory service authentication

You must configure the directory service authentication for the domain. To set up multi-realm authorization, you must run the Create a Directory Configuration wizard one time for each realm.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

See these FileNet P8 help topics for additional information:

- **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **How to** → **Configure for multiple realms**
- **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory service providers**.

To configure directory service authentication:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values for this task, filter by **EM: Create a Directory Configuration** in the **Installation or Configuration Program** column.

2. In the FileNet P8 dialog box, click **Connect**.

3. In the Create P8 Domain screen, enter the name for a new FileNet P8 domain and click **Continue** to start the Create a Directory Configuration wizard.

   If an error occurs, check the application server log file on the machine where Content Engine is deployed:

Application server type	Path to log file
IBM WebSphere Application Server	*WAS_install_path*/AppServer/profiles/ *profile_name*/logs/server_name/ SystemOut.log

4. Complete the configuration screens by using the values in your worksheet.
5. 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 if you are using Enterprise Manager for the first time and have not defined any Global Configuration Data (GCD) administrators.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring permissions for a FileNet P8 domain

You must configure the permissions for the FileNet P8 domain. To complete this configuration, you add your authorized users and groups to the list of Global Configuration Data (GCD) administrators.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To configure permissions for a a FileNet P8 domain:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **EM: Configure New Domain Permissions** in the **Installation or Configuration Program** column.

2. Complete the configuration screens by using the values in your worksheet.
3. In the Configure New Domain Permissions message box, click **OK**.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Verifying the trace log file location

You must verify that trace log file location after you deploy Content Engine.

Perform the following steps to check the path to the trace log file, p8_server_trace.log.

1. Start IBM FileNet Enterprise Manager by double-clicking the **Enterprise Manager** icon on the desktop, or by choosing **Start** → **All Programs** → **FileNet P8 Platform** → **Enterprise Manager**.
2. Navigate to the level (domain, site, virtual server, or server instance) at which trace logging is configured, right-click the node at that level, and choose **Properties**.
3. Click the **Trace Control** tab.
4. If the current value of the **Log File Output Location** parameter is not consistent with the version of the application server, then edit the value as needed.

The following table shows default locations, at the server instance level, for the trace log file, starting at *install_root*, the installation path for the application server.

Application Server Type	Path to Trace Log File
IBM WebSphere Application Server	*install_root* /profiles/*profile_name*/ FileNet/*server_name*/ p8_server_trace.log

5. Click **OK**.

Configuring storage devices for upgrades

When you upgrade from Content Engine version 3.5.2, you must configure your existing file store devices.

Complete the procedure for the devices in your environment.

"Configuring storage devices for migration to UNIX"
When you upgrade from Content Engine version 3.5.2 and migrate to a UNIX server, you must prepare the UNIX server to mount your existing file store devices.

"Preparing NetApp SnapLock volumes for upgrade"
If your existing IBM FileNet P8 domain includes a SnapLock fixed content device, you must allow access to your fixed content via NFS instead of, or in addition to, CIFS.

Configuring storage devices for migration to UNIX

When you upgrade from Content Engine version 3.5.2 and migrate to a UNIX server, you must prepare the UNIX server to mount your existing file store devices.

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 3.5.2 file store devices.

- If you are migrating your Content Engine version 3.5.2 file storage devices to UNIX, you must copy your existing storage directories and files to the UNIX machine and ensure that the *ce_os_user* account has permissions for the new directories and folders.

- If the version 3.5.2 file storage areas remain on the Windows NTFS devices, or if you intend to use the Content Engine Upgrader tool rather than Content Engine Web Upgrade Tool to upgrade your object stores (regardless of whether you leave the storage directories on NTFS or the UNIX file system), create an NFS mount from the Content Engine version 4.5.1 server to the NTFS mount.

Related tasks
"Installing or upgrading Content Engine" on page 17
You can install or upgrade Content Engine, Configuration Manager, .NET Clients (including IBM FileNet Enterprise Manager), Content Engine Upgrader, and FileNet Deployment Manager either interactively or silently.

Preparing NetApp SnapLock volumes for upgrade

If your existing IBM FileNet P8 domain includes a SnapLock fixed content device, you must 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 4.5.1 of IBM FileNet P8.

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.

Because all pre-4.0.0 versions of IBM FileNet P8 support only CIFS, all NetApp storage volumes used by Content Engine version 3.5.2 use NTFS security style. You must change the security style 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* to UNIX:

```
 telnet NAFiler
qtree security /vol/vol1/sa1 unix
```

After the qtree command executes, all files created by Content Engine on a UNIX platform will have UNIX security attributes.

3. Map UNIX users and groups to Window equivalents.

Because Content Engine version 3.5.2 support CIFS, rather than NFS, all existing files on a NetApp volume have NTFS security attributes. The users and groups with access rights to these files are defined by Windows.

To allow the version 3.5.2 files to remain accessible, you must create a mapping between the new UNIX account for version 4.5.1 of Content Engine and the old Windows account for version 3.5.2.

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.
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 might 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 3.5.2. *x* of Content Engine Server executes) and FNCE_UNIX_User (the UNIX user account under which version 4.5.1 of Content Engine Server executes).

a. Log on to the machine where version 3.5.2 of the Content Engine is installed.

b. Connect to the root volume on the target NetApp filer by 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:
   ```

4. Return to the parent topic and complete the task.

Related information

➡ Product documentation for IBM FileNet P8 Platform
Download the IBM FIleNet P8 Hardware and Software Requirements, as well as
PDF's of the rest of the IBM FileNet P8 Platform documentation.

Creating the data sources for an object store

Each object store in your site requires its own JDBC data sources. You can
configure the data sources by using the Configuration Manager graphical user
interface or the command line.

You created the JDBC data source for the Global Configuration Data (GCD) earlier
in the configuration process. You must add the data sources now for at least one
object store that you want to upgrade. You can create additional object store data
sources as needed either immediately after you create the initial object store data
sources or after you have confirmed your Content Engine configuration.

Restriction: Configuration Manager cannot create a data source that uses the same
name as an existing data source. If you manually configured a data source by
using your web application server administration tools, and you want to reuse the
name, then you must manually delete the existing data source before you attempt
to create the data source in Configuration Manager. See your web application
server documentation for more information.

"Creating the initial object store data sources by using the graphical user
interface"
You must create the JDBC data sources for each object store that Content
Engine uses. Content Engine uses the data source information to connect to and
update the object store database. By default, each profile includes a task for
creating the data sources for a single object store.

"Creating additional object store data sources by using the graphical user
interface" on page 60
You must create the JDBC data sources for each object store that Content
Engine uses. Content Engine uses the data source information to connect to and
update the object store database. By default, each profile includes a task for
creating the data sources for a single object store, and you can add additional
tasks as needed for each additional object store.

"Creating the data sources by using the command line" on page 61
You must create the JDBC data sources for each object store that Content
Engine uses. You must generate, edit, and execute a new `configurejdbc-os.xml`
file for each object store in your environment.

Creating the initial object store data sources by using the graphical user interface

You must create the JDBC data sources for each object store that Content Engine
uses. Content Engine uses the data source information to connect to and update
the object store database. By default, each profile includes a task for creating the
data sources for a single object store.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ‣ **Filter** ‣ **AutoFilter** command
   is enabled. To view only Configuration Manager values for this task, filter by

CM: Configure Object Store JDBC Data Sources (object store 1) in the **Installation or Configuration Program** column.

2. Start or stop the application server, depending on its type.

Table 12. Required application server state

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

3. If your configuration profile is not open in Configuration Manager, open it by selecting **File** → **Open** and navigating to your *profilename*.cfgp file.

4. Right-click **Configure Object Store JDBC Data Sources** in the profile pane, and select **Edit Selected Task**.

5. Enter the property values for your data sources by using the values in your installation worksheet.

6. Select **File** → **Save** to save your changes.

7. Optional: WebSphere and WebLogic only. Click **Test Database Connection** to test the connection to the database by using the database user name, database server name, database name, port number, and password that you provided. The test does not create the data sources.

8. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, right-click **Configure Object Store JDBC Data Sources (disabled)** in the profile pane, and choose **Enable Selected Task** from the context menu.

9. Apply the JDBC property settings by right-clicking **Configure Object Store JDBC Data Sources** in the profile pane, and selecting **Run Task**. Running the configuration task can take several minutes. The task execution status messages are displayed in the Console pane below the bootstrap properties.

Creating additional object store data sources by using the graphical user interface

You must create the JDBC data sources for each object store that Content Engine uses. Content Engine uses the data source information to connect to and update the object store database. By default, each profile includes a task for creating the data sources for a single object store, and you can add additional tasks as needed for each additional object store.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

The default configuration profile contains one task for configuring a single object store. You must create and run a new task for each additional object store in that Content Engine uses.

To configure the JDBC settings for additional object stores:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CM: Configure JDBC data sources (object store 2)** in the **Installation or Configuration Program** column.

2. Start or stop the application server, depending on its type.

Table 13. Required application server state

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

3. If your configuration profile is not open in Configuration Manager, open it by select **File** → **Open** and navigating to your *profilename*.cfgp file.

4. Create a task for the additional object store.

   a. Right-click any *task name* in the profile pane, such as **Configure Object Store JDBC Data Sources**, and select **Add New Task** → **Configure JDBC Data Sources**. The **New_Configure JDBC Data Sources** task is added to the profile.

   b. Right-click **New_Configure JDBC Data Sources**, and select **Rename Task**.

   c. Enter a useful name for the task, such as **Configure Object Store 2 JDBC Data Sources**.

   d. Click **OK**.

5. Right-click the **Configure Object Store 2 JDBC Data Sources** task in the profile pane, and select **Edit Selected Task**.

6. Provide the property values for the JDBC data sources for the object store by using the values in your worksheet for your additional object store.

7. Select **File** → **Save** to save your changes.

8. Optional: WebSphere and WebLogic only. Click **Test Database Connection** to test the connection to the database by using the database user name, database server name, database name, port number, and password that you provided. The test does not create the data sources.

9. Ensure that the task is enabled. When the task is disabled, the task name includes the text **(disabled)**. To enable the task, right-click **Configure Object Store 2 JDBC Data Sources (disabled)** in the profile pane, and choose **Enable Selected Task** from the context menu.

10. Apply the JDBC property settings by right-clicking **Configure Object Store 2 JDBC Data Sources** in the profile pane, and selecting **Run Task**. Running the configuration task can take a few minutes. The task execution status messages are displayed in the Console pane below the bootstrap properties.

11. Optional: Repeat this procedure as needed to add the data sources for each object store in your environment.

Creating the data sources by using the command line

You must create the JDBC data sources for each object store that Content Engine uses. You must generate, edit, and execute a new configurejdbc-os.xml file for each object store in your environment.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

If you generated all of the configuration profiles at the same time, you have already created an initial configurejdbc-os.xml file for the initial object store data sources. If you have an existing configurejdbc-os.xml file, you can generate another file for each additional object store. Each additional file that you add is named configurejdbc-os.*n*.xml. You can generate multiple configurejdbc-os.*n*.xml files as needed, depending on the number of object stores in your environment. For best results, create an additional file for each object store.

1. Open your completed Installation and Upgrade Worksheet file.

Tip: In the worksheet file, verify that the **Data ⟶ Filter ⟶ AutoFilter** command is enabled. To view only Configuration Manager values for this task, filter by **CM: Configure Object Store JDBC Data sources (object store 1)** in the **Installation or Configuration Program** column.

2. Log on to the application server as *config_mgr_user*, the user who runs Configuration Manager.

3. If you did not generate all the configuration files at the same time or if you must create another file, generate the `configurejdbc-os.`*n*`.xml` file by running the following command:

```
configmgr_cl generateConfig -appserver app_server_type -db  db_type
  -task configurejdbcos -profile myprofile
```

where:

-appserver *appserver_name*
> The **-appserver** *appserver_type* specifies the type of application server and must be WebSphere, WebLogic, or JBoss.

-db *database_type*
> The **-db** *database_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the `configurejdbcgcd` or `configurejdbcos` option. This parameter specifies the type of database to be used by Content Engine and must be `mssql`, `oracle`, `oracle_rac`, `db2`, or `db2zos`.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*`/tools/configure/profiles` directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2"` or `opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2`.
> - The absolute path to the profile input file, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"` or `opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg`.
>
> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"`.

4. Use a text editor to open the `configurejdbc-os.`*n*`.xml` file and edit it as follows:
   a. Provide the entries that are required for your environment by using the values in your worksheet.
   b. Replace each occurrence of `****INSERT VALUE****` with a value appropriate for your site. See the descriptions in the file for more information.
   c. Verify that the default values for the remaining properties are correct for your site.
   d. Set the **enabled** attribute value in the **<configuration>** tag to `true` so that you can run the configuration task in 6 on page 63.

e. Save your edits.

5. Run the **storepasswords** command to encrypt and store the required passwords.

   `configmgr_cl storepasswords -profile myprofile`

6. Run the following command to execute the `configurejdbc.n.xml` file:

   `configmgr_cl execute -task ConfigureJDBC -profile myprofile`

7. Optional: Check the completion status by running the following command:

   `configmgr_cl checkStatus -task configurejdbcos -task configurejdbc.n.xml -profile myprofile`

   Where `configurejdbc.n.xml` is the task file for the object store.

8. Repeat 3 on page 62 through 6 as needed for each additional object store.

Upgrading Content Search Engine software from version 3.5.2

If you used the full-text search feature (CBR) in version 3.5.2, you upgrade by installing and configuring version 4.5.1 of Content Search Engine (Autonomy K2). After the upgrade, you recreate index areas, called *K2 collections*, to replace the version 3.5.2 content search indexes.

Important: If version 4.5.1 of Content Search Engine is already part of your existing version 4.5.1 Content Engine system, omit the tasks in this section.

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 3.5.2 to 4.5.1, you can install and configure the Autonomy K2 product before or after you upgrade Content Engine. After upgrading Content Search Engine, you create new index areas, and specify classes and property templates for those classes (if so desired). Then you run a classes index job, which is part of the upgrade of Content Engine.

1. "Production environment upgrades" on page 64
   Upgrade of Content Search Engine and related Content Engine data must occur as an in-place upgrade. As such, the data is migrated to version 4.5.1 and is no longer usable in the previous version of IBM FileNet P8.

2. "Content Search Engine software upgrade" on page 64
   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.

3. "Upgrading Content Search Engine (UNIX)" on page 65
   You must install the K2 Master Administration Server software first. The Master Administration Server dashboard is the central control point for configuring Content Search Engine for single-server or multiserver installations.

4. "Updating the style file of user fields" on page 68
   You need to edit the Autonomy K2 style file of user fields to add a data field definition for storing the partitioning value of a document.

5. "Configuring Content Search Engine" on page 68
   You can configure services required on the K2 Master Administration Server, and on additional Administration Servers you might install for IBM FileNet P8 Content Search Engine. All servers are configured through the Master Administration Server Dashboard.

6. "Configuring Autonomy K2 for version 4.5.1 index areas" on page 72
   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.

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 version 4.5.1 and is no longer
usable in the previous version of IBM FileNet P8.

To perform an upgrade in a migration scenario, create a duplicate of your FileNet
P8 system by performing a backup and then a restore. Using the restored system,
complete the migration from version 3.5.2 to 4.5.1. 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 stand-alone 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.

Important: 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).

- Autonomy K2 was previously known as Verity, and you will see Verity still used
  in many of the interfaces described in the following procedures.
- If your Content Engine runs on Windows, then the machines in your Content
  Search Engine configuration must run on Windows. If your Content Engine runs
  on UNIX, then the machines in your Content Search Engine configuration must
  run on UNIX, but it to doesn't have to be the same UNIX 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, 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 place the `style.stp` file into the `stylefiles` directory to list words you do not want full-text indexed, for example, short words such as a, the, or and. However, any word in the stop word file cannot be searched. See the K2 documentation for more details. To create a stop word file you can typically copy the `vdk30.stp` file from either the main K2 install directory or the foreign language locales package over to the main `stylefile` directory, and then rename the file `style.stp`. You must do this copy operation before you create collections.

  Because indexing is case sensitive, the `style.stp` file should include lower-case and upper-case versions of words in the stop word list. For example, use "the", "The" and "THE" if you think that all of these might be encountered in a search.

Upgrading Content Search Engine (UNIX)

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

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install interactively on UNIX, X-Terminal must be installed.

You must designate each installation as either a Master Administration Server software installation or an Administration Server software installation. Only one K2 Master Administration Server can be installed for each Content Engine domain.

Remember: Most processes for the Autonomy K2 software will run as *k2_os_user*. However, the vspget process must run as root.

To install Content Search Engine on UNIX:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Search Engine values, filter by **CSE Installer** in the **Installation or Configuration Program** column.

2. Log on to the machine as *k2_os_user*.

3. Set the following environment variable and place the entry in the `.profile` file for *k2_os_user*.

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

4. Start the interactive or silent Content Search Engine installation:

1. "Upgrading Content Search Engine interactively (UNIX)" on page 66
   To upgrade Content Search Engine interactively, you must run an installation program wizard.

2. "Upgrading Content Search Engine silently (UNIX)" on page 66
   When you upgrade Content Search Engine silently, you run a command in the Content Search Engine software package that references parameters in an installation text file.

3. "Setting environment variables for Content Search Engine (UNIX)" on page 67
   Content Search Engine requires environment variables to be added to the
   .profile file of the Autonomy K2 operating system user.
4. "Starting or stopping the autonomy K2 services on UNIX" on page 67
   You can manually start or stop the Autonomy K2 services by issuing different
   commands according to your environment.

Related concepts

➡ Content Search Engine accounts
For details on accounts and required permissions, see "Creating accounts for
Content Search Engine" in *Plan and Prepare Your Environment for IBM FileNet P8*.

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Content Search Engine interactively (UNIX):

To upgrade Content Search Engine interactively, you must run an installation
program wizard.

An X terminal must be installed on the Content Search Engine server to run the
interactive installation program.

To upgrade Content Search Engine:
1. Run the installation program:

Platform	Command
AIX	P8CSE-4.5.1.0-AIX.BIN

2. Complete the installation screens using the values in your worksheet.
3. Review the following Content Search Engine log files for installation errors.
   The files are located in the verity directory of the installation path:
   - *cse_install_path* /verity/cse_install_log_4_5_1.txt
   - *cse_install_path* /verity/vconfig.log
4. Access the host machine and log on as a user with root privileges.
5. Enter the following commands to set the setuid bit for the vspget program so
   that the service runs as root:
   chown root *verity_install_path*/k2/*_platform*/bin/vspget
   chmod u+s *verity_install_path*/k2/*_platform*/bin/vspget
   Replace *_platform* with the following directory:

Platform	Directory
AIX	_rs6k43

Upgrading Content Search Engine silently (UNIX):

When you upgrade Content Search Engine silently, you run a command in the
Content Search Engine software package that references parameters in an
installation text file.

To upgrade Content Search Engine silently:
1. Access the IBM FileNet Content Search Engine installation package.

2. Edit the `CSE_silent_install.txt` file for your installation using the values from your worksheet.
3. Run the following command: Run the following command:

Platform	Command
AIX	`P8CSE-4.5.1.0-AIX.BIN -f` `CSE silent install.txt -i silent`

4. Review the following Content Search Engine log files for installation errors. The files are located in the `verity` directory of the installation path:
   - *cse_install_path* `/verity/cse_install_log_4_5_1.txt`
   - *cse_install_path* `/verity/vconfig.log`
5. Access the host machine and log on as a user with root privileges.
6. Enter the following commands to set the setuid bit for the vspget program so that the service runs as root:

   `chown root` *verity_install_path*`/k2/`*_platform*`/bin/vspget`

   `chmod u+s` *verity_install_path*`/k2/`*_platform*`/bin/vspget`

   Replace *_platform* with the following directory:

Platform	Directory
AIX	_rs6k43

Setting environment variables for Content Search Engine (UNIX):

Content Search Engine requires environment variables to be added to the `.profile` file of the Autonomy K2 operating system user.

To set environment variables:

In the `.profile` file for *k2_os_user*, set the `PATH` , `SHLIB_PATH`, `LIBPATH`, and `LD_LIBRARY_PATH` environment variables (depending on your operating system), each on its own line, without line breaks.

Important: If your `.profile` file has a line containing `set -u` that appears before the lines involving `PATH` , `SHLIB_PATH`, `LIBPATH`, or `LD_LIBRARY_PATH`, you must do one of the following edits before setting the environment variables:
- Remove the line containing `set -u` from the `.profile` file.
- Move the line containing `set -u` to a point in the `.profile` file below the lines involving `PATH`, `SHLIB_PATH`, `LIBPATH`, or `LD_LIBRARY_PATH`.

AIX

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

Starting or stopping the autonomy K2 services on UNIX:

You can manually start or stop the Autonomy K2 services by issuing different commands according to your environment.

To manually start or stop the Autonomy K2 services, use the following commands, according to your environment. (Each command is to be typed on one line.)

Option	Description
AIX	Start services:
	`nohup` *`verity_install_path`*`/k2/` `_rs6k43/bin/k2adminstart &`
	Stop services:
	`verity_install_path``/k2/` `_rs6k43/bin/k2adminstop`

Updating the style file of user fields

You need to edit the Autonomy K2 style file of user fields to add a data field definition for storing the partitioning value of a document.

To edit the style file:

1. Navigate to the directory in the Content Search Engine software that contains the `style.ufl` file:

   UNIX *`verity_install_path`*`/data/stylesets/FileNet_FileSystem_PushAPI`

2. Open the Autonomy K2 style file of user fields, `style.ufl`, for editing.

3. If it is not already present, add the CEMP definition at the bottom of the file:

   ```
   data-table: fce
   {
   fixwidth: cemp_partition_date 4 unsigned-integer
   }
   ```

4. Save your edits and close the file.

Configuring Content Search Engine

You can configure services required on the K2 Master Administration Server, and on additional Administration Servers you might install for IBM FileNet P8 Content Search Engine. All servers are configured through the Master Administration Server Dashboard.

- When you name the servers that you create, it is a best practice to indicate the type of server in the server name. Otherwise, when you configure Content Engine by using Enterprise Manager, it can be difficult to determine which server to select for an action. For example, use *`server_name`*`_broker` to indicate 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.

  – 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.

- In the following procedure, import the IBM FileNet Styleset only once, regardless how many K2 Administration Server machines you install.

To configure Content Search Engine:

1. (Optional) 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: *server_name* _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, *server_name* _broker.

g. Click **Finish**.

7. Import the IBM FileNet Styleset.

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 the server to which you will import the Styleset.
- Source Path:

  UNIX *verity_install_path*/data/stylesets/
  FileNet_FileSystem_PushAPI

e. Click **Import**.

Tip: 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 used by Content Search Engine.

8. Navigate to the source path you specified in the preceding step and complete the following substeps:

a. Open the file `style.ufl` for editing.

b. If it is not already present, add the CEMP definition at the bottom of the file:

```
data-table: fce
{
fixwidth: cemp_partition_date 4 unsigned-integer
}
```

c. Save and close the file.

9. 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

Restriction: LDAP Ticket Servers are not currently supported.

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**.

10. 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.
- Password: Enter the authentication password.

d. Click **Modify**.

The K2 software 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.

11. Launch the K2 Dashboard and log on.

12. 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 choose **Full Restart**.

b. 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.

13. 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**.

14. 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 of the page (bottom right).

d. Click **K2 Servers** 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.

   Important: If you have a multiserver 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 **K2 Brokers** 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.

15. 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 choose **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.

1. "Configuring 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.

Related concepts

⟶ Content Search Engine accounts
For details on accounts and required permissions, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring 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.

Configuring Autonomy K2 for version 4.5.1 index areas

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

1. "Collections directory requirements for Content Search Engine"
   The task of creating a collections directory includes setting up security permissions to allow read and write access.
2. "Completing the upgrade" on page 74
   Completing the upgrade of Content Search Engine is part of the process for upgrading Content Engine.

Related concepts

⟶ Content Search Engine accounts
For details on accounts and required permissions, see *Plan and Prepare Your Environment for IBM FileNet P8*.

⟶ Installation and Upgrade Worksheet
For more information about parameter values, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Collections directory requirements for Content Search Engine:

The task of creating a collections directory includes setting up security permissions to allow read and write access.

Important: 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 using the network (NFS, PCNFS, or CIFS) causes stability problems under load and corrupt your collections. Any existing configurations that contain non-local collections directories must be reconfigured.

To create a collections directory:

1. Create a directory in which to 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*.
   This path must be local to the index server that is assigned to write collections.
2. Create a temp directory (VerityIndexArea.TempDirectoryPath) which is used by the K2 Index Server and Content Engine Server during operations.
   Content Engine and the Autonomy K2 servers, the K2 Administration Server and Content Engine must be on the same machine or 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.

Option	Description
UNIX	1. Open the following K2 configuration file in a text editor (default path shown): *verity_install_path*/k2/common/ verity.cfg
	2. Modify the next available alias settings by adding the collections path. For example, change alias6, mapping6, and dirmode6 to the following settings:
	alias6=path1 mapping6=/*Collections_Directory_Path* dirmode6=wr
	The Collections_Directory_Path must be a local path and not a mount point.

4. Set file store access. Each file store that is full text indexed must be accessible by the Autonomy K2 server that performs the full text indexing. Permissions on the file store must be set the same as the permissions on the collections directories. Having the same permissions allows 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.

Although we are mapping remote collections, which are not supported, indexing is not run on these indexes. The mapping is used to close the collections.

Completing the upgrade:

Completing the upgrade of Content Search Engine is part of the process for upgrading Content Engine.

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.

Updating the CFS-IS settings

If Content Federation Services for FileNet Image Services (CFS-IS) is in your version 3.5.2 system, then you must enable the FileNet Image Services dispatcher. You must also modify the default map between CFS-IS and the object stores in the FileNet P8 domain.

To update the CFS-IS settings:

1. Start Enterprise Manager if it is not already running.
2. Right-click the Enterprise Manager root node in the navigation tree (left pane), and se;ect **Properties**.
3. Click the **IS Import Agent** tab, select the **Enable Dispatcher** check box, and then click **OK**.
4. Start the FileNet Image Services Remote Admin Console by selecting **Start** ➔ **Programs** ➔ **FileNet P8 Platform** ➔ **IS Remote Admin Console**.
5. For each object store that participates in CFS-IS federation, complete the following substeps:

a. Select **Applications** → **IS Catalog Export Tool**. The Default Mappings window displays the FileNet Image Services document class that is mapped to a FileNet P8 domain/object store combination.

b. Select the *document class name* and click **Remove**, and then click **OK**.

c. Modify the f_ce_domain_guid and f_ce_domain_name columns in the ce_id_map table on the IS database to match the new Content Engine version 4.5.1 domain GUID and domain name. For example:

```
UPDATE ce_id_map SET f_ce_domain_guid = 'new_domain_guid' ,
f_ce_domain_name = 'new_domain_name'
where f_ce_os_name = 'object_store_name';
```

Use Enterprise Manager to determine the *new_domain_guid* and *new_domain_name* values. Right-click the *P8 Domain node*, and select **Properties**. The ID property is the domain GUID.

6. Restart FileNet Image Services:

- To restart FileNet Image Services by using the graphical user interface, start FileNet FileNet Image Services Task Manager and click **Restart**.

- To restart FileNet Image Services from a command line, run the following commands:

```
initfnsw -y stop
killfnsw -DyAS
initfnsw start
```

7. For each object store that participates in CFS-IS federation, configure the FileNet Image Services documentation class for export to Content Engine.

a. On the main Remote Admin Console screen, select **Applications** → **IS Catalog Export Tool**.

b. From the **CE Configured IS Doc Classes** menu, select the IS document class whose index properties you want to export to Content Engine.

c. From the Available CE Object Store Mappings pane, select the Content Engine domain and object store that you want to export to Content Engine.

d. Click **Set Default Map**.

8. Stop and start Content Engine.

9. Verify your property mappings.

a. Start Enterprise Manager if it is not already running.

b. Right-click the *object store node* in the navigation tree (left pane), and select **Properties**.

c. Click the **Content Federation** tab.

d. Select your *class mapping name* and click **Add/Edit**.

e. Select **Current Class Mappings**, and click **Map Properties**.

f. Verify your property mappings.

g. Click **OK**.

h. Close Enterprise Manager.

Upgrading Content Engine data

You must upgrade your current Content Engine data (GCD, addons, fixed content devices, object stores, and related information).

There are two tools available to you for upgrading Content Engine data: the IBM FileNet Content Engine Upgrader tool and IBM FileNet Content Engine Upgrader Web Tool. The Content Engine Upgrader tool has both a graphical user interface (GUI) and a command-line interface (CLI). Content Engine Web Upgrade Tool has

only a graphical user interface. Choose a tool for upgrading Content Engine data based upon the criteria in the following table:

Interface to be used	Current state of the object stores	Upgrade tool to use
GUI	All object stores are at version 3.5.2.	Content Engine Web Upgrade Tool or Content Engine Upgrader tool
GUI	Some object stores at version 4.5.1 were upgraded using Content Engine Upgrader tool; other object stores in the same domain are still at version 3.5.2.	Content Engine Upgrader tool
GUI	Some object stores at version 4.5.1 were upgraded using Content Engine Web Upgrade Tool; other object stores in the same domain are still at version 3.5.2	Content Engine Web Upgrade Tool
CLI	Not applicable	Content Engine Upgrader tool

Whichever tool you use, keep in mind the following guidelines as you prepare to upgrade the Content Engine data:

- It is not necessary to check in checked-out documents before running either upgrade tool.
- If the maximum heap size of the JVM is 1 GB or more, do not run an upgrade tool on the same machine where any other major application (such as a database) is running unless the machine has at least 2 GB of RAM.
- If you use Content Engine Web Upgrade Tool to upgrade your Content Engine data, you must deploy it on the same Java Virtual Machine where version 4.5.1 of Content Engine is deployed.
- If you are upgrading an IBM FileNet P8 environment with many object stores, use the following approach:
  - Upgrade fewer than 20 object stores at a time.
  - For best performance when Content Engine is deployed on a web application server with a 32-bit JVM, restrict each version 4.5.1 FileNet P8 domain to 50 or fewer object stores. If the version 3.5.2 FileNet P8 domains contain more than 50 object stores, consider partitioning them into multiple 4.5.1 FileNet P8 domains during the upgrade process.
  - Ensure that the version 3.5.2 object stores with the same basic set of system objects (for file stores, fixed file stores, and related objects) are upgraded into the same 4.5.1 FileNet P8 domain.
- The time to upgrade object stores from version 3.5.2 depends on the number of documents the object store contains, the database type, and the resources of the machine on which the database system runs. If you are upgrading version 3.5.2 object stores that contain more than a million documents or a million custom objects, the Upgrader tool can take more than an hour to complete the upgrade.
- For non-English support relating to collections, see *IBM FileNet P8 Non-English Support Guide*.
1. "Upgrading Content Engine data using Content Engine Web Upgrade Tool" on page 77

Use Content Engine Web Upgrade Tool to validate and upgrade your version 3.5.2 FileNet P8 domain and object stores.

2. "Upgrading Content Engine data using Content Engine Upgrader tool" on page 87

Use Content Engine Upgrader tool to upgrade your version 3.5.2 FileNet P8 domain and object stores if you cannot use Content Engine Web Upgrade Tool.

Related information

➡️ Product documentation for IBM FileNet P8 Platform
Download the IBM FIleNet P8 Non-English Support Guide, as well as PDF's of the rest of the IBM FileNet P8 Platform documentation.

Upgrading Content Engine data using Content Engine Web Upgrade Tool

Use Content Engine Web Upgrade Tool to validate and upgrade your version 3.5.2 FileNet P8 domain and object stores.

Content Engine Web Upgrade Tool runs on the machine where Content Engine is deployed. You run the tool in two phases: validation and then upgrade.

The validation phase determines whether your object stores and related data are ready for upgrade. You can select all the object stores for validation in one step, or you can select subsets of them for validation in several steps. During the first validation step, even if you have not selected any object stores, Content Engine Web Upgrade Tool validates (by default) the GCD, domain addons, and fixed content devices. During a validation step Content Engine Web Upgrade Tool does not stop at errors; it proceeds until it has validated all the selected object stores. If an object store fails validation, you must edit its associated data and then run the validation again until it succeeds.

The upgrade phase upgrades the successfully validated object stores. You can select all the object stores for upgrading in one step, or you can select subsets of them for upgrading in several steps. The first time you run the upgrade step, even if you have not selected any object stores, Content Engine Web Upgrade Tool upgrades (by default) the GCD, domain addons, and fixed content devices. During an upgrade step Content Engine Web Upgrade Tool stops at the first error. If an object store upgrade fails, you must edit its associated data and then upgrade it again until it succeeds.

"Deploying Content Engine Web Upgrade Tool" on page 78
You need to deploy Content Engine Web Upgrade Tool before you can upgrade your Content Engine data.

"Backing up the version 4.5.1 GCD database" on page 79
You need to back up your version 4.5.1 GCD database before you start upgrading Content Engine data so that in case the upgrade fails, you can recover your original system environment.

"Starting Content Engine Web Upgrade Tool" on page 80
To use Content Engine Web Upgrade Tool, you browse to the machine where the tool is running.

"Upgrading a version 3.5.2 FileNet P8 domain into a new version 4.5.1 FileNet P8 domain" on page 81
You have a new version 4.5.1 FileNet P8 domain and are upgrading version 3.5.2 object stores into it.

"Upgrading a version 3.5.2 FileNet P8 domain into an existing version 4.5.1 FileNet P8 domain" on page 84

You have an existing version 4.5.1 FileNet P8 domain (which already contains object stores) and are upgrading version 3.5.2 object stores into it.

Deploying Content Engine Web Upgrade Tool:

You need to deploy Content Engine Web Upgrade Tool before you can upgrade your Content Engine data.

You must deploy Content Engine Web Upgrade Tool on the application server where you installed version 4.5.1 of the Content Engine software. Content Engine Web Upgrade Tool is available as of version 4.5.1 FP2 of Content Engine.

"Deploying Content Engine Web Upgrade Tool (WebSphere Application Server)"
You need to deploy IBM FileNet Content Engine Web Upgrade Tool on the application server where version 4.5.1 of Content Engine is deployed.

Deploying Content Engine Web Upgrade Tool (WebSphere Application Server):

You need to deploy IBM FileNet Content Engine Web Upgrade Tool on the application server where version 4.5.1 of Content Engine is deployed.

If Content Engine is deployed on a managed node, you must deploy Content Engine Web Upgrade Tool on the that node. If Content Engine is deployed in a cluster, you must deploy Content Engine Web Upgrade Tool in that cluster.

To deploy IBM FileNet Content Engine Web Upgrade Tool on WebSphere Application Server:

1. Log in as *ce_install_user* to the application server machine where you installed version 4.5.1 FP2 or later of the Content Engine software.
2. If you are going to upgrade a large object store or a large number of object stores, it is a best practice to deploy Content Engine Web Upgrade Tool to a Java Virtual Machine that has a heap size of at least 512 MB. If the heap size is less than 512 MB, increase the size as needed.
3. Navigate to the directory that contains the `upgrader.war` file, depending on the operating system where Content Engine is installed:

Option	Description
UNIX	`/opt/IBM/FileNet/ContentEngine/lib`

4. In the WebSphere Application Server administrative console, navigate to **Applications** → **Enterprise Applications** and click **Install**.
5. (WebSphere 6.1) Specify the path to the `upgrader.war` file, set the context root to `/upgrader`, and complete the following substeps:
   a. In the step to select installation options, set the application name to `upgrader`.
   b. If your deployment environment is WebSphere Application Server Network Deployment, select the node or cluster where Content Engine is deployed; otherwise, skip past the step to map modules to servers and the step to map virtual hosts to Web modules.
   c. Skip to the summary. Click **Finish**, and then click **Save**.
6. (WebSphere 7.0) In the Install New Application page, specify the path to the `upgrader.war` file, click **Next**, and then click **OK** to confirm the path. Then complete the following substeps:

a. In the step to select installation options, set the application name to `upgrader`.

b. If your deployment environment is WebSphere Application Server Network Deployment, select the node or cluster where Content Engine is deployed; otherwise, skip past the step to map modules to servers and the step to map virtual hosts to Web modules.

c. In the step to map context roots for Web modules, set the context root to `/upgrader`.

d. Skip to the summary. Click **Finish**, and then click **Save**.

7. Ensure that the class loader policy is set to Multiple by completing the following substeps on each server where you are deploying the `upgrader.war` file:

a. In the WebSphere Application Server administrative console, navigate to **Servers** → **WebSphere application servers** and click the name of the server where you just deployed the `upgrader.war` file.

b. Click the **Configuration** tab and under Server-specific Application Settings set the Classloader policy to Multiple.

8. Synchronize nodes if needed.

9. If your deployment environment is WebSphere Application Server Network Deployment, complete the following substeps:

a. Open for editing the deployed version of the `config.properties` file. You can this file at one of the following locations.

Operating system	Path to config.properties file
UNIX	`/opt/IBM/WebSphere/AppServer/profiles/` *`profile_name`*`/installedApps/`*`cell_name`*`/` `upgrader_war.ear/upgrader.war/WEB-INF/` `classes`
Windows	`C:\Program Files\IBM\WebSphere\` `AppServer\profiles\`*`profile_name`*`\` `installedApps\`*`cell_name`*`\` `upgrader_war.ear\upgrader.war\WEB-INF\` `classes`

b. Find the line in the file that starts with `SOAPConnectorAddress` and change the port (the value after the `=`) to the port number that the deployment manager uses.

10. Navigate back to **Applications** → **Enterprise Applications**, select upgrader, and then click **Start**.

Related concepts

⇨ Installation and Upgrade Worksheet
For more information about the Content Engine parameter values, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Backing up the version 4.5.1 GCD database:

You need to back up your version 4.5.1 GCD database before you start upgrading Content Engine data so that in case the upgrade fails, you can recover your original system environment.

To back up your version 4.5.1 GCD database:

User your database administration tools to back up the database that contains the GCD.

Starting Content Engine Web Upgrade Tool:

To use Content Engine Web Upgrade Tool, you browse to the machine where the tool is running.

To start Content Engine Web Upgrade Tool:
1. Log on to a machine where version 8.0 of Microsoft Internet Explorer is installed.
2. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ⟶ **Filter** ⟶ **AutoFilter** command is enabled. To view only Content Engine Web Upgrade Tool values, filter by **CE Web Upgrade Tool** in the **Installation or Configuration Program** column.
3. From the machine you logged on to, browse to the Web site `http://`*ce_server_name*`:`*port_number*`/upgrader`. *ce_server_name* is the host name of the machine where you run the tool, and *port_number* is the WSI port used by the application server on which Content Engine is deployed. (You cannot browse to the Web site using the HTTPS protocol.)

   Important: If Content Engine is in a cluster or farm environment, you can run the tool from any node in the environment. Browse directly to the node where the tool is running, rather than using a load balance URL.
4. At the Log In window, log in as *gcd_admin*, the GCD administrator user. This user must also be an object store administrator with full control permissions on each object store to be upgraded.

   Important: Do not run more than one instance of Content Engine Web Upgrade Tool at a time.

   If this is your first login for the FileNet P8 domain that you are upgrading, continue at step 5. Otherwise, Content Engine Web Upgrade Tool displays the Upgrade Information window, where you can resume validating and upgrading the object stores.
5. Enter the location of the GCD file `sysinit.dat` on the Windows machine where version 3.5.2 of Content Engine is installed. If you previously copied `sysinit.dat` to another location to ensure that it is accessible from a UNIX Content Engine server, you can specify that location instead. Do not rename the file: otherwise, Content Engine Web Upgrade Tool prompts you for a file named `sysinit.dat`.

   Important: Do not delete the GCD file on the Windows machine where version 3.5.2 of Content Engine is installed.
6. Verify that a file named `upgrade4512.xml` exists. To find the file, navigate to the directory shown in the working directory entry in the right panel of the Upgrade Information page. Then, navigate within that directory to the `FileNet/upgrader` directory. You might need to edit this file to fix errors that occur during the validation or upgrade phases.
7. Verify that a file named `upgrade_trace-`*timestamp*`.log` is in the `FileNet/upgrader/logs` directory. Content Engine Web Upgrade Tool creates this file anew each time you log in, and it moves the file from the previous login to the `FileNet/upgrader/logs/backup` directory. This file records all the validation and upgrade activity.

Upgrading a version 3.5.2 FileNet P8 domain into a new version 4.5.1 FileNet P8 domain:

You have a new version 4.5.1 FileNet P8 domain and are upgrading version 3.5.2 object stores into it.

"Validating object stores"
Validation determines whether your object stores and other Content Engine data are ready to be upgraded. You cannot upgrade your object stores until they are successfully validated.

"Upgrading object stores" on page 82
You need to upgrade your object stores before they can be made available for use in version 4.5.1 FP2 (or later) of Content Engine.

Validating object stores:

Validation determines whether your object stores and other Content Engine data are ready to be upgraded. You cannot upgrade your object stores until they are successfully validated.

To validate object stores:
1. For each fixed content device in the left panel of the Upgrade Information window that needs info (as indicated by its icon in the left panel), select the device and specify the required parameter values in the right panel.

   Important: The SnapLockPassword parameter value is not stored. If you log off Content Engine Web Upgrade Tool and log back on, you must reenter the value.
2. If Hitachi [SnapLock] or IBM [Tivoli] are listed among the fixed content devices in the left panel, complete the following substeps. Otherwise, Content Engine Web Upgrade Tool does not process them:
   a. Choose **Tools** → **Edit XML**.
   b. In the `upgrade4512.xml` file, replace all occurrences of `Unsupported` with `UpgradeFinished` to make Content Engine Web Upgrade Tool ignore these unsupported fixed content devices and fixed file stores (if any) in the object stores you upgrade.
   c. Click **Save**.
3. Select the object stores to be validated. Initially, all the object stores are selected for validation. You can validate object stores in any order.

   If you want to validate object stores individually, select them from the left panel of Content Engine Web Upgrade Tool. If you want to select all the object stores at once, choose **Tools** → **Select All Object Stores**.

   You can clear your selections by choosing **Tools** → **Clear All Object Stores.** (The **Clear All Object Stores** menu choice is unavailable until you select at least one object store.)
4. For each object store that you selected to be validated, complete the following substeps:
   a. In the right panel, specify the required values from your worksheet.
   b. If the object store uses fixed file stores, set the values for the root directories to their current locations. (The initial values that appear in these fields are from `sysinit.dat` and are not correct for version 4.5.1 of Content Engine.)

c. When you click **Save**, a prompt reminds you to be sure that you chose the correct data sources for the object store. It also reminds to you to verify that you can read from and write to the file stores of the object stores from the Content Engine server.

5. Click **Validate** to validate the object stores that you selected.

   Each time you click **Validate**, Content Engine Web Upgrade Tool creates the file `validation.log` in the directory `FileNet/upgrader/logs`, and moves the (time-stamped) log file from the previous validation (if any) into the directory `FileNet/upgrader/logs/backup`. Both directories are relative to the location specified in the working directory on the Upgrade Information page.

   Validation errors, such as missing or incorrect parameters, open in the right panel and are also written to `validation.log`.

6. Select **View** › **Validation Log** to check the `validation.log` file for errors. The file opens in a scrollable window. Errors are also displayed in the right panel.

   You can correct some validation errors within the Content Engine Web Upgrade Tool interface; other errors can be corrected only manually or by using another interface, such as IBM FileNet Enterprise Manager.

7. To correct a validation error within Content Engine Web Upgrade Tool, complete the following substeps:

   a. If the error indicates unspecified or invalid JDBC or JDBC XA data sources, choose the correct data source or fix the data source at the application server.

   b. If the error is in the file store path, specify the correct path.

   c. If the error is in the file store path, ensure that the file store exists, and that Content Engine can read and write to the file store. (Content Engine Web Upgrade Tool checks that the destination directory associated with the file store exists and that the inbound directory is empty.)

   d. Click **Save**.

   e. Click **Validate** and return to step 6.

8. If the validation log indicates an invalid object store state because of pre-upgrade errors (for example, conflicting classes or properties or pending queue items or publishing requests), log on to version 3.5.2 of IBM FileNet Enterprise Manager and fix the error within the Content Engine 3.5.2 domain.

9. If you successfully validated an object store, you can upgrade it. To validate more object stores, return to step 3 on page 81.

Upgrading object stores:

You need to upgrade your object stores before they can be made available for use in version 4.5.1 FP2 (or later) of Content Engine.

To upgrade object stores:

1. Select the object stores to be upgrade.

   If you want to upgrade object stores individually, select them from the left panel of Content Engine Web Upgrade Tool. If you want to select all the object stores at once, choose **Tools** › **Select All Object Stores**. You can upgrade object stores in any order. In either case, the first set of Content Engine object store that you select for upgrade includes (by default) the GCD, domain addons, and fixed content devices.

   You can clear your selections by choosing **Tools** › **Clear All Object Stores.** (The **Clear All Object Stores** menu choice is unavailable until you select at least one object store.

Tip: The upgrade log file can get long. For ease in viewing the log, upgrade only one object store at a time.

2. Click **Upgrade**. The right panel displays a log showing the progress of the upgrade.

   Each time you click **Upgrade**, Content Engine Web Upgrade Tool creates a file named `upgrade-running.log` in the `FileNet/upgrader/logs` directory, and copies the previous `upgrade-running.log` file to the `upgrader_trace-timestamp.log` file in the `FileNet/upgrader/logs/backup` directory.

   After the upgrade starts, you can pause it by clicking **Pause**. The upgrade pauses as soon as it finishes upgrading whatever object store node it is currently processing. To resume the upgrade, click **Resume**.

   If you exit the browser in which Content Engine Web Upgrade Tool is running, an upgrade that is in progress continues in the background. If you log back in, you see the upgrade status at its current state in the right panel.

   Content Engine Web Upgrade Tool indicates which object store it is upgrading by changing the appearance of the corresponding icon in the left panel. The tool also changes the appearance of each icon to indicate that it has completed the upgrade of an object store.

   Content Engine Web Upgrade Tool also puts the most recent upgrade report file `upgrade.htm` into the directory `FileNet/upgrader/reports` and moves the (time-stamped) report file from the previous upgrade into the directory `FileNet/upgrader/reports/backup`. Both directories are relative to the location specified in the working directory on the Upgrade Information page.

3. If Content Engine Web Upgrade Tool detects an error in upgrading an object store, it halts. To correct the error, complete the following substeps:

   a. Choose **View → Upgrade Log** and search the upgrade log to determine where the error occurred.

   b. Close the error log in the right panel and click the node in the left panel where the error occurred.

   c. If the error can be fixed without requiring you to restore the version 4.5.1 GCD database (in case of a domain upgrade failure) or the object store database (in case of an object store upgrade failure), fix the condition that caused the error and then click **Reenable Upgrade** next to the failed node to reset the node to the ready state.

   d. If the error indicates that you need to restore the version 4.5.1 GCD database to its state before the upgrade, click **Reenable Upgrade** at the top of the right panel for each domain node to reset it to the ready state. Use your database tools to restore the database to its state before the upgrade of Content Engine from version 3.5.2.

   e. If the error indicates that you need to restore the object store database to its state before the upgrade, click **Reenable Upgrade** at the top of the right panel to reset all nodes of the object store to ready state. Use your database tools to restore the database to its state before the upgrade of Content Engine from version 3.5.2. If the error occurred during the Upgrade Addons phase because of classes or properties that were removed from the object store and imported by a System Addon, reimport the Addon XML file to replace the missing classes or properties. Then, log in to version 4.5.1 FP2 of IBM FileNet Enterprise Manager and refresh the object stores node. If the object store that failed is visible, delete it.

   f. Click **Bring Online** for each object store selected for upgrade that Content Engine Upgrade Tool did not attempt to upgrade (because of the error that stopped processing) if you need the version 3.5.2 object stores available while you are resolving the object store or domain error.

g. Return to step 2 on page 83.

4. If the upgrade succeeds, and you have additional object stores to upgrade, return to step 1 on page 82.

   Important: When you have successfully upgraded an object store, Content Engine Web Upgrade Tool prevents you from trying to upgrade the object store again.

5. When you have upgraded all the object stores, choose **View** → **Upgrade Report** to see a summary of the results of the upgrades.

Upgrading a version 3.5.2 FileNet P8 domain into an existing version 4.5.1 FileNet P8 domain:

You have an existing version 4.5.1 FileNet P8 domain (which already contains object stores) and are upgrading version 3.5.2 object stores into it.

To upgrade a version 3.5.2 domain into an existing version 4.5.1 FileNet P8 domain:

1. Resolve any conflicts between the GCD objects (such as marking sets and fixed content devices) and object store names in the version 4.5.1 domain with the corresponding GCD objects and object store names in the version 3.5.2 domain.

2. Choose **File** → **Start New Upgrade** to start an upgrade of version 3.5.2.

3. Enter the location of the GCD file sysinit.dat on the Windows machine where version 3.5.2 of Content Engine is installed. If you previously copied sysinit.dat to another location to ensure that it is accessible from a UNIX Content Engine server, you can specify that location instead. Do not rename the file: otherwise, Content Engine Web Upgrade Tool prompts you for a file named sysinit.dat.

   Important: Do not delete the GCD file on the Windows machine where version 3.5.2 of Content Engine is installed.

4. Verify that a file named upgrade4512.xml exists. To find the file, navigate to the directory shown in the working directory entry in the right panel of the Upgrade Information page. Then, navigate within that directory to the FileNet/upgrader directory. You might need to edit this file to fix errors that occur during the validation or upgrade phases.

5. Verify that a file named upgrade_trace-*timestamp*.log is in the FileNet/upgrader/logs directory. Content Engine Web Upgrade Tool creates this file anew each time you log in, and it moves the file from the previous login to the FileNet/upgrader/logs/backup directory. This file records all the validation and upgrade activity.

   "Upgrading object stores"
   You need to upgrade your object stores before they can be made available for use in version 4.5.1 FP2 (or later) of Content Engine.

   "Validating object stores" on page 86
   Validation determines whether your object stores and other Content Engine data are ready to be upgraded. You cannot upgrade your object stores until they are successfully validated.

Upgrading object stores:

You need to upgrade your object stores before they can be made available for use in version 4.5.1 FP2 (or later) of Content Engine.

To upgrade object stores:

1. Select the object stores to be upgrade.

   If you want to upgrade object stores individually, select them from the left panel of Content Engine Web Upgrade Tool. If you want to select all the object stores at once, choose **Tools → Select All Object Stores**. You can upgrade object stores in any order. In either case, the first set of Content Engine object store that you select for upgrade includes (by default) the GCD, domain addons, and fixed content devices.

   You can clear your selections by choosing **Tools → Clear All Object Stores.** (The **Clear All Object Stores** menu choice is unavailable until you select at least one object store.

   Tip: The upgrade log file can get long. For ease in viewing the log, upgrade only one object store at a time.

2. Click **Upgrade**. The right panel displays a log showing the progress of the upgrade.

   Each time you click **Upgrade**, Content Engine Web Upgrade Tool creates a file named `upgrade-running.log` in the `FileNet/upgrader/logs` directory, and copies the previous `upgrade-running.log` file to the `upgrader_trace-`*timestamp*`.log` file in the `FileNet/upgrader/logs/backup` directory.

   After the upgrade starts, you can pause it by clicking **Pause**. The upgrade pauses as soon as it finishes upgrading whatever object store node it is currently processing. To resume the upgrade, click **Resume**.

   If you exit the browser in which Content Engine Web Upgrade Tool is running, an upgrade that is in progress continues in the background. If you log back in, you see the upgrade status at its current state in the right panel.

   Content Engine Web Upgrade Tool indicates which object store it is upgrading by changing the appearance of the corresponding icon in the left panel. The tool also changes the appearance of each icon to indicate that it has completed the upgrade of an object store.

   Content Engine Web Upgrade Tool also puts the most recent upgrade report file `upgrade.htm` into the directory `FileNet/upgrader/reports` and moves the (time-stamped) report file from the previous upgrade into the directory `FileNet/upgrader/reports/backup`. Both directories are relative to the location specified in the working directory on the Upgrade Information page.

3. If Content Engine Web Upgrade Tool detects an error in upgrading an object store, it halts. To correct the error, complete the following substeps:

   a. Choose **View → Upgrade Log** and search the upgrade log to determine where the error occurred.

   b. Close the error log in the right panel and click the node in the left panel where the error occurred.

   c. If the error can be fixed without requiring you to restore the version 4.5.1 GCD database (in case of a domain upgrade failure) or the object store database (in case of an object store upgrade failure), fix the condition that caused the error and then click **Reenable Upgrade** next to the failed node to reset the node to the ready state.

   d. If the error indicates that you need to restore the version 4.5.1 GCD database to its state before the upgrade, click **Reenable Upgrade** at the top of the right panel for each domain node to reset it to the ready state. Use your database tools to restore the database to its state before the upgrade of Content Engine from version 3.5.2.

e. If the error indicates that you need to restore the object store database to its state before the upgrade, click **Reenable Upgrade** at the top of the right panel to reset all nodes of the object store to ready state. Use your database tools to restore the database to its state before the upgrade of Content Engine from version 3.5.2. If the error occurred during the Upgrade Addons phase because of classes or properties that were removed from the object store and imported by a System Addon, reimport the Addon XML file to replace the missing classes or properties. Then, log in to version 4.5.1 FP2 of IBM FileNet Enterprise Manager and refresh the object stores node. If the object store that failed is visible, delete it.

f. Click **Bring Online** for each object store selected for upgrade that Content Engine Upgrade Tool did not attempt to upgrade (because of the error that stopped processing) if you need the version 3.5.2 object stores available while you are resolving the object store or domain error.

g. Return to step 2 on page 85.

4. If the upgrade succeeds, and you have additional object stores to upgrade, return to step 1 on page 85.

   Important: When you have successfully upgraded an object store, Content Engine Web Upgrade Tool prevents you from trying to upgrade the object store again.

5. When you have upgraded all the object stores, choose **View** → **Upgrade Report** to see a summary of the results of the upgrades.

Validating object stores:

Validation determines whether your object stores and other Content Engine data are ready to be upgraded. You cannot upgrade your object stores until they are successfully validated.

To validate object stores:

1. For each fixed content device in the left panel of the Upgrade Information window that needs info (as indicated by its icon in the left panel), select the device and specify the required parameter values in the right panel.

   Important: The SnapLockPassword parameter value is not stored. If you log off Content Engine Web Upgrade Tool and log back on, you must reenter the value.

2. If Hitachi [SnapLock] or IBM [Tivoli] are listed among the fixed content devices in the left panel, complete the following substeps. Otherwise, Content Engine Web Upgrade Tool does not process them:

   a. Choose **Tools** → **Edit XML**.

   b. In the `upgrade4512.xml` file, replace all occurrences of `Unsupported` with `UpgradeFinished` to make Content Engine Web Upgrade Tool ignore these unsupported fixed content devices and fixed file stores (if any) in the object stores you upgrade.

   c. Click **Save**.

3. Select the object stores to be validated. Initially, all the object stores are selected for validation. You can validate object stores in any order.

   If you want to validate object stores individually, select them from the left panel of Content Engine Web Upgrade Tool. If you want to select all the object stores at once, choose **Tools** → **Select All Object Stores**.

You can clear your selections by choosing **Tools** → **Clear All Object Stores.** (The **Clear All Object Stores** menu choice is unavailable until you select at least one object store.)

4. For each object store that you selected to be validated, complete the following substeps:

   a. In the right panel, specify the required values from your worksheet.

   b. If the object store uses fixed file stores, set the values for the root directories to their current locations. (The initial values that appear in these fields are from `sysinit.dat` and are not correct for version 4.5.1 of Content Engine.)

   c. When you click **Save**, a prompt reminds you to be sure that you chose the correct data sources for the object store. It also reminds to you to verify that you can read from and write to the file stores of the object stores from the Content Engine server.

5. Click **Validate** to validate the object stores that you selected.

   Each time you click **Validate**, Content Engine Web Upgrade Tool creates the file `validation.log` in the directory `FileNet/upgrader/logs`, and moves the (time-stamped) log file from the previous validation (if any) into the directory `FileNet/upgrader/logs/backup`. Both directories are relative to the location specified in the working directory on the Upgrade Information page.

   Validation errors, such as missing or incorrect parameters, open in the right panel and are also written to `validation.log`.

6. Select **View** → **Validation Log** to check the `validation.log` file for errors. The file opens in a scrollable window. Errors are also displayed in the right panel.

   You can correct some validation errors within the Content Engine Web Upgrade Tool interface; other errors can be corrected only manually or by using another interface, such as IBM FileNet Enterprise Manager.

7. To correct a validation error within Content Engine Web Upgrade Tool, complete the following substeps:

   a. If the error indicates unspecified or invalid JDBC or JDBC XA data sources, choose the correct data source or fix the data source at the application server.

   b. If the error is in the file store path, specify the correct path.

   c. If the error is in the file store path, ensure that the file store exists, and that Content Engine can read and write to the file store. (Content Engine Web Upgrade Tool checks that the destination directory associated with the file store exists and that the inbound directory is empty.)

   d. Click **Save**.

   e. Click **Validate** and return to step 6.

8. If the validation log indicates an invalid object store state because of pre-upgrade errors (for example, conflicting classes or properties or pending queue items or publishing requests), log on to version 3.5.2 of IBM FileNet Enterprise Manager and fix the error within the Content Engine 3.5.2 domain.

9. If you successfully validated an object store, you can upgrade it. To validate more object stores, return to step 3 on page 86.

Upgrading Content Engine data using Content Engine Upgrader tool

Use Content Engine Upgrader tool to upgrade your version 3.5.2 FileNet P8 domain and object stores if you cannot use Content Engine Web Upgrade Tool.

"Preparing the Content Engine Upgrader tool" on page 88
Before using the Content Engine Upgrader tool to update your Content Engine data, you must install the JDBC driver and edit the upgrader batch file.

"Backing up the version 4.5.1 GCD database"
You need to back up your version 4.5.1 GCD database before you start
upgrading Content Engine data so that in case the upgrade fails, you can
recover your original system environment.

Preparing the Content Engine Upgrader tool:

Before using the Content Engine Upgrader tool to update your Content Engine
data, you must install the JDBC driver and edit the upgrader batch file.

Backing up the version 4.5.1 GCD database:

You need to back up your version 4.5.1 GCD database before you start upgrading
Content Engine data so that in case the upgrade fails, you can recover your
original system environment.

To back up your version 4.5.1 GCD database:

User your database administration tools to back up the database that contains the
GCD.

Updating the database storage area parameters

After you have upgraded your object stores, you need to initialize the parameter
values for the associated database storage areas.

Having upgraded the object stores, you now need to set the file count and the file
bytes parameter values for the associated database storage areas. These values
need to match the values for the database storage areas before the object stores
were upgraded.

To update the database storage area parameters:

Run the Content Consistency Checker tool to update the file count and the file
bytes parameter values for the database storage areas. You update the values by
running the Update Storage Area Statistics function in the tool. For details, see
System Administration → **Content Engine Administration** → **Content storage** →
Storage area consistency → **How to...** → **Update storage area statistics**.

Completing Content Search Engine upgrade from 3.5.2

After you upgrade Content Engine data, you must create new content-search index
areas in Enterprise Manager. Then you can enable indexing for designated classes
and properties and reindex your stored data. The reindexing process creates new
Autonomy K2 collections.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

This procedure covers the minimum setup and configuration steps to get
content-based retrieval (CBR) configured and running with Autonomy K2 for an
upgrade from 3.5.2. For more detail on content-based retrieval and Content Engine,
see the IBM FileNet P8 help topic **System Administration** → **Content Engine
Administration** → **Content-based retrieval** → **How to...** → **Configure CBR**.

Important: Where machine name variables are required, IP addresses will not
validate. In these cases, you must enter the host name for the machine.

To complete the Content Search Engine upgrade from version 3.5.2:

1. Start Enterprise Manager and log on as the Global Configuration Data (GCD) Administrator.
2. Create a Verity Domain Configuration (VDC).
   a. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Enterprise Manager values for this task, filter by **EM: Verity Domain Configuration Tab** in the **Installation or Configuration Program** column.
   b. Right-click **Enterprise Manager [domain]** in the Enterprise Manager tree and select **Properties**.
   c. Click the **Verity Domain Configuration** tab.
   d. 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. The default port number is 9950.
      - 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.
   e. Click **Create Configuration** to 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**.

   Important: 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. A given Content Engine server does not call multiple Broker Servers or merge associated search results. See the appropriate steps of the Configuring Content Search Engine task for instructions on how to create Broker Servers and assign Search Servers.
4. To access the Verity Locale Configuration Guide, complete the following substeps.

   Important: For information about how the K2 product uses locales and about performance and indexing options, review the Verity Locale Configuration Guide in the Autonomy-supplied documentation set.
   a. Open the Autonomy K2 Dashboard and click **Help**.
   b. Click **Library** (top right of the screen) to access the Autonomy K2 product guides. Each guide is available in HTML or PDF format.
5. To assign a CBR locale to the object store for which you want to enable indexing, complete the following substeps:

a. Right-click the object store in the left pane of Enterprise Manager and choose **Properties** from the context menu.

b. In the Properties dialog box, click the **CBR** tab.

c. Specify one of the following locales (which are installed by default with Autonomy K2 Content Search Engine software) in the CBR Locale box:

- uni (multi-language and slower indexing)
- englishv (enhanced word-stemming and faster indexing)
- english (limited capabilities and fast indexing)

If you click **Set to Default CBR Locale**, the CBR locale is set to *uni*.

d. Click **Apply** or **OK**.

6. Create an index area for version 4.5.1 Content Search Engine collections:

a. Open your completed Installation and Upgrade Worksheet file.

Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Enterprise Manager values for this task, filter by **EM: Create an Index Area** in the **Installation or Configuration Program** column.

b. In the Enterprise Manager tree view, expand the **Object Stores** container.

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

d. Complete the Create an Index Area wizard by using the values from your worksheet.

Important: In the wizard, make sure you set the template type to FileNet_FileSystem_PushAPI.

For more information about creating indexes, see the IBM FileNet P8 help topic **System Administration → Content Engine Administration → Content-based retrieval → How to... → Create Verity index area**.

7. Re-enable content-based retrieval for class definitions by activating the CBR Enabled flag for each class:

a. Right-click *class_name* for the class that you want to configure in the Enterprise Manager tree and click **Properties**.

b. Select the **CBR Enabled** check box, and click **OK**.

c. Click **Yes** to confirm.

8. For each class property for which you want to enable content-based retrieval, complete the following substeps. The properties you enable must include those that you indexed for content-based retrieval in version 3.5.2 of Content Engine. Refer to the list of properties that you generated by SQL query command on each object store during the pre-upgrade phase of upgrading Content Engine.

a. Right-click the base class or subclass that contains the property or properties that you want to reindex and click **Properties** from the context menu.

b. Click the **Property Definitions** tab.

c. Click the string property you want to enable for indexing and then click **Edit**.

d. Select the CBR Enabled check box, and then click **OK**.

9. Reindex each class on which you enabled indexing in the preceding two steps by right-clicking the class and choosing one of the following:

- **All Tasks → Index Class for Content Search**.

- **All Tasks → Index Class for Content Search (include subclasses)**.

  For details, see the IBM FileNet P8 help topic **System Administration → Content Engine Administration → Content-based retrieval → How to... → Reindex**.

10. To monitor the reindexing progress, complete the following substeps:
    a. Select and right-click the object store node that contains the index areas that are being reindexed.
    b. Choose **All Tasks → Index Jobs Management**.
    c. Click **Refresh** to monitor the progress. Click **Close** when reindexing is complete, as indicated by the job status <TERMINATED NORMALLY>.

11. Because the version 3.5.2 Content Search Engine collections (indexes) are no longer used, you can save hard drive space by deleting the index directories and all subdirectories. For example, `D:\CE_storage\FS_object_store_name\ index`.

Related tasks

"Determining the CBR-enabled properties in an object store" on page 14
For each object store to be upgraded, you must determine the associated classes and properties that are currently enabled for content-based retrieval (CBR). You enable these properties after the upgrade is complete.

Upgrading Enterprise Manager

If you did not upgrade Enterprise Manager as part of a Content Engine upgrade, you can upgrade Enterprise Manager interactively or silently.

Be sure that you have available the Installation and Upgrade Worksheet that was filled out during your planning activities.

Important: You must keep one running version of Enterprise Manager version 3.5.2 until after you upgrade the Content Engine data. You must use Enterprise Manager version 3.5.2 to prepare your version 3.5.2 object stores for the Content Engine Upgrader tool.

To upgrade Enterprise Manager:
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CE Installer** in the **Installation or Configuration Program** column.
2. On the machine where you are upgrading Enterprise Manager, log on as a member of the Local Administrators group or the Power Users group.
3. Access the Content Engine software package.
4. Start the Enterprise Manager installation.

Table 14. Installation methods

Installation method	Description
To install interactively	1. Run the following command in the software package: `P8CE-4.5.1-Win.exe` 2. Complete the installation program by using the values in your installation worksheet.

Table 14. Installation methods (continued)

Installation method	Description
To install silently	1. Open the `CE_silent_install.txt` file in the software package for editing.
	2. 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`
	3. Save your edits.

	4. Run the following command in the software package:
	
	`P8CE-4.5.1-WIN.EXE -f`
	`CE_silent_install.txt -i silent`

Related concepts

↪ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading from version 4.0.1 or 4.5.0

Upgrading Content Engine and Content Search Engine from version 4.0.1 or 4.5.0 is a multi-step process.

Important: Before you start to upgrade the Content Engine software, ensure that you are at the correct version and fix pack level, as indicated in *Plan and Prepare Your Environment for IBM FileNet P8*.

1. "Completing pre-upgrade Content Engine configuration" on page 94
   You must complete several pre-upgrade procedures to prepare the Content Engine for an upgrade from version 4.0.1 or 4.5.0.

2. "Uninstalling Content Engine and Content Search Engine Client Updater (Version 4.0.1 and AIX)" on page 96
   If you are upgrading from Content Engine version 4.0.1 on AIX, where the root user installed Content Engine version 4.0.1 and the user who is performing the upgrade is non-root, then you must uninstall the existing version of the Content Engine and Content Search Engine Client Updater software.

3. "Upgrading Content Engine software from version 4.0.1 or 4.5.0" on page 98
   On UNIX application servers, you will upgrade Content Engine Server and install Configuration Manager. On Windows application servers, you will upgrade or install these components as needed: Content Engine Server, Configuration Manager, .NET Clients (including Enterprise Manager), Content Engine Upgrader, or FileNet Deployment Manager. You can run the upgrade interactively or silently.

4. "Upgrading Enterprise Manager" on page 100
   If you did not upgrade Enterprise Manager as part of a Content Engine upgrade, you can upgrade Enterprise Manager interactively or silently.

5. "Installing Content Engine software updates" on page 101
   You must apply any updates, fix packs, or interim fixes for the Content Engine software that have been released.

6. "Installing the latest Content Search Engine Client files on Content Engine servers" on page 101
   If your environment includes Content Search Engine, you must install the files that Content Engine uses to communicate with Content Search Engine.

7. "Granting directory permissions to the Configuration Manager user" on page 105
   You must grant the file and directory permissions required by *config_mgr_user*, which is the user who runs Configuration Manager.

8. "Configuring Content Engine instances upgraded from version 4.0.1 or 4.5.0" on page 105
   You can configure your upgraded Content Engine instances with Configuration Manager. You must upgrade the Content Engine bootstrap properties for the upgraded Content Engine instance. A single Content Engine application instance equates to one deployed application on your application server.

9. "Installing the latest Process Engine client files on Content Engine servers" on page 111
   You must upgrade the Process Engine Client files on all application server machines where Content Engine Server is deployed.

10. "Deploying upgraded Content Engine instances" on page 113
    You must deploy the upgradedContent Engine on the web application server. Deploying makes the Content Engine application available for use.

11. "Installing storage device source files" on page 116
    If your upgraded IBM FileNet P8 system includes Tivoli Storage Manager or EMC Centera devices, you must install files on the Content Engine server.

12. "Completing Content Engine post-deployment steps" on page 119
    You must complete several Content Engine post-deployment steps before you can put an IBM FileNet P8 system into production.

13. "Verifying the Content Engine deployment" on page 120
    You can verify that the Content Engine deployment was successful by accessing the IBM FileNet P8 System Health page.

14. "Checking the automatic upgrade status for Content Engine data" on page 122
    When you deploy the upgraded Content Engine application on the web application server and start the application, the existing Content Engine data upgrade begins. You can check the progress on the upgrade to confirm that all object stores have been upgraded successfully

15. "Installing the latest Content Engine client files on other IBM FileNet P8 servers (for staged upgrades)" on page 123
    When you upgrade Content Engine 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.

16. "Upgrading Content Search Engine software from version 4.0.1 or 4.5.0" on page 123
    If you used the full-text search feature (CBR) in version 4.0.1 or 4.5.0, to upgrade you must export your current K2 server configurations on each machine, install the new version of Content-Search Engine (Autonomy K2) software, and then import your saved configuration.

Related information

➥ General requirements for all FileNet P8 Platform systems
Check to make sure your software is at the correct fix pack level before you upgrade.

Completing pre-upgrade Content Engine configuration

You must complete several pre-upgrade procedures to prepare the Content Engine for an upgrade from version 4.0.1 or 4.5.0.

1. "Verifying that publishing requests have been completed"
   You must verify that all publishing requests have completed for each object store before you begin the upgrade from version 4.0.1 or 4.5.0.

2. "Backing up the database" on page 95
   You must back up the database before you begin the upgrade.

3. "Shutting down CFS Exporter" on page 95
   If your system includes version 4.0.0 or 4.5 of Content Federation Services (CFS), you must shut down the CFS Exporter before upgrading Content Engine.

4. "Preventing name conflicts with reserved words" on page 95
   You must verify that the object store names, custom classes, or custom properties in your current version do not conflict with reserved words before you begin the upgrade from version 4.0.1 or 4.5.0.

5. "Removing the deployed FileNetEngine application" on page 96
   You must remove the existing FileNetEngine application from the web application server to undeploy or uninstall the application EAR file. Removing the existing version ensures a clean deployment of the upgraded version 4.5.1 Content Engine EAR file.

6. "Clearing the web application server cache" on page 96
   You must clear the cache for the web application server to remove data for the Content Engine version 3.5.2 application to prevent conflicts with the upgraded version.

Verifying that publishing requests have been completed

You must verify that all publishing requests have completed for each object store before you begin the upgrade from version 4.0.1 or 4.5.0.

To verify that publishing requests complete:

1. Start Enterprise Manager if it is not already running.

2. Search for publishing requests in the queue by expanding the **Publishing** folder and clicking **Queue**. Then, click **View → All requests**.

3. Verify that the queue (right pane) of publishing requests is empty. If the queue is not empty, wait until the publish requests are processed. 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).

4. In the list view, under the object store icon, right-click the **Search Results** folder, and choose **New Search**.

5. In the Content Engine Query Builder window, click **Select From Table → PublishRequest**.

6. Retain the default settings and click **OK**.

7. Click **Yes** at the prompt for a WHERE clause. Any remaining event items in the queue are displayed in the Query Status window.

8. Remove unwanted items from the queue.
   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 window, click **Select From Table** → **PublishRequest**.
   c. Select the **Delete objects** check box in the **Actions** tab of the Search window.
   d. Click **OK**. Click **OK** again to confirm the deletion.
9. Repeat this procedure for each object store in your environment.

Backing up the database

You must back up the database before you begin the upgrade.

1. Use your existing database backup solution to back up each of the Content Engine object store databases.
2. Backup the Global Configuration Data GCD database. The default GCD database name is FNGCDDB.
3. Backup the database for each object store in your environment.

Shutting down CFS Exporter

If your system includes version 4.0.0 or 4.5 of Content Federation Services (CFS), you must shut down the CFS Exporter before upgrading Content Engine.

Leave the CFS Exporter shut down until you upgrade CFS to version 4.5.1. See *IBM FileNet Content Federation Services Installation and Upgrade Guide* for instructions.

Preventing name conflicts with reserved words

You must verify that the object store names, custom classes, or custom properties in your current version do not conflict with reserved words before you begin the upgrade from version 4.0.1 or 4.5.0.

Verify that your custom classes and properties do not conflict with the reserved words listed in "New Content Engine classes and properties" in *Plan and Prepare Your Environment for IBM FileNet P8* that describes how you must change any custom class or custom property names in your existing object stores that match the reserved names.

Upgrading from version 4.0.1 to 4.5.1

   Upgrades from version 4.0.1 use the automatic upgrade facility, which automatically renames any conflicting class or property symbolic name that it detects during the upgrade. The upgrade facility writes warning messages for each renamed class or property symbolic name to the p8_server_error.log file on the application server where Content Engine is deployed. It also writes these messages to the automatic upgrade status web page. If you know of naming conflicts with version 4.5.0 classes or properties, it is a best practice to resolve these conflicts before you upgrade to version 4.5.1 of Content Engine.

Upgrading from version 4.5.0 to 4.5.1

   Upgrades from version 4.5.0 use the automatic upgrade facility, which will not encounter any conflicting class or property symbolic names.

Related concepts

➡ New Content Engine classes and properties
Learn about new reserved names and how to prevent naming conflicts during an upgrade.

Removing the deployed FileNetEngine application

You must remove the existing FileNetEngine application from the web application server to undeploy or uninstall the application EAR file. Removing the existing version ensures a clean deployment of the upgraded version 4.5.1 Content Engine EAR file.

To remove the deployed FileNetEngine application:

Undeploy or uninstall the Content Engine EAR file.

Option	Description
WebSphere Application Server	Use the administrative console to undeploy the FileNetEngine application, `Engine-ws.ear`.

Clearing the web application server cache

You must clear the cache for the web application server to remove data for the Content Engine version 3.5.2 application to prevent conflicts with the upgraded version.

To clear the cache:

Delete the cache files for you web application server.

Option	Description
WebSphere Application Server	`Program Files/IBM/WebSphere/AppServer/ profiles/profilename/temp/ profilenameserver1/FileNetEngine/*.*` Where: • *profilename* is profile for Content Engine. • *server_name* is the name of the Content Engine server instance.

Uninstalling Content Engine and Content Search Engine Client Updater (Version 4.0.1 and AIX)

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root user installed Content Engine version 4.0.1 and the user who is performing the upgrade is non-root, then you must uninstall the existing version of the Content Engine and Content Search Engine Client Updater software.

You can run the uninstall program interactively or silently.

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root user installed Content Engine version 4.0.1 and the user who is performing the upgrade is non-root, then you must uninstall the existing version of the Content Engine software. You can uninstall the software interactively.

"Uninstalling Content Engine silently (AIX)"
If you are upgrading from Content Engine version 4.0.1 on AIX, where the root
user installed Content Engine version 4.0.1 and the user who is performing the
upgrade is non-root, then you must uninstall the existing version of the Content
Engine software. You can uninstall the software silently.

"Uninstalling Content Search Engine Client Updater interactively (version 4.0.1
and AIX)" on page 98
If you are upgrading from Content Engine version 4.0.1 on AIX, where the root
user installed Content Engine version 4.0.1 and the user who is performing the
upgrade is non-root, then you must uninstall the existing version of the Content
Search Engine Client Updater software. You can uninstall the software
interactively.

"Uninstalling Content Search Engine Client Updater silently (version 4.0.1 and
AIX)" on page 98
If you are upgrading from Content Engine version 4.0.1 on AIX, where the root
user installed Content Engine version 4.0.1 and the user who is performing the
upgrade is non-root, then you must uninstall the existing version of the Content
Search Engine Client Updater software. You can uninstall the software
interactively.

Uninstalling Content Engine interactively (version 4.0.1 and AIX)

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root
user installed Content Engine version 4.0.1 and the user who is performing the
upgrade is non-root, then you must uninstall the existing version of the Content
Engine software. You can uninstall the software interactively.

To uninstall the software

1. Log on as *ce_install_user*, which is the user who installed the current version of
   Content Engine, to the server where Content Engine is installed.
2. Shut down the web application server.
3. Navigate to the _uninst2 directory in the current version of Content Engine.
4. Run the following command:

   `uninstaller.bin`

   The uninstall might take a while to complete.

Uninstalling Content Engine silently (AIX)

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root
user installed Content Engine version 4.0.1 and the user who is performing the
upgrade is non-root, then you must uninstall the existing version of the Content
Engine software. You can uninstall the software silently.

To uninstall the software

1. Log on as *ce_install_user*, which is the user who installed the current version of
   Content Engine, to the server where Content Engine is installed.
2. Shut down the web application server.
3. Navigate to the _uninst2 directory in the current version of Content Engine.
4. Edit the sample response file.
   a. Open the `CE_silent_uninstall.txt` sample response file in a text editor.
   b. Edit the file using the instructions in the file.
   c. Save your changes.
5. Run the following command:

```
./uninstaller.bin -options path_to_response_file -silent
```

where *path_to_response_file* is the path to the CE_silent_uninstall.txt that you edited. The uninstall might take a while to complete.

Uninstalling Content Search Engine Client Updater interactively (version 4.0.1 and AIX)

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root user installed Content Engine version 4.0.1 and the user who is performing the upgrade is non-root, then you must uninstall the existing version of the Content Search Engine Client Updater software. You can uninstall the software interactively.

To uninstall the software

1. Log on as *ce_install_user*, which is the user who installed the current version of Content Engine, to the server where Content Engine is installed.
2. Shut down the web application server.
3. Navigate to the _uninst2 directory in the current version of Content Engine.
4. Run the following command:
   ```
   uninstaller.bin
   ```

   The uninstall might take a while to complete.

Uninstalling Content Search Engine Client Updater silently (version 4.0.1 and AIX)

If you are upgrading from Content Engine version 4.0.1 on AIX, where the root user installed Content Engine version 4.0.1 and the user who is performing the upgrade is non-root, then you must uninstall the existing version of the Content Search Engine Client Updater software. You can uninstall the software interactively.

To uninstall the software

1. Log on as *ce_install_user*, which is the user who installed the current version of Content Engine, to the server where Content Engine is installed.
2. Shut down the web application server.
3. Navigate to the _uninst2 directory in the current version of Content Engine.
4. Edit the sample response file.
   a. Open the CE_silent_uninstall.txt sample response file in a text editor.
   b. Edit the file by using the instructions in the file.
   c. Save your changes.
5. Run the following command:
   ```
   ./uninstaller.bin -options path_to_response_file -silent
   ```

   where *path_to_response_file* is the path to the CE_silent_uninstall.txt that you edited. The uninstall might take a while to complete.

Upgrading Content Engine software from version 4.0.1 or 4.5.0

On UNIX application servers, you will upgrade Content Engine Server and install Configuration Manager. On Windows application servers, you will upgrade or install these components as needed: Content Engine Server, Configuration Manager, .NET Clients (including Enterprise Manager), Content Engine Upgrader, or FileNet Deployment Manager. You can run the upgrade interactively or silently.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

FileNet Deployment Manager is an optional tool that 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.

> "Upgrading Content Engine software interactively (UNIX)"
> You can upgrade the Content Engine software interactively. On UNIX application servers, you will upgrade Content Engine Server and Configuration Manager
>
> "Upgrading Content Engine software silently (UNIX)"
> You can upgrade the Content Engine software silently. On UNIX application servers, you will upgrade Content Engine Server and Configuration Manager.

Related concepts

↪ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Content Engine software interactively (UNIX)

You can upgrade the Content Engine software interactively. On UNIX application servers, you will upgrade Content Engine Server and Configuration Manager

To upgrade Content Engine software interactively:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Installer** in the **Installation or Configuration Program** column.

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

Platform	Command
AIX	P8CE-4.5.1-AIX.BIN

3. Complete the wizard screens by using the values in your worksheet.
4. When the installation completes, click **Done** and check for errors in the Content Engine *ce_install_path*/ce_install_log_4_5_1.txt error log file, where *ce_install_path* is the location where Content Engine is installed.

Related concepts

↪ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Content Engine software silently (UNIX)

You can upgrade the Content Engine software silently. On UNIX application servers, you will upgrade Content Engine Server and Configuration Manager.

To upgrade Content Engine software silently:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Installer** in the **Installation or Configuration Program** column.

2. Edit the `CE_silent_install.txt` file by providing the values from your worksheet.
3. Run one of the following commands in the software package, depending on your operating system:

Platform	Command
AIX	`P8CE-4.5.1-AIX.BIN -f CE_silent_install.txt -i silent`

4. When the installation completes, check for errors in the Content Engine error log file *ce_install_path*/`ce_install_log_4_5_1.txt`, where *ce_install_path* is the location where Content Engine is installed.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Enterprise Manager

If you did not upgrade Enterprise Manager as part of a Content Engine upgrade, you can upgrade Enterprise Manager interactively or silently.

Be sure that you have available the Installation and Upgrade Worksheet that was filled out during your planning activities.

To upgrade Enterprise Manager:
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ➔ **Filter** ➔ **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CE Installer** in the **Installation or Configuration Program** column.
2. On the machine where you are upgrading Enterprise Manager, log on as a member of the Local Administrators group or the Power Users group.
3. Access the Content Engine software package.
4. Start the Enterprise Manager installation.

Table 15. Installation methods

Installation method	Description
To install interactively	1. Run the following command in the software package: `P8CE-4.5.1-Win.exe` 2. Complete the installation program by using the values in your installation worksheet.

Table 15. Installation methods (continued)

Installation method	Description
To install silently	1. Open the `CE_silent_install.txt` file in the software package for editing.
	2. 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`
	3. Save your edits.
	4. Run the following command in the software package: `P8CE-4.5.1-WIN.EXE -f CE_silent_install.txt -i silent`

Related concepts

⯈ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing Content Engine software updates

You must apply any updates, fix packs, or interim fixes for the Content Engine software that have been released.

If no Content Engine software updates are available, skip this procedure.

To install the Content Engine software updates:

Open the readmes for the Content Engine software updates, and perform the installation procedures in the readmes on each Content Engine instance.

Installing the latest Content Search Engine Client files on Content Engine servers

If your environment includes Content Search Engine, you must install the files that Content Engine uses to communicate with Content Search Engine.

Important: With the initial release of Content Engine 4.5.1, the Content Search Engine 4.5.1 client files are included. Therefore, you are not required to manually install these client files on a new Content Engine server. However, if fix packs of these client files become available, install them manually before you use Configuration Manager to configure and deploy Content Engine. For instructions on how to obtain the latest Content Search Engine Client software updates, access the IBM FileNet P8 Platform support site.

You can manually install Content Search Engine Client files either interactively or silently:

1. "Removing the existing Content Search Engine Client 4.0.1 files from Content Engine Server machines" on page 102
   Remove any previously installed Content Search Engine Client 4.0.1 files from each Content Engine Server machine.

2. "Collecting current Content Engine deployment data"
   Content Engine deployment data provides confirmation that installation of
   Content Search Engine Client files on Content Engine Server machines has
   succeeded.
3. "Installing Content Engine Search Client files" on page 103
   You can install the Content Search Engine Client files interactively or silently.
   These files make it possible for the Content Search Engine server component to
   communicate with the Content Engine server component.

Removing the existing Content Search Engine Client 4.0.1 files from Content Engine Server machines

Remove any previously installed Content Search Engine Client 4.0.1 files from each
Content Engine Server machine.

You must explicitly remove any previously installed version 4.0.1 of the Content
Search Engine Client files because version 4.5.1 of the Content Search Engine Client
installer does not completely remove them. If you have no version 4.0.1 files, then
return to the parent topic and complete the task.

Remove the version 4.0.1 Content Search Engine Client files, as appropriate for
your operating system:

Option	Description
UNIX	Navigate to */InstallDevice*/FileNet/ ContentEngine/_uninst/cse, run ./uninstaller.bin, and follow the dialog boxes to remove the files.

Collecting current Content Engine deployment data

Content Engine deployment data provides confirmation that installation of Content
Search Engine Client files on Content Engine Server machines has succeeded.

To collect current Content Engine deployment data:

1. Enter the following address in a Web browser: http://
   ce_app_server_machine_name:*port* /FileNet/Engine *ce_app_server_machine_name*
   is the host name of the application server machine on which you are deploying
   Content Engine. *port* is one of the default port numbers used by the application
   server:

Option	Description
Application Server Type	Default Port Number
WebSphere	9080

2. Record the following data on the Web page,

Option	Description
Key	Value
Startup Message	Version number of Content Engine
Process Engine	Version number of Process Engine Client files

Installing Content Engine Search Client files

You can install the Content Search Engine Client files interactively or silently. These files make it possible for the Content Search Engine server component to communicate with the Content Engine server component.

"Installing Content Search Engine Client files interactively"
You can install the Content Search Engine Client files interactively by running the installation program on all application server machines where Content Engine Server is installed.

"Installing Content Search Engine Client files silently" on page 104
You can install the Content Search Engine Client files silently by editing a silent input file and running the installation program from a command line. You must install Content Search Engine Client files on all application server machines where Content Engine Server is installed.

Installing Content Search Engine Client files interactively:

You can install the Content Search Engine Client files interactively by running the installation program on all application server machines where Content Engine Server is installed.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install Content Search Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Search Engine Client values, filter by **CSE Client Installer** in the **Installation or Configuration Program** column.

2. 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.

3. On the machine where Content Engine is to be deployed, log on as *ce_install_user*.

4. Locate the Content Search Engine Client installation software. It must be the same version as your Content Search Engine Server software. You must install Content Search Engine and Content Engine software on the same operating system. For example, if Content Engine runs on Windows, locate the Content Search Engine Client installation software that runs on Windows.

5. Start the installation program:

Operating System	Install Program
AIX	P8CSE-CLIENT-4.5.1-AIX.BIN

6. Complete the Content Search Engine Client installation screens by using the values in your worksheet.

7. Review the following Content Search Engine Client log files:
   - UNIX - *cse_client_install_path*/CSEClient_install_log_4.5.1.0.log

Related concepts

⤷ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing Content Search Engine Client files silently:

You can install the Content Search Engine Client files silently by editing a silent
input file and running the installation program from a command line. You must
install Content Search Engine Client files on all application server machines where
Content Engine Server is installed.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

To install Content Search Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data ⟶ Filter ⟶ AutoFilter** command
   is enabled. To view only Content Search Engine Client values, filter by **CSE
   Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these
   updates might be required with other components and expansion products,
   contact your service representative.

3. On the machine where Content Engine is to be deployed, log on as
   ce_install_user.

4. Locate the Content Search Engine Client installation software. It must be the
   same version as your Content Search Engine Server software. You must install
   Content Search Engine and Content Engine software on the same operating
   system. For example, if Content Engine runs on Windows, locate the Content
   Search Engine Client installation software that runs on Windows.

5. Find the `CSE_silent_install.txt` file within the directory that contains the
   Content Search Engine Client installation software. Then edit the file to specify
   the location where you want to install the Content Search Engine Client files.

   The default installation paths are as follows:

   UNIX `/opt/IBM/FileNet/ContentEngine/CSEClient`

6. Start the silent installation of Content Search Engine Client files by running the
   following command.

Operating System	Install Program
AIX	`P8CSE-CLIENT-4.5.1-AIX.BIN —f` `CSEClient_silent_install.txt —i silent`

7. Review the following Content Search Engine Client log files:
   - UNIX - *cse_client_install_path*/`CSEClient_install_log_4.5.1.0.log`

Related concepts

➦ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Granting directory permissions to the Configuration Manager user

You must grant the file and directory permissions required by *config_mgr_user*, which is the user who runs Configuration Manager.

To grant permissions to the Configuration Manager user:

1. Log on as *ce_install_user* to the server where you installed Content Engine.
2. Grant permissions to the *config_mgr_user* user for the executable file that you intend to use:

Option	Description
UNIX - graphical user interface	Grant execute permission to `configmgr`
UNIX - command line	Grant execute permission to `configmgr_cl`

3. Grant write permission to the directory where you want Configuration Manager to place the configuration XML files that it creates.

   If you do not specify the profile directory when you run Configuration Manager, grant write permission on the default directory `ce_install_path/tools/configure` and all its files and subdirectories.
4. Log off the Content Engine server and log on again as *config_mgr_user*.

Configuring Content Engine instances upgraded from version 4.0.1 or 4.5.0

You can configure your upgraded Content Engine instances with Configuration Manager. You must upgrade the Content Engine bootstrap properties for the upgraded Content Engine instance. A single Content Engine application instance equates to one deployed application on your application server.

Configuration is a multiple step process. You must provide information about your Content Engine application environment, then you apply the settings by running the configuration tasks, and then you deploy the application. You can configure multiple instances before you deploy any of them, or you can configure and deploy one instance at a time.

Remember the following points when you prepare to configure Content Engine:
- (WebSphere Application Server only) For best results, configure no more than one Content Engine instance in a WebSphere profile.
- If you need an accessible software version of Configuration Manager for people with disabilities, use the command-line version of Configuration Manager.
- If your FileNet P8 domain uses multiple non-managed application servers, then you use a copy of a single EAR file with the bootstrap settings on all the servers. After you complete the tasks on the initial server, you copy the EAR file with the bootstrap settings to the additional servers.

  "Configuring upgraded instances by using the graphical user interface" on page 106

You can configure a Content Engine instance by using the graphical user interface version of Configuration Manager.

"Configuring upgraded instances by using the command line" on page 109
You can configure an upgraded Content Engine instance on a web application server by using the command-line version of Configuration Manager.

Configuring upgraded instances by using the graphical user interface

You can configure a Content Engine instance by using the graphical user interface version of Configuration Manager.

See the appendix "Configuration Manager user interface reference" on page 222 for complete information about using the graphical user interface. If you need an accessible version of Configuration Manager, use the command-line interface instead of the graphical user interface.

1. "Starting the Configuration Manager graphical user interface"
   You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.
2. "Changing the password save preference"
   By default, any passwords that you enter in the graphical user interface in Configuration Manager are not saved to a file. The passwords are stored in memory as long as you have the profile open in Configuration Manager. You can change the password save preference to save the passwords each time that you save your changes to a profile.
3. "Creating an upgrade configuration profile" on page 107
   To upgrade Content Engine Server, you must use Configuration Manager to create an upgrade configuration profile. The profile contains deployment, authentication, and database-connection information.
4. "Configuring the upgraded Content Engine instance" on page 108
   You must provide the bootstrap property values to upgrade Content Engine.

Starting the Configuration Manager graphical user interface:

You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.

To start Configuration Manager:

1. 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:

Option	Description
UNIX	Run this command:
	ce_install_path/tools/configure/configmgr

   The first time that you start Configuration Manager, the Welcome is displayed.
2. Select one of the links in the Welcome to learn more or to start working in a wizard, or close the Welcome by clicking the **X** in the tab at the upper left. You can reopen the Welcome later, as needed, from the **Help** menu.

Changing the password save preference:

By default, any passwords that you enter in the graphical user interface in Configuration Manager are not saved to a file. The passwords are stored in

memory as long as you have the profile open in Configuration Manager. You can change the password save preference to save the passwords each time that you save your changes to a profile.

When you close the profile, the passwords are erased from memory. Each time that you start Configuration Manager or open a saved profile, you must specify any passwords required by the tasks and for the application server properties before you run the tasks; otherwise, the tasks won't run successfully. If your site security requirements permit you to save passwords to a file, you can change the password save preference setting.

To change the password save preference:
1. Click **Window** → **Preferences**.
2. Complete one of the following actions:

Option	Description
To save passwords to file	Select the **Save all passwords to file when saving a task or profile** check box.
To prevent writing passwords to file	Clear the **Save all passwords to file when saving a task or profile** check box.

3. Click **OK**.

Creating an upgrade configuration profile:

To upgrade Content Engine Server, you must use Configuration Manager to create an upgrade configuration profile. The profile contains deployment, authentication, and database-connection information.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Create Upgrade Profile** in the Installation or Configuration Program column.
2. Start the Create Upgrade Configuration Profile wizard by using one of the following methods:
   - Select **File** → **New Upgrade Profile**.
   - Click the **Create a profile for an upgrade installation** icon in the toolbar.
3. Complete the configuration screens by referring to the values in your worksheet.
4. The wizard screens that are displayed depend on whether an existing Servers.xml file was found in the *ce_install_path*\ContentEngine directory, where *ce_install_path* is the location where Content Engine is installed.

Option	Description
Servers.xml contains 1 server	The upgrade wizard reads the server information from the Servers.xml file. You must name the profile and provide a location for the profile. If the Servers.xml file does not contain information on the deployment type, you must select a deployment type. The wizard creates a profile with the Configure Bootstrap Properties task and the Deploy Application task.
Servers.xml contains more than 1 server	The upgrade wizard reads the server information from the Servers.xml file. You must select a single server to upgrade from the list of servers that was found in Servers.xml. You must name the profile and provide a location for the profile. If the Servers.xml file does not contain information on the deployment type, you must select a deployment type. The wizard creates a profile for the server that you selected with the Configure Bootstrap Properties task and the Deploy Application task.
Servers.xml was not found	The upgrade wizard cannot determine any information about your existing server. You must name the profile, provide a location for the profile, and provide the application server properties. The wizard creates a profile with the Configure Bootstrap Properties task and the Deploy Application task.

5. The upgrade profile is displayed as an icon in the left pane with icons for the Configure Bootstrap and Deploy Application tasks.

Related concepts

▶ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring the upgraded Content Engine instance:

You must provide the bootstrap property values to upgrade Content Engine.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ⁓ **Filter** ⁓ **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Upgrade Bootstrap** in the Installation or Configuration Program column.

2. With your upgrade profile open, right-click **Configure Bootstrap Properties** in the profile pane, and select **Edit Selected Task**.
3. In the **Bootstrap operation** field, select **Upgrade**.

4. In the **Bootstrapped EAR file location** field, enter the fully qualified path to the currently deployed Content Engine EAR file.

5. Select **File → Save** to save your changes.

6. Apply the bootstrap property settings by right-clicking **Configure Bootstrap Properties** in the profile pane, and selecting **Run Task**. Running the configuration task can take a few minutes. The task execution status messages are displayed in the Console pane below the bootstrap properties.

Related concepts

⬆ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring upgraded instances by using the command line

You can configure an upgraded Content Engine instance on a web application server by using the command-line version of Configuration Manager.

Ensure that you have the required accounts and appropriate permissions before you begin.

To configure the upgraded Content Engine instance:

1. Log on to the machine where Content Engine is installed as *config_mgr_user*, which is the Configuration Manager user.

2. Change 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 **generateupgrade** command to generate the upgrade profile. Enter the command on one line without line breaks.

   ```
   configmgr_cl generateupgrade [-appserver app_server_type] -server server_name
    [-deploy deploy_type] -profile myprofile [-silent] [-force] [-help]
   ```

   Where:

   -appserver *appserver_type*
       The **-appserver** *appserver_type* parameter is required only if the *ce_install_path*/Servers.xml file does not exist or does not contain any servers. This parameter specifies the type of application server and must be one of the following values: WebSphere, WebLogic, or JBoss.

   -server *server_name*
       The **-server** *server_name* parameter indicates the server to create the upgrade profile for. Use the **listservers** command to display a list of valid *server_name* values before you run the **generateupgrade** command.

   -deploy *deploy_type*
       The **-deploy** *deploy_type* parameter is required only if the *ce_install_path*/Servers.xml file does not exist or does not contain the deployment type for the selected server. This parameter specifies the type of Content Engine deployment. If a value for the deployment type exists in the Servers.xml file for the selected server, then the value from the Servers.xml file is used instead the value you specify on the command line.

       Select one of the following values: standard, cluster, or netdeploy (network deployment).

Specify `standard` if you are deploying Content Engine to a stand-alone (that is, a server that is neither managed nor clustered) WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server.

Specify `cluster` if you are deploying Content Engine to a WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server cluster

Specify `netdeploy` if you are deploying Content Engine to a managedWebSphere Application Server instance. That is, you are using Network Deployment to manage individual servers that are not necessarily in a cluster.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
> - The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ ContentEngine/tools/configure/profiles/ce_was_tiv_db2/ ce_was_tiv_db2.cfg.
>
> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg".

-silent
> The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-help
> The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

The `upgrade.xml` file and one directory named for each existing Content Engine instance is created in the profile directory.

Important: If the **generateupgrade** command returns a message saying that an upgrade profile already exists, skip to 4.

4. (IBM WebSphere Application Server or Oracle WebLogic Server only) Set the application server administrator properties:
   a. Navigate to the *myprofile*/*mycontent_engine_instance* directory.
   b. Open the `applicationserver.xml` file in a text editor.

c. Set the value for the `ApplicationServerAdminUsername` property.

   ```
   <propertyname="ApplicationServerAdminUsername">
    <value>administrator_user_name </value>
   </property>
   ```

   d. Save your changes and close the file.

   e. Run the **storepasswords** command, which prompts you (respond with `yes` or `no`) to choose whether to encrypt and store the web application server administrator password to the `applicationserver.xml` file. If you choose to not save the password, you will need to edit `applicationserver.xml` by specifying the clear text of the password for the ApplicationServerAdminPassword property.

   ```
   configmgr_cl storepasswords -profile myprofile
   ```

   f. Repeat this step as needed to edit the application server properties for each Content Engine instance directory that was created.

5. At a command prompt, run the **execute** command.

   ```
   configmgr_cl execute -task configurebootstrap -profile myprofile
   ```

   Status messages are displayed in the console.

6. Optional: At a command prompt, check the completion status of the task by running the **checkStatus** command.

   ```
   configmgr checkStatus -task configurebootstrap -profile myprofile
   ```

Related concepts

"Configuration Manager command-line reference" on page 238
Configuration Manager can be run from a command line. This section covers the syntax for the command-line version of Configuration Manager.

Specifying IBM FileNet P8 accounts for upgrades
For details on accounts and required permissions, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing the latest Process Engine client files on Content Engine servers

You must upgrade the Process Engine Client files on all application server machines where Content Engine Server is deployed.

"Upgrading the latest Process Engine client files on Content Engine servers interactively"
The upgrade wizard provides an interactive way to upgrade the Process Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

"Upgrading the latest Process Engine client files on Content Engine servers silently" on page 112
You can upgrade the Process Engine Client software on the Content Engine server by using a silent input text file and running the upgrade program from the command line. Use the values in your installation and upgrade worksheet to edit the silent input text file before you run the upgrade.

Upgrading the latest Process Engine client files on Content Engine servers interactively

The upgrade wizard provides an interactive way to upgrade the Process Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To upgrade the Process Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. On the machine where Content Engine is to be deployed, log on as *ce_upgrade_user*.

4. 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.

5. Expand the (TAR or ZIP) Process Engine Client install software.

6. Shut down all instances of Enterprise Manager and any other Content Engine client applications, such as Application Engine.

7. Run the command for the installation program according to your operating system.

Operating System	Install Program
AIX	P8PE-CLIENT-4.5.1-AIX.BIN

8. Complete the Process Engine client install screens.

Related concepts

☞ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading the latest Process Engine client files on Content Engine servers silently

You can upgrade the Process Engine Client software on the Content Engine server by using a silent input text file and running the upgrade program from the command line. Use the values in your installation and upgrade worksheet to edit the silent input text file before you run the upgrade.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install the Process Engine Client files silently:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. On the machine where Content Engine is to be deployed, log on as
   ce_upgrade_user.
4. 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.
5. Expand the (TAR or ZIP) Process Engine Client install software.
6. Shut down all instances of Enterprise Manager and any other Content Engine
   client applications, such as Application Engine.
7. In the expanded install software, open the file `PEClient_silent_install.txt`
   and edit it as follows: Change the `Variable_CheckboxCE` line to the following
   line:

   `-V Variable_CheckboxCE="true"`
8. Save your edits.
9. Run one of the commands in the following table to perform the silent install:

Operating System	Install Program
AIX	`P8PE-CLIENT-4.5.1-AIX.BIN -silent` `-options "PEClient_silent_install.txt"`

Related concepts

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Deploying upgraded Content Engine instances

You must deploy the upgradedContent Engine on the web application server.
Deploying makes the Content Engine application available for use.

Restriction: Use the command-line version of Configuration Manager if any of
these conditions is true:

- You need an accessible software version of Configuration Manager for people
  with disabilities to use.

  "Deploying upgraded instances by using the graphical user interface"
  You can deploy an upgraded Content Engine instance on the web application
  server by using the graphical user interface version of Configuration Manager.
  Deploying makes the Content Engine application available for use.

  "Deploying upgraded instances by using the command line" on page 115
  You can deploy an upgraded Content Engine instance on a web application
  server by using the command-line version of Configuration Manager.
  Deploying makes the Content Engine application available for use.

Deploying upgraded instances by using the graphical user interface

You can deploy an upgraded Content Engine instance on the web application
server by using the graphical user interface version of Configuration Manager.
Deploying makes the Content Engine application available for use.

Make sure that the application server is running or stopped, depending on its type:

Table 16. Required Application Server State

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

For more information about the properties and values you set in the Configuration Manager, roll your mouse over the property name to view the hover help for the property. See the appendix "Configuration Manager user interface reference" on page 222 for complete information about using the graphical user interface.

To deploy Content Engine:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Configuration Manager values, filter by **CM: Deploy Application** in the Installation or Configuration Program column. For general information about the worksheet file, such as filtering details, see "Using the installation and upgrade worksheet" in *Plan and Prepare Your Environment for IBM FileNet P8*.

2. Log on to the application server machine as *config_mgr_user*, the Configuration Manager user.

3. 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:

Option	Description
UNIX	*ce_install_path*/tools/configure/configmgr

4. If you upgraded your application server and then redeployed Content Engine before upgrading Content Engine, click the **Edit Application Server Properties** icon and correct any obsolete parameter values for your application server type.

5. Select **File** → **Open Configuration Profile**, and use the information in your worksheet to specify the values for the parameters in the wizard screens.

6. In the Open Configuration Profile wizard, enter the path to the profile for the Content Engine instance, and then click **Finish**.

7. Right-click the **Deploy Application** task in the profile pane (left pane), and select **Edit Selected Task**.

8. Provide the property values for your deployment, by using the values in your worksheet.

9. Select **File** → **Save**.

10. Enable the task. By default, the **Deploy Application** task is disabled. When the task is disabled, the task name includes the text **(Disabled)**. To enable the task, select **Deploy Application (Disabled)** in the profile pane, and then either right-click and choose **Enable Selected Task** from the context menu, or click the **Enable the Selected Task** icon in the task toolbar.

11. Right-click the **Deploy Application** task in the left pane, and select **Run Task**.

    Running the deploy task can take a few minutes. The task execution status messages are displayed in the Console pane below the deploy application properties.

12. Repeat this procedure as needed for each Content Engine instance.

Related concepts

➡ Specifying IBM FileNet P8 accounts for upgrades
For details on accounts and required permissions, see *Plan and Prepare Your Environment for IBM FileNet P8*.

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Deploying upgraded instances by using the command line

You can deploy an upgraded Content Engine instance on a web application server by using the command-line version of Configuration Manager. Deploying makes the Content Engine application available for use.

Ensure that the application server is running or stopped:

Table 17. Required Application Server State

Option	Description
WebSphere Application Server	Start the application server if it is not already running.

See "Configuration Manager command-line reference" on page 238 for syntax details.

To deploy a Content Engine instance from the command line:

1. Log on to the application server machine as the *config_mgr_user* user, which is the Configuration Manager user.
2. If you upgraded the version of your application server, or if you moved it to a new directory or device and then redeployed Content Engine before upgrading Content Engine to version 4.5.1 from version 4.0.1 or 4.5.0, 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.
3. Set the current directory to *ce_install_path*/tools/configure, where *ce_install_path* is the path where you installed Content Engine.
4. Run the following command to apply the settings in the `deployapplication.xml` file:

   `configmgr execute -task deployapplication -profile myprofile`

   where the *myprofile* value can be one of the following items:
   - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
   - The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
   - The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ce_was_tiv_db2.cfg"`.

5. Optional: Run the following command to check the completion status of the DeployApplication task.

   `configmgr checkstatus -task deployapplication -profile *myprofile*`

6. Repeat this procedure as needed to deploy any additional Content Engine instances.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing storage device source files

If your upgraded IBM FileNet P8 system includes Tivoli Storage Manager or EMC Centera devices, you must install files on the Content Engine server.

"Installing Tivoli Storage Manager client and adding native API library paths (WebSphere)"
The client software and the native API library paths enable you to use Tivoli Storage Manager to access a fixed content device.

"Installing or updating EMC Centera SDK library files" on page 117
EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.

Installing Tivoli Storage Manager client and adding native API library paths (WebSphere)

The client software and the native API library paths enable you to use Tivoli Storage Manager to access a fixed content device.

1. "Installing Tivoli Storage Manager client"
   Install the Tivoli Storage Manager client software on each application server where Content Engine is deployed.

2. "Copying the Tivoli Storage Manager API libraries to additional servers" on page 117
   If you are running a Content Engine server farm, the Tivoli Storage Manager API libraries must be on each server in the farm.

3. "Creating a shared library definition for Tivoli Storage Manager native API library files" on page 117
   The shared library definition specifies the location of the Tivoli Storage Manager JAR file.

Installing Tivoli Storage Manager client:

Install the Tivoli Storage Manager client software on each application server where Content Engine is deployed.

1. Download the Tivoli Storage Manager client software from the IBM Support site at http://www.ibm.com/support/docview.wss?uid=swg24019757.

2. Complete the platform-specific installation instructions included with each Tivoli Storage Manager client download package. Record the path where you install Tivoli Storage Manager client because you will need to specify the path as the DSMI directory when you create a Tivoli fixed content device in FileNet Enterprise Manager. If you are installing Tivoli Storage Manager client on

multiple UNIX hosts, the installation path must be the same on each host. On AIX, install Tivoli Storage Manager client at `/usr/tivoli/tsm/client/ba/bin`.

Copying the Tivoli Storage Manager API libraries to additional servers:

If you are running a Content Engine server farm, the Tivoli Storage Manager API libraries must be on each server in the farm.

Copy the entire `tsm100` directory structure from the Content Engine installation directory to each of the servers in the farm. It is a best practice to use the same directory structure on each server in the farm. For example:

Option	Description
UNIX	`/opt/IBM/FileNet/ContentEngine/tsm100`

Creating a shared library definition for Tivoli Storage Manager native API library files:

The shared library definition specifies the location of the Tivoli Storage Manager JAR file.

1. Log on to the WebSphere administrative console.
2. Create a shared library definition according to the type of deployment (stand-alone or clustered), the version of WebSphere, and the operating system on which WebSphere runs, as shown in the following substeps. For the detailed procedure on creating the shared library, visit http://www.ibm.com/software/webservers/appserv/was/library/.
   a. Specify a Node scope for the library.
   b. In a server farm, if you installed the `tsm100` directory in different locations, choose Server scope and add a Shared Library entry for each server in your server farm.
   c. Provide a name for the shared library, for example `TSMAPLIB`.

Installing or updating EMC Centera SDK library files

EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.

Update the EMC Centera SDK library files if Centera fixed content devices are in your existing environment or if Centera fixed content devices will be added to your IBM FileNet P8 environment.

1. "Installing EMC Centera SDK version 3.2 library files"
   EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.
2. "Configuring EMC Centera SDK environment variables for version 3.2" on page 118
   You must specify values for the environment variables so that Content Engine can access the EMC Centera SDK library files.

Installing EMC Centera SDK version 3.2 library files:

EMC Center SDK library files enable access between the IBM FileNet P8 environment and EMC Centera fixed content devices.

1. Log on to the Content Engine Server machine as *ce_install_user*.

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

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

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.

Option	Description
UNIX	install.sh

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.

Configuring EMC Centera SDK environment variables for version 3.2:

You must specify values for the environment variables so that Content Engine can access the EMC Centera SDK library files.

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:

Option	Description
UNIX	setCenteraLibPath.sh

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`

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

4. Stop and start the application server.

Completing Content Engine post-deployment steps

You must complete several Content Engine 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, and then perform the remaining verification procedures.

When you deploy Content Engine after upgrading from version 4.0.1 or 4.5.0, an upgrader servlet automatically upgrades the Global Configuration Data (GCD) database and the object stores.

"Completing post-deployment steps (WebSphere Application Server)"
If you are using IBM WebSphere Application Server federated user repositories, you must specify the primary administrative user for the web application server.

Completing post-deployment steps (WebSphere Application Server)

If you are using IBM WebSphere Application Server federated user repositories, you must specify the primary administrative user for the web application server.

To set the administrative user

1. Log in to the WebSphere administrative console.

2. Navigate to **Security** → **Secure administration, applications, and infrastructure**, and click **Configure**.
3. If you have not already done so, specify a unique user (short name) for the **Primary administrative user name**.

   Important: This name must exist in one of the realms and must be unique.
4. Specify the **Server user identity**. You can specify **Automatically generated server identity**, or specify a user name that exists in one of the repositories. This name must exist in one of the realms and must be unique.
5. Save your changes to the master configuration.
6. Restart the web application server where Content Engine is deployed, as follows:
   - For a stand-alone server, stop and start the application server.
   - For a network deployment on WebSphere Application Server Network Deployment, stop and start the application server whereContent Engine is deployed, including the Deployment Manager and all managed nodes.
   - For a cluster, stop and start the cluster.

Verifying the Content Engine deployment

You can verify that the Content Engine deployment was successful by accessing the IBM FileNet P8 System Health page.

1. Verify the state of Content Engine deployment by browsing to the IBM FileNet P8 System Health page:

   `http://server:port/P8CE/Health`

   where:

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

   port is the WSI port used by the Web application server on the machine where Content Engine is deployed.

   The following table lists an example address for your application server:

Table 18. Example FileNet P8 System Health page address

Application Server Type	Web Page Address
IBM WebSphere Application Server	`http://server:9080/P8CE/Health`

2. Verify that the FileNet P8 System Health page contains the Content Engine instance host name and port number. The FileNet P8 System Health page provides status for the items in the following table.

   Tip: At this point, red icons appear on the page at the left of those entries that do not yet exist. In particular, if you are upgrading from version 3.5.2, a red icon appears before the GCD, directory configurations, object stores, storage areas, and sites entries. Also, if you are upgrading from version 3.5.2, no domain name appears yet in the Domain entry.

Table 19. FileNet P8 System Health page contents

Section	Description
Domain	Displays the FileNet P8 domain name if a domain was found.

Table 19. FileNet P8 System Health page contents (continued)

Section	Description
Global Configuration Database	Verifies that the GCD contains a valid domain object, that the XA and non-XA data sources are defined and have unique names, and that the bootstrap user name and password are defined. If any of these verification tests fail, then the failed icon is displayed.
Directory Configurations	Verifies that at least one directory service is configured and lists the number of configured directory service providers.
PE Connection Points	Displays the number of Process Engine connection points.
PE Isolated Regions	Displays the number of Process Engine isolated regions.
Fixed Content Devices	Lists the number of fixed content devices (FCDs) For each FCD listed, there is at least one associated fixed storage area that is in the open state.
Object Stores	Verifies that at least one object store exists and lists the number of object stores.
Storage Areas	Verifies that at least one storage area is defined, lists the number of storage areas, and the status for each storage area, such as online or offline.
Content Cache Areas	Displays the number of content cache areas.
Sites	Verifies that at least one site is defined and lists the number of sites. For each site listed, at least one virtual server or server instance exists.

3. Optional: Bookmark the FileNet P8 System Health page Web page address in your browser for later use.
4. Browse to the FileNet Engine Web page:

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

   where:

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

   port is the HTTP port used by the application server where Content Engine is deployed.

   The following table lists an example address for your application server:

Table 20. Example FileNet Engine Web page address

Application Server Type	Web Page Address
IBM WebSphere Application Server	http://myserver:9080/FileNet/Engine

5. Verify that the FileNet Engine Web page contains the following information:
   - Verify the value in the **Startup Message** key. The Content Engine build and version (for example, dap451.097), must match the build and version in the `ce_install_path`/`ce_version.txt` file, where *ce_install_path* is the location where the Content Engine software is installed.
   - Verify the value in the **Process Engine** key. The Process Engine Client build and version (for example, pui451.075), must match the build and version in the `version.txt` file in the TAR or ZIP file of the Process Engine Client installation software.
   - Verify that the values for JDBC driver, server instance, operating system, and other properties match the values that you entered when you configured Content Engine.
6. Optional: Bookmark the FileNet Engine Web page in your browser for later use.

Checking the automatic upgrade status for Content Engine data

When you deploy the upgraded Content Engine application on the web application server and start the application, the existing Content Engine data upgrade begins. You can check the progress on the upgrade to confirm that all object stores have been upgraded successfully

To check the status of the automatic data upgrade:

1. Enter the URL for the automatic upgrade status page into your browser.
   `http://server_name:port/FileNet/AutomaticUpgradeStatus`
2. Review the status for the FileNet P8 domain or object store.

Table 21. Automatic upgrade status descriptions

Status	Description
Ready	The FileNet P8 domain or object store is ready for use. No automatic upgrade work was performed since theContent Engine server was started.
Completed	The FileNet P8 domain or object store is ready for use. Automatic upgrade work was performed after the Content Engine server was started. Click **Completed** to view the detailed status for the object store upgrade process. In a multi-server configuration, automatic upgrade work might be in progress on additional servers.
In Progress	The FileNet P8 domain or object store is not available. Automatic upgrade work is currently being performed. The Schema Revision column in the Status for object stores section displays the upgrade phase that was completed. Content Engine version 4.5.1 uses schema revision 15.1.23. Click **Completed** to view the detailed status for the object store upgrade process.
Not Available	The FileNet P8 domain or object store is not available. The Content Engine server cannot determine the status of the object store. The server is still starting the Content Engine application and cannot access the object store to determine its status, or the object store did not load the object store cache. Wait a while for the server to finish starting. If the server cannot load the object store cache, then the object store cannot be upgraded automatically. The automatic upgrade does not work for upgrades from version 3.5.2 to version 4.5.1.

Asynchronous Upgrade In Progress	Object stores only. The object store is available. Automatic upgrade work was completed and additional asynchronous upgrade work is in progress. The work involves tasks that take a long time to complete, usually due to the large number of documents or custom objects in the object store.

Installing the latest Content Engine client files on other IBM FileNet P8 servers (for staged upgrades)

When you upgrade Content Engine 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:

- "Upgrading the Content Engine client files on Process Engine servers" on page 148
- "Installing the latest Content Engine client files on Application Engine servers" on page 169

Important:

- You must install the Content Engine client files for any expansion products that use Content Engine Client files, such as:
  - IBM FileNet Workplace XT
  - IBM InfoSphere Enterprise Records
  - IBM FileNet Business Process Framework
- You must redeploy any components that are web-application-server-based, such as Application Engine, Workplace XT, and InfoSphere Enterprise Records, after you upgrade the Content Engine Client files on the associated machines.

Upgrading Content Search Engine software from version 4.0.1 or 4.5.0

If you used the full-text search feature (CBR) in version 4.0.1 or 4.5.0, to upgrade you must export your current K2 server configurations on each machine, install the new version of Content-Search Engine (Autonomy K2) software, and then import your saved configuration.

1. "Production environment upgrades" on page 124
   Upgrading of Content Search Engine and related Content Engine data must occur as an in-place upgrade. As such, the data is migrated to 4.5.1 and is no longer usable in the previous version of the IBM FileNet P8 Platform.

2. "Upgrading Content Search Engine"
   Upgrade Content Search Engine on all machines running Autonomy K2
   software. Upgrade the K2 Master Administration Server first, before upgrading
   the K2 Administration Servers.

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 4.5.1 and is no longer
usable in the previous version of the IBM FileNet P8 Platform.

The index areas are upgraded at the same time as the Content Engine data is
upgraded.

To perform an upgrade in a migration scenario, create a duplicate of your IBM
FileNet P8 system by performing a backup and then a restore. Using the restored
system, complete the migration to 4.5.1. 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.

Upgrading Content Search Engine

Upgrade Content Search Engine on all machines running Autonomy K2 software.
Upgrade the K2 Master Administration Server first, before upgrading the K2
Administration Servers.

If you have a multi-server configuration, do not start any K2 services until after all
machines have been upgraded.

"Upgrading Content Search Engine software (UNIX)"
To upgrade Content Search Engine, you must upgrade the underlying
Autonomy K2 software.

"Updating the style file of user fields" on page 129
You must edit the Autonomy K2 style file of user fields to add a data field
definition for storing the partitioning value of a document.

Upgrading Content Search Engine software (UNIX):

To upgrade Content Search Engine, you must upgrade the underlying Autonomy
K2 software.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

Important: If you have a multiserver configuration, do not start any K2 services
until after all Content Search Engine servers have been upgraded. You will need to
upgrade the Master Administration Server first.

"Upgrading Content Search Engine interactively (UNIX)" on page 125
To upgrade Content Search Engine interactively, you must run an installation
program wizard.

"Upgrading Content Search Engine silently (UNIX)" on page 127
When you upgrade Content Search Engine silently, you run a command in the
Content Search Engine software package that references parameters in an
installation text file.

Related concepts

↪ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Content Search Engine interactively (UNIX):

To upgrade Content Search Engine interactively, you must run an installation program wizard.

An X terminal must be installed on the Content Search Engine server to run the interactive installation program.

To upgrade Content Search Engine:
1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Content Engine values, filter by **CSE Installer** in the **Installation or Configuration Program** column.
2. Access the Autonomy K2 machine and log on as *k2_os_user*.
3. To stop the Autonomy K2 Administration Server service and the Tomcat application server, run the following command.

Option	Description
AIX	`verity_install_path/k2/_rs6k43/bin/k2adminstop`

4. Copy the IBM FileNet Content Search Engine installation package to the server.
5. Change directory to the location of the IBM FileNet Content Search Engine installation package.
6. Run the following commands:

   `backup_k2.sh "verity_install_path" "temporary_location_path"`

   where:

   "verity_install_path" is the path to the `verity` directory in the current installation path.

   "temporary_location_path" is the location in which you want to temporarily back up your Autonomy K2 configuration files, outside of the current installation path, and not used by any other application. Files from this backup are used to restore your existing configuration after you complete the software upgrade.

   For example, `backup_k2.sh "/opt/IBM/FileNet/ContentEngine/verity" "/opt/temp"`.
7. If K2 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. Execute the following command to determine which Verity processes are currently running on the machine:

   `ps -eaf | grep k2`
9. For each Verity process that is running, terminate it by executing the following command:

   `kill -9 pid`

where *pid* is the process ID.

10. (AIX only) Run the command **slibclean** as the root user to remove all unused libraries from shared memory.

11. Access the IBM FileNet Content Search Engine installation package.

12. Run the installation program.

Platform	Command
AIX	P8CSE-4.5.1.0-AIX.BIN

13. Complete the installation screens by using the values from your worksheet.

14. Review the following Content Search Engine log files for installation errors. The files are located in the `verity` directory of the installation path.
    - *verity_install_path*/verity/cse_install_log_4_5_1.txt
    - *verity_install_path*/verity/vconfig.log

15. To stop the Autonomy K2 Administration Server service and the Tomcat application server, run the following command.

Option	Description
AIX	*verity_install_directory*/k2/_rs6k43/bin/ k2adminstop

16. Change directory to the location of the IBM FileNet Content Search Engine installation package.

17. To restore your configuration from the backup files, run the following command:

    restore_k2.sh *"temporary_location_path"* *"verity_install_path"*

    where:

    "temporary_location_path" is the temporary path you entered earlier in this procedure when you ran the `backup_k2.sh` file.

    "verity install path" is the path to the `verity` directory in the Autonomy K2 installation path.

    For example, restore_k2.bat "/opt/temp" *"verity_install_path"*.

18. Add the following lines to the restored file *verity_install_path*/data/
stylesets/FileNet_FileSystem_PushAPI/style.ufl:

    ```
    data-table: fce
    {
    fixwidth: cemp_partition_date 4 unsigned-integer
    }
    ```

    Important: The purpose of this edit is to enable use of the Content Search Engine partitioning feature, new in version 4.5.1.

19. Run the following commands to set the setuid bit for the vspget program so that the service runs as root.

    Important: Most processes for the Autonomy K2 software run as *k2_os_user*. However, the vspget process must run as root. For details on required accounts and related permissions, see "Specify IBM FileNet P8 accounts" in *Plan and Prepare Your Environment for IBM FileNet P8*.
    - chown root *verity_install_path*/k2/*_platform*/bin/vspget
    - chmod u+s *verity_install_path*/k2/*_platform*/bin/vspget

    where *_platform* is the following directory.

Platform	Directory
AIX	_rs6k43

20. If you moved collections earlier in this procedure, restore your *collections directory* and *temp directory* to their original locations.

21. For a single-server configuration, start the K2 services. For multiserver configurations, start the services only after all servers have been upgraded.

   To start the Verity K2 Administration Server service and the Tomcat service, use the follow command.

Option	Description
AIX	nohup *verity_install_path*/k2/_rs6k43/ bin/k2adminstart

22. Access the K2 Dashboard and verify that the services from your previous installation are restored and running.

23. Delete the *temporary_location_path* directory of the backup you ran earlier.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Content Search Engine silently (UNIX):

When you upgrade Content Search Engine silently, you run a command in the Content Search Engine software package that references parameters in an installation text file.

To upgrade Content Search Engine silently:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Search Engine values, filter by **CSE Installer** in the **Installation or Configuration Program** column.

2. Access the Autonomy K2 machine and log on as *k2_os_user*.

3. To stop the Autonomy K2 Administration Server service and the Tomcat application server, run the following command.

Option	Description
AIX	*verity_install_path*/k2/_rs6k43/bin/ k2adminstop

4. Copy the IBM FileNet Content Search Engine installation package to the server.

5. (Master Administration Server only) Change directory to the location of the IBM FileNet Content Search Engine installation package.

6. Run the following commands:

   backup_k2.sh "*verity_install_path*" "*temporary_location_path*"

   where:

   "verity_install_path" is the path to the verity directory in the current installation path.

"temporary_location_path" is where you temporarily back up your Autonomy K2 configuration files, outside of the current installation path. *"temporary_location_path"*t must not and not used by any other application. Files from this backup are used to restore your existing configuration after you complete the software upgrade.

For example, `backup_k2.sh "/opt/IBM/FileNet/ContentEngine/verity/k2" "/opt/temp".`

7. If K2 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. Execute the following command to determine which Verity processes are currently running on the machine:

    `ps -eaf | grep k2`

9. For each Verity process that is running, terminate it by executing the following command:

    `kill -9 `*`pid`*

    where *pid* is the process ID.

10. (AIX only) Run the command **slibclean** as the root user to remove all unused libraries from shared memory.

11. Access the IBM FileNet Content Search Engine installation package.

12. Edit `CSE_silent_install.txt` to reflect the appropriate responses for your installation by using the values in your worksheet. You must install to the same path as your previous installation.

13. Run the following command:

    `./P8CSE-4.5.1.0-`*`UNIX_type`*`.BIN -f CSE_silent_install.txt -i silent`

14. Review the following Content Search Engine log files for installation errors. The files are located in the `verity` directory of the installation path.

    *   *verity_install_path*`/cse_install_log_4_5_1.txt`
    *   *verity_install_path*`/vconfig.log`

15. To stop the Autonomy K2 Administration Server service and the Tomcat application server, run the following command.

Option	Description
AIX	*verity_install_path*`/k2/_rs6k43/bin/` `k2adminstop`

16. Change directory to the location of the IBM FileNet Content Search Engine installation package.

17. To restore your configuration from the backup files, run the following command:

    `restore_k2.sh `*`"temporary_location_path"`*` `*`"verity_install_path"`*

    where:

    "temporary_location_path" is the temporary path you entered earlier in this procedure when you ran the `backup_k2.sh` file.

    "verity install path" is the path to the `verity` directory in the Autonomy K2 installation path.

    For example, `restore_k2.bat "/opt/temp" "/opt/IBM/FileNet/ContentEngine/verity".`

18. Add the following lines to the restored file *verity_install_path*`/data/`
`stylesets/FileNet_FileSystem_PushAPI/style.ufl`:

```
data-table: fce
{
fixwidth: cemp_partition_date 4 unsigned-integer
}
```

Important: The purpose of this edit is to enable use of the Content Search Engine partitioning feature, new in version 4.5.1.

19. Run the following commands to set the setuid bit for the vspget program so that the service runs as root.

    Important: Most processes for the Autonomy K2 software run as *k2_os_user*. However, the vspget process must run as root. For details on required accounts and related permissions, see "Specify IBM FileNet P8 accounts" in *Plan and Prepare Your Environment for IBM FileNet P8*.
    ```
    chown root verity_install_path/k2/_platform/bin/vspget chmod u+s
    verity_install_path/k2/_platform/bin/vspget
    ```
    where: *_platform* is the following directory.

Platform	Directory
AIX	_rs6k43

20. If you moved collections earlier in this procedure, restore your *collections directory* and *temp directory* to their original locations.
21. For a single-server configuration, start the K2 services. For multiserver configurations, start the services only after all servers have been upgraded.

    To start the Verity K2 Administration Server service and the Tomcat service, use the follow command.

Option	Description
AIX	nohup *verity_install_path*/k2/_rs6k43/ bin/k2adminstart

22. Access the K2 Dashboard and verify that the services from your previous installation are restored and running.
23. Delete the *temporary_location_path* directory of the backup you ran earlier.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Updating the style file of user fields:

You must edit the Autonomy K2 style file of user fields to add a data field definition for storing the partitioning value of a document.

To edit the style file:
1. Navigate to the directory in the Content Search Engine software that contains the `style.ufl` file:

    UNIX *verity_install_path*/data/stylesets/FileNet_FileSystem_PushAPI
2. Open the Autonomy K2 style file of user fields, `style.ufl`, for editing.
3. If it is not already present, add the CEMP definition at the bottom of the file:

```
data-table: fce
{
fixwidth: cemp_partition_date 4 unsigned-integer
}
```

4. Save your edits and close the file.

Upgrading and configuring Process Engine

Upgrade the Process Engine by performing several pre-upgrade steps, installing the Process Engine software and performing multiple post-upgrade steps.

Important: Before you start to upgrade the Process Engine software, ensure that you are at the correct version and fix pack level, as indicated in *Plan and Prepare Your Environment for IBM FileNet P8*.

The pre-upgrade and post-upgrade steps you perform will depend on which version of Process Engine software you are upgrading from.

1. "Completing 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.
2. "Upgrading Process Engine" on page 139
   Upgrade Process Engine software either interactively or silently.
3. "Upgrading the Content Engine client files on Process Engine servers" on page 148
   Installing or upgrading the Content Engine Client files on all Process Engine enables communication between the Content Engine and the Process Engine. For upgrades from 4.0, you must first remove the Content Engine Client files before upgrading them.
4. "Completing post-upgrade Process Engine configuration" on page 152
   You must perform several procedures to complete the upgrade of all Process Engine data and objects. The steps vary depending on what release you are upgrading from.
5. "Installing Process Engine software updates" on page 159
   After installing Process Engine, software updates such as fix packs or interim fixes might also be required.
6. "Installing the latest Process Engine client files on other IBM FileNet P8 servers (for staged upgrades)" on page 159
   When you upgrade Process Engine server software, you must also upgrade the associated Process Engine Client files installed on other machines running FileNet P8 components, for example, on Content Engine and Application Engine servers.

Related information

➡ General requirements for all FileNet P8 Platform systems
Check to make sure your software is at the correct fix pack level before you upgrade.

Completing 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.

"Completing pre-upgrade steps for upgrades from Process Engine Version 3.5" on page 132
Verify that prerequisites have been completed and run a number of pre-upgrade steps before upgrading Process Engine software from Version 3.5.

"Completing pre-upgrade steps for upgrades from Process Engine 4.0 or 4.5" on page 137

Verify that prerequisites have been completed and run a number of pre-upgrade steps before upgrading Process Engine software from Version 4.0 or 4.5.

Completing pre-upgrade steps for upgrades from Process Engine Version 3.5

Verify that prerequisites have been completed and run a number of pre-upgrade steps before upgrading Process Engine software from Version 3.5.

- Verify that you have a current backup of the Process Engine and Process Engine database.
- Verify that you have run the **vwverify** command on all regions to repair any inconsistencies between database queues and rosters.
- Verify that you have reconciled the Process Engine user security information. This reconciliation and repair of inconsistencies discussed above were a part of the steps described in Preparing Process Engine for upgrade in the Planning guide.
- Determine when to run the upgrade SQL script for Oracle databases. This script must be run by using one of the following methods:
  - Run the script manually before you run the Process Engine installation program.
  - Run the script automatically from the Process Engine installation program, which allows the installation program to prompt for the sys password for Oracle in an xterm window.
  - Run the script automatically from the Process Engine installation program, by running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local 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.

  Important: 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 by using the procedures in "Completing post-upgrade Process Engine configuration" on page 152.
- 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 Process Analyzer is installed, complete the steps in Verifying that all Process Engine 3.5 events have been transmitted and Verifying that all Process Analyzer 3.5 events have been published before purging event logs and statistics records.

1. "Updating all isolated regions" on page 133
   Update all isolated regions before upgrading the Process Engine software by performing a transfer of the upgrade.cdl file. This assures that there are no pending transfers.
2. "Completing pre-upgrade steps for Process Analyzer installations" on page 133
   For Process Analyzer production installations, perform the following tasks before upgrading Process Engine software from version 3.5 to 4.5.1.

3. "Purging event logs and statistics records" on page 135
   Optionally, purge event log and statistics records before you upgrade to reduce the time necessary for the database upgrade to complete.

4. "Stopping all Process Engine-related services and applications (UNIX)" on page 136
   Stop all Process Engine-related services and applications before upgrading from version 3.5.

Related tasks

➡ Preparing Process Engine for upgrade
For more information about preparing to upgrade, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Updating all isolated regions

Update all isolated regions before upgrading the Process Engine software by performing a transfer of the `upgrade.cdl` file. This assures that there are no pending transfers.

1. Restart the Process Engine software by entering the following command:
   `initfnsw -y restart`

2. Save any custom settings for CE_Operations including the login and password for the adaptor and any user exposed data fields.

3. Change directories to the following location of the *upgrade.cdl* file, depending on your operating system:

Option	Description
UNIX	`/fnsw/local/sd`

4. Initiate a transfer on all working isolated regions by entering the following command: `vwtfer -o upgrade.cdl -r XXX`

   where -r XXX indicates a region number for a single working region.

5. At the prompt, log on as *pe_service_user*.

6. Repeat for every isolated region.

7. Restore the custom settings for CE_Operations saved before running the transfer. You will need to restore these custom settings again after competing the steps in "Updating the Process Engine security" on page 154.

Completing pre-upgrade steps for Process Analyzer installations

For Process Analyzer production installations, perform the following tasks before upgrading Process Engine software from version 3.5 to 4.5.1.

1. "Verifying that all Process Analyzer 3.5 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.

2. "Verifying that all Process Engine 3.5 events have been transmitted to Process Analyzer" on page 134
   The steps in this procedure apply only if you are using a production Process Analyzer, they are not necessary for a simulation Process Analyzer.

Verifying that all Process Analyzer 3.5 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.

Verifying that all Process Engine 3.5 events have been transmitted to Process Analyzer:

The steps in this procedure apply only if you are using a production Process Analyzer, they are not necessary for a simulation Process Analyzer.

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 by 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 run a query on the Process Engine database by using the converted date/time value as one parameter in the SQL query.

To verify that all Process Engine 3.5 events have been transmitted to Process Analyzer 3.5:

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. Complete 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, run 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. Enter the following at the **vwtool** prompt:

   ```
   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 ran on Process Analyzer VMAEDM in 3a on page 134. (for example, 09/11/2006 16:23:59). Your input must match the format of the current system mask from 3d 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, run 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, by 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:

physical_table name is the table name from 3k.

The *PAInstallDate* is the number returned in 3f.

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.

Purging event logs and statistics records

Optionally, purge event log and statistics records before you upgrade to reduce the time necessary for the database upgrade to complete.

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 reduce the amount of time necessary for the upgrade to complete. Purging these records can take a significant amount of time, so plan this activity accordingly.

Run the command by logging on with the following service user name:
pe_service_user.
The following examples show the **vwlog** command syntax. If you use Process Analyzer, do not use the -P option when you purge logging records.

Table 22. vwlog command syntax

Description	Command
The following command removes the statistics from the specified isolated region.	`vwlog -X -r region number`
The following command removes all log records from the isolated region. Use this command only if all workflows terminated and you no longer need tracking or milestone information.	`vwlog -L -r region number`

You can use other parameters for the **vwlog** utility to select log records for deletion that meet specific conditions, such as log records for terminated workflows, for tracker related records, and so on. For more information, see the FileNet P8 help topics under **System Administration** ⟶ **Process Engine** ⟶ **Administration** ⟶ **Administrative tools** ⟶ **vwlog**.

Stopping all Process Engine-related services and applications (UNIX)

Stop all Process Engine-related services and applications before upgrading from version 3.5.

1. Log on as root.
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 P8 software is shut down:
   ```
   killfnsw -DAyS
   ```

   or
   ```
   killfnsw -D -A -y -S
   ```
5. (AIX only) Run the following command:
   ```
   slibclean
   ```

6. Run the following command to look for any java processes that are still running:

```
ps -ef | grep java | grep VW
```

7. Kill any FileNet P8-related running java processes.

Completing pre-upgrade steps for upgrades from Process Engine 4.0 or 4.5

Verify that prerequisites have been completed and run a number of pre-upgrade steps before upgrading Process Engine software from Version 4.0 or 4.5.

- Verify that you have a current backup of the Process Engine and Process Engine database.
- Verify that you have run **vwverify** on all regions to repair any inconsistencies between database queues and rosters.
- Determine when to run 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.
  - automatically, from the Process Engine installation program, allowing the installation program to prompt for the sys password for Oracle in an xterm window.
  - automatically, from the Process Engine installation program, running silently by using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

  "Updating all isolated regions"
  Update all isolated regions before upgrading the Process Engine software by performing a transfer of the `upgrade.cdl` file. This assures that there are no pending transfers.

  "Purging event logs and statistics records" on page 138
  Optionally, purge event log and statistics records before you upgrade to reduce the time necessary for the database upgrade to complete.

  "Stopping all Process Engine-related services and applications (Unix)" on page 138
  Stop all Process Engine-related services and applications before upgrading from version 4.0 or 4.5.

Related tasks

➥ Preparing Process Engine for upgrade
For more information about preparing to upgrade, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Updating all isolated regions

Update all isolated regions before upgrading the Process Engine software by performing a transfer of the `upgrade.cdl` file. This assures that there are no pending transfers.

1. Restart the Process Engine software by entering the following command:

```
initfnsw -y restart
```

2. Save any custom settings for CE_Operations including the login and password for the adaptor and any user exposed data fields.

3. Change directories to the following location of the *upgrade.cdl* file, depending on your operating system:

Option	Description
UNIX	`/fnsw/local/sd`

4. Initiate a transfer on all working isolated regions by entering the following command: `vwtfer -o upgrade.cdl -r XXX`

   where -r XXX indicates a region number for a single working region.

5. At the prompt, log on as *pe_service_user*.

6. Repeat for every isolated region.

7. Restore the custom settings for CE_Operations saved before running the transfer. You will need to restore these custom settings again after competing the steps in "Completing the upgrade from Process Engine 4.0, 4.5.0, or higher" on page 157.

Purging event logs and statistics records

Optionally, purge event log and statistics records before you upgrade to reduce the time necessary for the database upgrade to complete.

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 reduce the amount of time necessary for the upgrade to complete. Purging these records can take a significant amount of time, so plan this activity accordingly.

Run the command by logging on with the following service user name: *pe_service_user*.
The following examples show the **vwlog** command syntax. If you use Process Analyzer, do not use the `-P` option when you purge logging records.

Table 23. vwlog command syntax

Description	Command
The following command removes the statistics from the specified isolated region.	`vwlog -X -r region number`
The following command removes all log records from the isolated region. Use this command only if all workflows terminated and you no longer need tracking or milestone information.	`vwlog -L -r region number`

You can use other parameters for the **vwlog** utility to select log records for deletion that meet specific conditions, such as log records for terminated workflows, for tracker related records, and so on. For more information, see the FileNet P8 help topics under **System Administration → Process Engine → Administration → Administrative tools → vwlog**.

Stopping all Process Engine-related services and applications (Unix)

Stop all Process Engine-related services and applications before upgrading from version 4.0 or 4.5.

1. Log on as root.

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 Process Engine software is shut down.

   ```
   killfnsw -DAyS
   ```

   or

   ```
   killfnsw -D -A -y -S
   ```

6. (AIX only) Run the following command:

   ```
   slibclean
   ```

7. Run the following command to look for any java processes that are still running:

   ```
   ps -ef | grep java | grep VW
   ```

8. Kill any FileNet P8-related running java processes.

Upgrading Process Engine

Upgrade Process Engine software either interactively or silently.

"Upgrading Process Engine interactively"
Several Process Engine upgrade screens will display only if the upgrade is from Process Engine version 3.5. Those screens are noted in the Installation and Upgrade Worksheet.

"Upgrading Process Engine silently" on page 143
Modify the silent installation response file according to the values in the installation worksheet before upgrading the software.

Upgrading Process Engine interactively

Several Process Engine upgrade screens will display only if the upgrade is from Process Engine version 3.5. Those screens are noted in the Installation and Upgrade Worksheet.

You will find references to logging on as the root and fnsw users within the procedures to upgrade Process Engine software. 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.

Important: 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*.

To verify the database connection:

1. "Verifying the Process Engine database connection" on page 140
   Verify the ability to connect to the Process Engine database. The Process Engine installation program will connect to the database during software installation.

2. "Upgrading Process Engine interactively" on page 141
   You will find references to logging on as the root and fnsw users within the

procedures to upgrade Process Engine software. 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.

3. "Setting the f_maint and f_sw passwords" on page 142
Set or reset the passwords for the f_sw and f_maint users,. This will verify that the encrypted version matches the database or operating system version of the passwords.

4. "Restoring any custom modifications for root and fnsw users (UNIX)" on page 142
Process Engine installation creates 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.

5. "Editing the /etc/inittab file" on page 143
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.

Verifying the Process Engine database connection

Verify the ability to connect to the Process Engine database. The Process Engine installation program will connect to the database during software installation.

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

Verifying the Process Engine database connection (Oracle):

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

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

To verify the Process Engine database connection:
1. Run the following at a command prompt on the Process Engine:
```
su - oracle -c "sqlplus"
```
2. Enter one of the following commands at the SQLPlus prompt, as follows:
   - If the Process Engine pre-installation SQL scripts will be run from the Process Engine installation program by prompting for the sys password, type the following command:
   ```
   sys/password as sysdba
   ```
   - If the SQL scripts will be run from the Process Engine installation program by using operating system authentication, type the following command:
   ```
   / as sysdba
   ```
3. At the prompt, enter the following SQL command:
```
SQL> select instance_name, host_name, version
from v$instance;
```
The following represents an example of the information returned:

INSTANCE_NAME

HOST_NAME

VERSION

----

p8dbshr HQVWBUCS 10.2.0.2.0

where:

- *p8dbshr* is the instance ORACLE_SID.
- *HQVWBUCS* is the database server name.
- *10.2.0.2.0* is the Oracle server version.

Upgrading Process Engine interactively

You will find references to logging on as the root and fnsw users within the procedures to upgrade Process Engine software. 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.

Several Process Engine upgrade screens will display only if the upgrade is from Process Engine version 3.5. Those screens are noted in the Installation and Upgrade Worksheet.

"Upgrading the Process Engine software interactively (UNIX)"
You can upgrade the Process Engine by using the Process Engine installation wizard.

Upgrading the Process Engine software interactively (UNIX):

You can upgrade the Process Engine by using the Process Engine installation wizard.

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Process Engine values, filter by **PE Installer** in the **Installation or Configuration Program** column.

2. Log on to the Process Engine server as the root user. If you plan to run the SQL scripts from the Process Engine installation program, you must log on as a database administrator.

3. Navigate to the Process Engine software package.

4. From the console, launch the appropriate P8PE-4.5.1-*platform*.bin installation program and wait for the files to finish unpacking.

Platform	Command
AIX	P8PE-4.5.1-AIX.BIN

5. 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. Reboot when prompted.

6. Check for any .log files that were generated when the upgrade program ran. The files could be in any of the following locations, depending on whether the upgrade was successful and, if not, where errors occurred. Correct any errors or failures indicated before proceeding to the next step.

   - /fnsw/local/logs/PE (if the install completes successfully)
   - /fnsw/tmp_installer/PE451_setup.log (if the install has errors)

- /fnsw/local/logs/wizard/*date_time.log* and other subdirectories under /fnsw/local/logs

7. Log off as the root user and log on as *fnsw.*

Related concepts

→ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8.*

Setting the f_maint and f_sw passwords

Set or reset the passwords for the f_sw and f_maint users,. This will verify that the encrypted version matches the database or operating system version of the passwords.

> "Setting the f_maint and f_sw passwords (Oracle and Microsoft SQL Server)"
> 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 Microsoft SQL Server database. The encrypted password and the database user's passwords must match.

Setting the f_maint and f_sw passwords (Oracle and Microsoft SQL Server):

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 Microsoft SQL Server database. The encrypted password and the database user's passwords must match.

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

You must change the passwords for the f_maint and f_sw users. For Oracle and Microsoft SQL Server databases, both the encrypted file and the database passwords are updated.

On UNIX, several Process Engine utilities are configured with setuid *fnsw.* (or its alias), and will therefore run with that user's X Window settings, regardless of which user is currently logged on.

1. Log on as the fnsw user.
2. Start the Database Server Connect application by running the following command: Xdbconnect -r
3. Change the primary password for the users f_sw and f_maint to match the database password.
4. Exit the application.

Restoring any custom modifications for root and fnsw users (UNIX)

Process Engine installation creates 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

Editing 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.

Table 24. Changes to make to the /etc/inittab file

Operating system:	Changes to make:
AIX	Change:
	`rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1`
	to read:
	`:rcfnsw:2:once:/etc/rc.initfnsw 2>&1 | alog -tboot > /dev/console 2>&1`

Upgrading Process Engine silently

Modify the silent installation response file according to the values in the installation worksheet before upgrading the software.

Important: 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*.

1. "Verifying the Process Engine database connection" on page 144
   Verify the ability to connect to the Process Engine database. The Process Engine installation program will connect to the database during software installation.

2. "Upgrading Process Engine silently" on page 145
   You will find references to logging on as the root and fnsw users within the procedures to upgrade Process Engine software. 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.

3. "Setting the f_maint and f_sw passwords" on page 146
   Set or reset the passwords for the f_sw and f_maint users,. This will verify that the encrypted version matches the database or operating system version of the passwords.

4. "Restoring any custom modifications for root and fnsw users (UNIX)" on page 147
   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.
5. "Editing the /etc/inittab file" on page 147
   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.

Verifying the Process Engine database connection

Verify the ability to connect to the Process Engine database. The Process Engine installation program will connect to the database during software installation.

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

Verifying the Process Engine database connection (Oracle):

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

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

To verify the Process Engine database connection:
1. Run the following at a command prompt on the Process Engine:
   ```
   su - oracle -c "sqlplus"
   ```
2. Enter one of the following commands at the SQLPlus prompt, as follows:
   - If the Process Engine pre-installation SQL scripts will be run from the Process Engine installation program by prompting for the sys password, type the following command:
     ```
     sys/password as sysdba
     ```
   - If the SQL scripts will be run from the Process Engine installation program by using operating system authentication, type the following command:
     ```
     / as sysdba
     ```
3. At the prompt, enter the following SQL command:
   ```
   SQL> select instance_name, host_name, version
   from v$instance;
   ```
   The following represents an example of the information returned:
   INSTANCE_NAME

   HOST_NAME

   VERSION

   ----

   p8dbshr HQVWBUCS 10.2.0.2.0
   where:

- *p8dbshr* is the instance ORACLE_SID.
- *HQVWBUCS* is the database server name.
- *10.2.0.2.0* is the Oracle server version.

Upgrading Process Engine silently

You will find references to logging on as the root and fnsw users within the procedures to upgrade Process Engine software. 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.

Modify the silent installation response file according to the values in the installation worksheet before upgrading the software.

"Upgrading the Process Engine software silently (UNIX)"
Take the following steps to silently upgrade Process Engine.

"Encrypting passwords for Process Engine" on page 146
Several passwords are required to configure Process Engine components. To accommodate your security requirements, you can encrypt these passwords before you enter them into the configuration files.

Upgrading the Process Engine software silently (UNIX):

Take the following steps to silently upgrade Process Engine.

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Process Engine values, filter by **PE Installer** in the **Installation or Configuration Program** column.

2. Log on to the Process Engine server as the root user. If you plan to run the SQL scripts from the Process Engine installation program, you must log on as a database administrator.

3. Edit the `PE_silent_install.txt` file to reflect the appropriate responses for your installation. All passwords in the response file must be encrypted.

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

5. Navigate to the temporary directory on the local disk.

6. From the console, launch the appropriate `P8PE-4.5.1-`*platform*`.bin` installation program.

Platform	Command
AIX	`P8PE-4.5.1-AIX.BIN -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 for any `.log` files that were generated when the upgrade program ran. The files could be in any of the following locations, depending on whether the upgrade was successful and, if not, where errors occurred. Correct any errors or failures indicated before proceeding to the next step.
   - `/fnsw/local/logs/PE` (if the install completes successfully)
   - `/fnsw/tmp_installer/PE451_setup.log` (if the install has errors)
   - `/fnsw/local/logs/wizard/`*date_time.log* and other subdirectories under `/fnsw/local/logs`

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

Related concepts

↪ Encrypting passwords for Process Engine
For more information about encrypting passwords for Process Engine, see the *Plan and Prepare Your Environment for IBM FileNet P8*.

Related tasks

"Encrypting passwords for Process Engine"
Several passwords are required to configure Process Engine components. To accommodate your security requirements, you can encrypt these passwords before you enter them into the configuration files.

Encrypting passwords for Process Engine:

Several passwords are required to configure Process Engine components. To accommodate your security requirements, you can encrypt these passwords before you enter them into the configuration files.

Ensure that Java is installed and that the location to your Java runtime environment (JRE) is in your PATH environment variable.

1. Copy the following files from the `\Tools` directory on the installation media for Process Engine software to a local drive:
   - `fnencryptutils` - an executable .jar file
   - `RunEncryptApp` - a batch file

2. Run one of the following executable files to start the password application:

Option	Description
UNIX	RunEncryptApp.sh

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.

Related tasks

"Upgrading the Process Engine software silently (UNIX)" on page 145
Take the following steps to silently upgrade Process Engine.

Setting the f_maint and f_sw passwords

Set or reset the passwords for the f_sw and f_maint users,. This will verify that the encrypted version matches the database or operating system version of the passwords.

"Setting the f_maint and f_sw passwords (Oracle and Microsoft SQL Server)"
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 Microsoft SQL Server database. The encrypted password and the database user's passwords must match.

Setting the f_maint and f_sw passwords (Oracle and Microsoft SQL Server):

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 Microsoft SQL Server database. The encrypted password and the database user's passwords must match.

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

You must change the passwords for the f_maint and f_sw users. For Oracle and Microsoft SQL Server databases, both the encrypted file and the database passwords are updated.

On UNIX, several Process Engine utilities are configured with setuid *fnsw.* (or its alias), and will therefore run with that user's X Window settings, regardless of which user is currently logged on.

1. Log on as the fnsw user.
2. Start the Database Server Connect application by running the following command: `Xdbconnect -r`
3. Change the primary password for the users f_sw and f_maint to match the database password.
4. Exit the application.

Restoring 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

Editing 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.

Table 25. Changes to make to the /etc/inittab file

Operating system:	Changes to make:
AIX	Change:
	`rcfnsw:2:once:/etc/rc.initfnsw 2>&1 |` `alog -tboot > /dev/console 2>&1`
	to read:
	`:rcfnsw:2:once:/etc/rc.initfnsw 2>&1 |` `alog -tboot > /dev/console 2>&1`
HPUX	Change:
	`rcfn:2:once:/etc/rc.initfnsw 2>&1 |` `alog -tboot > /dev/console 2>&1`
	to read:
	`#rcfn:2:once:/etc/rc.initfnsw 2>&1 |` `alog -tboot > /dev/console 2>&1`
Solaris	Change:
	`fn:3:wait:/bin/sh /etc/rc.initfnsw` `</dev/console >/dev/console 2>&1`
	to read:
	`#fn:3:wait:/bin/sh /etc/rc.initfnsw` `</dev/console >/dev/console 2>&1`

Upgrading the Content Engine client files on Process Engine servers

Installing or upgrading the Content Engine Client files on all Process Engine
enables communication between the Content Engine and the Process Engine. For
upgrades from 4.0, you must first remove the Content Engine Client files before
upgrading them.

1. "Uninstalling Content Engine client 4.0.x instances"
   If you have the 4.0 version of the Content Engine client files installed, you must
   remove them before you run the upgrade installation to version 4.5.1. You must
   remove the 4.0 Content Engine client file for all Process Engine server
   machines.

2. "Upgrading the Content Engine client files" on page 149
   Installing or upgrading the Content Engine client files on all Process Engine
   enables communication between the Content Engine and the Process Engine.
   You can install these files interactively by using the installation wizard, or
   silently by using the command line and a silent input file.

Uninstalling Content Engine client 4.0.x instances

If you have the 4.0 version of the Content Engine client files installed, you must
remove them before you run the upgrade installation to version 4.5.1. You must
remove the 4.0 Content Engine client file for all Process Engine server machines.

"Uninstalling Content Engine Client 4.0 instances on UNIX"
You must use the uninstaller tools to remove the 4.0 version of the Content
Engine client files from your UNIX Process Engine server machines.

Uninstalling Content Engine Client 4.0 instances on UNIX

You must use the uninstaller tools to remove the 4.0 version of the Content Engine
client files from your UNIX Process Engine server machines.

Depending on the Content Engine client fix pack versions that you have installed, you might have more than one instance of the Content Engine client 4.0 installed. Additionally, the version 4.0 Content Engine client installer supported installing multiple instance of the Content Engine client on the same server. You must uninstall all instances before upgrading to the 4.5.1 version of Content Engine client.

1. Locate and uninstall all instances of the Content Engine client on the Process Engine server.
2. Run one of the following commands, depending on the fix pack version of Content Engine client that is installed.

Option	Description
Prior to fix pack P8CE-4.0.1-006	`/ce_client_install_path/ClientUpdater/` `_uninst/uninstaller.bin`
As of fix pack P8CE-4.0.1-006	`/ce_client_install_path/CEClient/` `_CEClientuninst/uninstaller`

3. In the Uninstall IBM FileNet Client Installer screen, click **Next** .
4. In the Select Components screen, select **Complete Uninstall**. Click **Next**.
5. In the Uninstall Complete screen, click **Done**.
6. Repeat the uninstallation steps for each installed instance of Content Engine client on this server.
7.

Upgrading the Content Engine client files

Installing or upgrading the Content Engine client files on all Process Engine enables communication between the Content Engine and the Process Engine. You can install these files interactively by using the installation wizard, or silently by using the command line and a silent input file.

"Upgrading the latest Content Engine client files on Process Engine servers interactively"
The installation wizard provides an interactive way to install the Content Engine client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

"Upgrading the latest Content Engine client files on Process Engine servers silently" on page 150
The command line method provides a way to silently install the Content Engine client files. You can use the values in your installation worksheet to edit the silent input text file before you run the installation.

Related concepts

⮕ Installation and Upgrade Worksheet
For more information about the Content Engine parameter values, see *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading the latest Content Engine client files on Process Engine servers interactively

The installation wizard provides an interactive way to install the Content Engine client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install the Content Engine client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data ⇒ Filter ⇒ AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. 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

   The default password for fnsw was BPMtemp1pzwd.

4. Verify that there is a current backup of Process Engine.

5. 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.

6. Access the IBM FileNet P8 Content Engine Client update software.

7. Run one of the commands in the table below, *CE_version* is the version of Content Engine you intend to run, for example, 4.5.1.

Operating System	Install Program
AIX	`P8CE-CLIENT-CE_version -AIX.BIN`

8. Complete the installation program wizard by using the values from your worksheet.

9. When the installation completes, check the Content Engine Client log file for errors. The path to the log file depends on the type of operating system on the machine where you installed the Content Engine Client files:

Operating System	Path to Log File
UNIX	`/opt/IBM/FileNet/CEClient/ceclient_install_log_4_5_1.txt`

Related concepts

⬐ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading the latest Content Engine client files on Process Engine servers silently

The command line method provides a way to silently install the Content Engine client files. You can use the values in your installation worksheet to edit the silent input text file before you run the installation.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install the Content Engine client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Client Installer** in the **Installation or Configuration Program** column.

2. To download the latest software updates, and to determine which of these updates might be required for use with other components and expansion products, contact your service representative.

3. 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

   The default password for fnsw was BPMtemp1pzwd.

4. Verify that there is a current backup of Process Engine.

5. 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.

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

7. 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.

8. 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:

9. 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.

   CE_version is the version of Content Engine you intend to run, for example, 4.5.1.

   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`

10. When the installation completes, check the Content Engine Client log file for errors. The path to the log file depends on the type of operating system on the machine where you installed the Content Engine Client files:

Operating System	Path to Log File
UNIX	`/opt/IBM/FileNet/CEClient/` `ceclient_install_log_4_5_1.txt`

11. When the installation completes, check the Content Engine Client log file for errors. The path to the log file depends on the type of operating system on the machine where you installed the Content Engine Client files:

Operating System	Path to Log File
UNIX	/opt/IBM/FileNet/CEClient/ ceclient_install_log_4_5_1.txt

Related concepts

⤷ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Completing post-upgrade Process Engine configuration

You must perform several procedures to complete the upgrade of all Process
Engine data and objects. The steps vary depending on what release you are
upgrading from.

"Completing the upgrade from Process Engine 3.5"
Complete the following tasks when upgrading Process Engine from Version 3.5.

"Completing the upgrade from Process Engine 4.0, 4.5.0, or higher" on page 157
Perform the following steps to update a Process Engine database to version
4.5.1. A database update will take place automatically by restarting the
software.

Completing the upgrade from Process Engine 3.5

Complete the following tasks when upgrading Process Engine from Version 3.5.

1. "Completing the upgrade from Process Engine 3.5 by database type"
    Complete the following tasks when upgrading Process Engine from Version 3.5.

2. "Removing existing routers in Process Task Manager" on page 154
    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.

3. "Updating the Process Engine security" on page 154
    Update the Process Engine security by using the **vwtaskman** tool.

4. "Migrating routers and updating isolated regions" on page 156
    Convert all routers to connection points and assign passwords to any existing
    isolated regions.

5. "Updating email notification" on page 157
    Enable email notification in order to use it. Process Engine also supports
    localized email notification.

6. "Backing up the database and restarting software" on page 157
    Back up and restart the database software after upgrading.

Completing the upgrade from Process Engine 3.5 by database type

Complete the following tasks when upgrading Process Engine from Version 3.5.

"Completing the upgrade from Process Engine 3.5 (Oracle)"
Complete the upgrade by initiating an automatic schema change in the
database.

Completing the upgrade from Process Engine 3.5 (Oracle):

Complete the upgrade by initiating an automatic schema change in the database.

Important: Throughout this procedure there are multiple software restarts. Run all restarts as documented. Do not start or restart Process Task Manager or other IBM FileNet software unless specifically instructed to do so in the procedures. When you restart Process Task Manager it accesses the database, which must not be done until the database has been upgraded.

To run steps on a UNIX operating system, the terminal must support X Window Systems and the DISPLAY environment variable must be set.

To update the Process Engine database objects:

1. After you update the Process Engine software, log onto the Process Engine server.

Option	Description
UNIX	Log on as fnsw.

2. Restart the Process Engine software by entering the following command:
   `initfnsw -y restart`

3. 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. Turn on tracing when prompted as to whether you want to initiate tracing to capture the changes made to a file. Accept the defaults to trace database access and transfer.

      Important: 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 and updates the schema accordingly.

      Important: If the upgrade fails at this point, you must restore the Process Engine database backup.

      See the FileNet Image Services error log to monitor the progress of the updates and ensure that no errors occur. Check the logs to verify that messages like the following examples are captured:

      ```
      2009/10/17 16:23:43.261 <fnsw> VW/Process (14952) ...
      [INFO] VW: Database upgrade successful to version 46,
      please follow instructions to perform the next step.

      2009/10/17 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.
```

4. Exit vwtool when you get a message that the procedure is complete.

5. Back up the Process Engine database. This backup can serve as a checkpoint if an error occurs later that requires you to restore the database.

6. Restart the Process Engine software. At a Windows command prompt or UNIX command line, enter the following command: `initfnsw -y restart`.

Removing existing routers in Process Task Manager

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.

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.

   You can do this immediately or wait until the corresponding connection points are created. On a Process Engine that is installed on the same server with an Application Engine, the routers cannot be removed until after the Application Engine and Process Engine client files have been upgraded. After 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 this procedure until you have deleted all Process Routers.

Updating the Process Engine security

Update the Process Engine security by using the **vwtaskman** tool.

Before you update the Process Engine security, ensure the following is true:
- Your directory server is running and correctly configured.
- Content Engine 4.5 is running.
- You have the CE_Operations custom settings saved in "Updating all isolated regions" on page 133.

When the Process Task Manager starts, a message is displayed that indicates that routers must be migrated. Routers will be migrated in a subsequent task. 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:

      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*. The service user name should be entered as a short name, not a distinguished name. See the IBM FileNet P8 help topic **System Administration → Enterprise-wide Administration → Process Task Manager → Process Engine → Configure the Process Engine → Security** for details on the user and groups.

d. Click **Apply.** If you get an error applying security settings, click **Close** on the message, correct the problem if noted, and start this procedure again.

Additional information is available in:

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 format to the 4.5 format.

a. Enable tracing when prompted.

b. 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. Monitor that log by executing:

`vl -t`

c. Address any errors that occurred run **vwtool** again. You might be prompted to ignore issues related to the user environment record upgrades. If all users did not successfully migrate to the 4.5 format, look at the messages in the error log and resolve the problems. Examples of the types of resolution required could include fixing a problem with the configuration of the Content Engine and its application server's access to the directory server or creating users in the underlying directory server itself.

Important: 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. If 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.

If **vwtool** fails with a shared memory error, configure contiguous free memory for Process Engine and try again.

3. Check the error log to verify that the database version number has been updated to 55. 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`

5. This restart of the software starts an automatic transfer of the upgrade.cdl file for all isolated regions. Monitor the syslogs for progress by executing:

```
vl -t
```

   If messages in the log indicate that any isolated region failed for phase 2 of the upgrade, the isolated region will be locked and the upgrade will proceed for other isolated regions. Make note of any region numbers where this phase of the upgrade failed.

6. Run the following substeps for any isolated region where the transfer failed.

   a. Start vwtool and enter the following at the prompt:

   ```
   reg XXX
   ```

   where:

   XXX indicates a region where the automatic transfer failed

   b. Enter at the vwtool prompt:

   ```
   unlockreg
   ```

   c. Repeat the unlock for every region that failed.

   d. Change directories to the following location of the *upgrade.cdl* file, depending on your operating system:

   UNIX `/fnsw/local/sd`

   e. Initiate a transfer on isolated regions that failed the automatic upgrade by entering the following command:

   ```
   vwtfer -o upgrade.cdl -r XXX
   ```

   where:

   -r XXX indicates a region where the automatic transfer failed

   At the prompt, log on as *pe_service_user*.

   f. Repeat the transfer on every region that failed.

   g. Restore the custom settings for CE_Operations saved before the upgrade in "Updating all isolated regions" on page 133.

Migrating routers and updating isolated regions

Convert all routers to connection points and assign passwords to any existing isolated regions.

1. "Configuring new regions in Process Task Manager"
   Configure new regions by using the Process Task Manager.

Configuring new regions in Process Task Manager:

Configure new regions by using the Process Task Manager.

1. Start Process Task Manager on the Process Engine server as follows, depending on your operating system:

Option	Description
UNIX	Enter the followng command on the command line: `vwtaskman`

   (UNIX) 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. The password you enter here must match the password entered for the region as it is defined in the Enterprise Manager.

   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:

Option	Description
UNIX	`/fnsw/local/logs/TM_daemon/` `PEDirectoryServerConnectionDebug.txt`

Updating email notification

Enable email notification in order to use it. Process Engine also supports localized email notification.

1. Add a language pack for the Default Authoring Locale.

   See the IBM FileNet P8 help topic *System 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.

2. Verify that the Default Authoring Locale is correct (it has defaulted to the operating system's locale).

3. Process Engine supports localized email notification. For details on configuring this, see the *IBM FileNet P8 Non-English Support Guide*.

Related information

⇨ Product documentation for IBM FileNet P8 Platform
Download the IBM FIleNet P8 Non-English Support Guide, as well as PDF's of the rest of the IBM FileNet P8 Platform documentation.

Backing up the database and restarting software

Back up and restart the database software after upgrading.

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. Restart the Process Engine software, as follows:

   At a Windows command prompt, or UNIX command line, type the following command:

   ```
   initfnsw -y restart
   ```

Completing the upgrade from Process Engine 4.0, 4.5.0, or higher

Perform the following steps to update a Process Engine database to version 4.5.1. A database update will take place automatically by restarting the software.

Before starting this procedure, ensure you have the CE_Operations custom settings saved in "Updating all isolated regions" on page 137. To update the database to version 4.5.1:

1. Run the following command to restart the Process Engine:

   ```
   initfnsw -y restart
   ```

2. Monitor the syslogs for progress by executing:

   ```
   vl -t
   ```

3. Verify that messages similar to the following are captured:

```
2009/07/27 15:47:14.632 <fnsw> VW/Process (3564.2552.12 0xdec.9f8)
 ... [INFO]VW (vwcemp):
Database upgrade successful to version 55
2009/07/27 15:47:14.679 <fnsw> VW/Process (3564.2552.12 0xdec.9f8)
 ... [INFO]VW (vwcemp):
Create report 'C:\FNSW_LOC\logs\peconfig\dbUpgradeRpt_55_20090727'
2009/07/27 15:47:17.240 <fnsw> VW/Process (3564.2552.12 0xdec.9f8)
 ... [INFO]VW (vwcemp):
Completed report 'C:\FNSW_LOC\logs\peconfig\dbUpgradeRpt_55_20090727'
2009/07/27 15:47:17.256 <fnsw> VW/Process (3564.2552.12 0xdec.9f8)
 ...VW (vwcemp):
Completed phase 1 upgrade
```

4. Continue to monitor the syslogs for progress by executing:

```
vl -t
```

The software will automatically be restarted again to complete phase 2 of the upgrade process. The automatic restart of the software occurs in the background and will not be reflected in the command window. If the software cannot be restarted, messages will be logged. Manually restart the software if the automatic restart fails.

5. This restart of the software starts an automatic transfer of the upgrade.cdl file for all isolated regions. Monitor the syslogs for progress by executing:

```
vl -t
```

6. Verify that messages similar to the following are captured:

```
2009/07/27 15:48:30.488 <fnsw> VW/Process (2040.3388.10 0x7f8.d3c)
 ... [INFO]VW (vwcemp):
PE upgrade phase 2 completed.
2009/07/27 15:48:30.503 <fnsw> VW/Process (2040.3388.10 0x7f8.d3c)
... [INFO]VW (vwcemp):
Enable PE daemon processes
2009/07/27 15:48:30.519 <fnsw> VW/Process (2040.3388.10 0x7f8.d3c)
... [INFO]VW (vwcemp):
Start PE daemon processes ...
```

If messages in the log indicate that any isolated region failed for phase 2 of the upgrade, the isolated region will be locked and the upgrade will proceed for other isolated regions. Make note of any region numbers where this phase of the upgrade failed.

7. Complete the following substeps for any isolated region where the transfer failed.

   a. Start vwtool and enter the following at the prompt:

   ```
   reg XXX
   ```

   where:

   XXX indicates a region where the automatic transfer failed

   b. Enter at the vwtool prompt:

   ```
   unlockreg
   ```

   c. Repeat the unlock for every region that failed.

   d. Change directories to the following location of the *upgrade.cdl* file, depending on your operating system:

   UNIX /fnsw/local/sd

   e. Initiate a transfer on isolated regions that failed the automatic upgrade by entering the following command:

   ```
   vwtfer -o upgrade.cdl -r XXX
   ```

where:

-r XXX indicates a region where the automatic transfer failed

At the prompt, log on as *pe_service_user*.

   f. Repeat the transfer on every region that failed.

   g. Restore the custom settings for CE_Operations saved before the upgrade in "Updating all isolated regions" on page 137.

1. "Backing up the database and restarting software"
   Back up and restart the database software after upgrading.

Backing up the database and restarting software

Back up and restart the database software after upgrading.

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. Restart the Process Engine software, as follows:

   At a Windows command prompt, or UNIX command line, type the following command:

   ```
   initfnsw -y restart
   ```

Installing Process Engine software updates

After installing Process Engine, software updates such as fix packs or interim fixes might also be required.

If you installed an instance into a high availability environment, perform the procedure on every node.

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 readme files for the Process Engine software updates and perform the installation procedures in the readme files on the Process Engine.

Installing 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 FileNet P8 components, for example, on Content Engine and Application Engine servers.

In a standard upgrade of your FileNet P8 system, you would install these client files in the course of upgrading each respective component.

However, if you are staging your FileNet P8 upgrade over a period of time and not upgrading one or more of the other FileNet P8 components at this time, you must still perform the Process Engine Client install on the other component machines.

- You will have this same client-installation requirement for any expansion products that use Process Engine Client files, such as:
  - IBM FileNet Workplace XT
  - IBM InfoSphere Enterprise Records
  - IBM FileNet Business Process Framework

- You must redeploy any components that are web-application-server based, such as Content Engine, Application Engine, Workplace XT, and FileNet Records Manager after you upgrade the Process Engine Client files on the associated machines.

Upgrading and configuring Application Engine

You can upgrade Application Engine from version 3.5 or version 4.0 to version 4.0.2. The procedures for upgrading Application Engine vary depending on the version that you are upgrading.

Important: Before you start to upgrade the Application Engine software, ensure that you are at the correct version and fix pack level, as indicated in *Plan and Prepare Your Environment for IBM FileNet P8*.

1. "Preparing for Application Engine upgrade" on page 162
   Before you upgrade your Application Engine installation, you must prepare your existing configuration for the upgrade. This preparation includes backing up existing files and undeploying your web application.

2. "Upgrading Application Engine" on page 164
   After you have prepared your existing configuration, you can begin the Application Engine upgrade. The tasks for upgrading Application Engine differ depending on whether you start from version 3.5 or version 4.0.

3. "Installing Application Engine software updates" on page 169
   If any required fix packs and interim fixes exist for Application Engine, you must install these software updates.

4. "Installing the latest Content Engine client files on Application Engine servers" on page 169
   The Content Engine Client software enables communication between the Application Engine and the Content Engine. Install the latest release or fix pack version of the Content Engine Client files on all Application Engine servers.

5. "Installing the latest Process Engine client files on Application Engine servers" on page 173
   The Process Engine Client software enables communication between the Application Engine and the Process Engine. You must install the Process Engine Client files on all Application Engine servers.

6. "Configuring Application Engine upgrades from 3.5" on page 175
   After you complete the Application Engine upgrade from 3.5, you must configure Application Engine files and your application sever to enable communication between Application Engine and Content Engine. You can choose the appropriate configuration tasks for your application server type and environment.

7. "Manually copying custom data" on page 188
   If your custom data was not retained in the upgrade, you must manually copy the data from the version 3.5 or 4.0.1 backup files (`deploy_backup`) to the upgraded directories.

8. "Deploying upgraded instances of Application Engine" on page 188
   After you upgrade and configure Application Engine, deploy your client application, Workplace, on your application server. It might be necessary to re-create the WAR or WAR and EAR files before you deploy.

9. "Completing post-upgrade Application Engine configuration" on page 192
   After you upgrade and deploy your Application Engine instances, you can complete additional configuration. Actions include modifying a configuration file, verifying Workplace sign-in, and establishing the Process Engine connection point.

Related information

➡ General requirements for all FileNet P8 Platform systems
Check to make sure your software is at the correct fix pack level before you
upgrade.

Preparing for Application Engine upgrade

Before you upgrade your Application Engine installation, you must prepare your
existing configuration for the upgrade. This preparation includes backing up
existing files and undeploying your web application.

- (Upgrades from 3.5 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*.
- Review the steps needed to retain the Application Engine 3.5 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 or
    4.0.1 configuration files.
  – During the upgrade, a number of existing configuration files will be moved to
    the newly created `AE_install_path`/Config/AE directory.
  – As part of the upgrade, the installation program automatically creates a
    backup of your existing Application Engine 3.5 or 4.0.1 configuration files,
    appending the suffix .old to the file names.
  – The existing version 3.5 or 4.0.1 `Actions.xml` and `web.xml` files will be merged
    with the 4.0.2 versions during the upgrade.
  – 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 and
Upgrade Guide*.

To prepare for the Application Engine upgrade:

1. Log on to the Application Engine server:

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

2. (Upgrades from 3.5 only) Verify that the Process Router is stopped.

   Tip: 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`

   b. Stop the router, if running.

      This is the router configured and started as part of the FileNet P8 Platform 3.5 installation. For more information, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide* task "Start the Process Router."

   c. Exit the Process Task Manager.

3. Back up, 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.

   Even though the installer automatically creates a backup of your existing Application Engine 3.5 or 4.0.1 configuration files, it is a best practice to back up the following directories to get a complete backup of your deployed system. You might need these files in later steps to retain your 3.5 or 4.0.1 settings. For example, when you copy modified files to the installed Workplace directory and perform the upgrade and configuration of Application Engine,

Option	Description
WebSphere Application Server	1. 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 Your path might vary if you created a new profile or renamed your WAR file. 2. Stop the Workplace application from the WebSphere administrative console. 3. Uninstall the Workplace application from **Enterprise Applications**. 4. Save the changes and stop the WebSphere Application Server. 5. Delete the temp Workplace directory (default: *app_engine_war*) from: *WAS_HOME*/profiles/default/temp/ *node_name*/*instance*/app_engine.war

4. Copy modified files to the installed Workplace directory. During the upgrade, the configuration files will be moved from the following directory:

   AE_install_path/Workplace/WEB-INF to

   AE_install_path/Config/AE

   If you have made modifications to any of these files directly in the deployed Workplace directory (app_engine.war for WebSphere Application Server or Workplace.war for JBoss) in your 3.5. or 4.0.1 environment, you must copy the 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.

   - `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

As part of the upgrade the installation program automatically creates a backup of your existing Application Engine 3.5 or 4.0.1 configuration files, appending the suffix .old to the file names.

- 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

Related tasks

"Manually copying custom data" on page 188
If your custom data was not retained in the upgrade, you must manually copy the data from the version 3.5 or 4.0.1 backup files (`deploy_backup`) to the upgraded directories.

Upgrading Application Engine

After you have prepared your existing configuration, you can begin the Application Engine upgrade. The tasks for upgrading Application Engine differ depending on whether you start from version 3.5 or version 4.0.

"Upgrading Application Engine from 3.5" on page 165
An upgrade from version 3.5 resembles a new install. You can choose to upgrade Application Engine interactively by using the upgrade wizard, or silently by using a silent input file.

"Upgrading Application Engine from 4.0" on page 167
In an upgrade from version 4.0 or later, the upgrade tool verifies that the currently installed version of Application Engine can be upgraded. You can

choose to upgrade Application Engine interactively by using the upgrade
wizard, or silently by using a silent input file.

Upgrading Application Engine from 3.5

An upgrade from version 3.5 resembles a new install. You can choose to upgrade
Application Engine interactively by using the upgrade wizard, or silently by using
a silent input file.

- If you plan to install and use the IBM FileNet Workplace XT product, you do
  not need to install Application Engine.
- To ensure proper functionality and performance, only install one instance of
  Application Engine per application server (or virtual machine or WebSphere
  Application Server LPAR). You can, however, deploy multiple instances of a
  single Application Engine version per application server.
- 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.
- You can install a new Application Engine 4.0.2 server within an existing 4.0 IBM
  FileNet P8 environment if your Content Engine is at version 401.006 or later and
  your Process Engine is at 403-000.001 or later. The Content Engine and Process
  Engine client files you install on the Application Engine must match the version
  of your Content Engine and Process Engine.

  "Upgrading Application Engine interactively"
  You can use the installation wizard to upgrade Application Engine. To complete
  the upgrade, provide the necessary parameter values to enable the Application
  Engine client application, Workplace to access the information that is managed
  by Content Engine.

  "Upgrading Application Engine silently" on page 166
  You can include the parameters for your Application Engine upgrade in a silent
  input file, and then run the upgrade from a command line.

Upgrading Application Engine interactively

You can use the installation wizard to upgrade Application Engine. To complete
the upgrade, provide the necessary parameter values to enable the Application
Engine client application, Workplace to access the information that is managed by
Content Engine.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

To install the Application Engine software interactively:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command
   is enabled. To view only Content Engine values, filter by **AE Installer** in the
   Installation or Configuration Program column.

2. Log on to the application server.

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

3. Access the IBM FileNet P8 Application Engine 4.0.2.0 installation software.
4. Start the installation program.

Platform	Command
AIX	P8AE-4.0.2.0-AIX.bin

5. Complete the Application Engine installer screens by using the values from your worksheet.

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

   Verify that no errors or failures were logged. Look for the ERROR designation at the start of a line. Correct any errors before you proceed.

Related concepts

⇨ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Application Engine silently

You can include the parameters for your Application Engine upgrade in a silent input file, and then run the upgrade from a command line.

To upgrade the Application Engine software silently:

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

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

2. Locate the IBM FileNet P8 Application Engine 4.0.2 installation software package, and copy the appropriate AE_silent_input.txt file to a local directory.

3. Follow the instructions in the silent input file to edit the file to record the appropriate responses for your installation.

4. Modify all instances of *AE_install_path* in the script as follows.

   UNIX

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

5. From a command prompt, navigate to and run the installer.

   UNIX

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

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

   Verify that no errors or failures were logged. Look for the ERROR designation at the start of a line. Correct any errors before you proceed.

Upgrading Application Engine from 4.0

In an upgrade from version 4.0 or later, the upgrade tool verifies that the currently
installed version of Application Engine can be upgraded. You can choose to
upgrade Application Engine interactively by using the upgrade wizard, or silently
by using a silent input file.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

- If you plan to install and use the IBM FileNet Workplace XT product, installing
  Application Engine is not required.
- 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.
- 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.

  "Upgrading Application Engine from 4.0 interactively"
  You can use the installation wizard to upgrade Application Engine. To complete
  the upgrade, provide the necessary parameter values to enable the Application
  Engine client application, Workplace, to access the information that is managed
  by Content Engine.

  "Upgrading Application Engine from 4.0 silently" on page 168
  You can include the parameters for your Application Engine upgrade in a silent
  input file, and then run the upgrade from a command line.

Related concepts

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Application Engine from 4.0 interactively

You can use the installation wizard to upgrade Application Engine. To complete
the upgrade, provide the necessary parameter values to enable the Application
Engine client application, Workplace, to access the information that is managed by
Content Engine.

To upgrade Application Engine:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** › **Filter** › **AutoFilter** command
   is enabled. To view only Application Engine values, filter by **AE Installer** in the
   Installation or Configuration Program column.

2. Log on to the server where Application Engine is installed:

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

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

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

3. Access the IBM FileNet Application Engine 4.0.2 installation software.
4. Start the installation program.

Platform	Command
AIX	P8AE-4.0.2.0-AIX.bin

5. Complete the Application Engine installation program screens by using the values in your worksheet.

   In addition to the prompts for system-specific information, 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.

6. 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.

Related concepts

⇨ Installation and upgrade worksheet

For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading Application Engine from 4.0 silently

You can include the parameters for your Application Engine upgrade in a silent input file, and then run the upgrade from a command line.

To install the Application Engine software silently:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ⇨ **Filter** ⇨ **AutoFilter** command is enabled. To view only Application Engine values, filter by **AE Installer** in the **Installation or Configuration Program** column.

2. Log on to the server where Application Engine is installed:

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

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

3. Edit the silent input file to reflect the appropriate responses for your upgrade by following the instructions in the file and the values in your worksheet.
4. From a command prompt, navigate to the installer and run it.

UNIX
```
./P8AE-4.0.2.0-operating system.bin -options
path_to_edited_input_file/AE_silent_input.txt -silent
```
5. View the `app_engine_install_log_4_0_2.txt` file, located in
   `AE_install_path`/AE/Logs.

   Verify that no errors or failures were logged. Look for the ERROR designation
   at the start of a line. Correct any errors before you proceed.

Related concepts

▶ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing Application Engine software updates

If any required fix packs and interim fixes exist for Application Engine, you must
install these software updates.

To install the Application Engine software updates:
1. To download the latest software updates, and to determine which of these
   updates might 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 Application Engine
   Client files using the subsequent tasks.

Installing the latest Content Engine client files on Application Engine servers

The Content Engine Client software enables communication between the
Application Engine and the Content Engine. Install the latest release or fix pack
version of the Content Engine Client files on all Application Engine servers.

Important: If you are performing a staged upgrade, and you are upgrading
Application Engine before you upgrade Content Engine, it is not necessary to
complete this task at this time. When you upgrade Content Engine to version 4.5.1,
you must then use this task to run the version 4.5.1 Content Engine Client installer
on your Application Engine, and redeploy the Workplace application.

Depending on the Content Engine Client versions that you have installed, you
might have more than one instance of the Content Engine Client version 4.0
installed. Additionally, the version 4.0 Content Engine Client installer supported
installing multiple instance of the Content Engine Client on the same server.
Uninstall all instances before upgrading to Content Engine Client version 4.5.1.

Before you upgrade the Content Engine Client files to version 4.5.1, you must
remove the version 4.0 Client files.

The installation wizard provides an interactive way to install the Content
Engine Client files. You can use the values you collected on your worksheet to
fill in the required value for each field on the wizard screens.

"Installing the Content Engine Client files silently" on page 172
The command line method provides a way to silently install the Content
Engine Client files. You can use the values in your installation worksheet to edit
the silent input text file before you run the installation.

Removing the Content Engine 4.0 Client files from Application Engine servers

Before you upgrade the Content Engine Client files to version 4.5.1, you must
remove the version 4.0 Client files.

If you are performing a staged upgrade, and are upgrading Application Engine
before you upgrade Content Engine, it is not necessary 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.

Perform the following procedure to identify and uninstall all Content Engine Client
version 4.0 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 might
have more than one instance of the CE Client 4.0 installed. Additionally, the
version 4.0 CE Client installer supported installing multiple instance of the CE
Client on the same server. You must uninstall all instances before upgrading to the
4.5 version of CE Client.

"Removing Content Engine 4.0 Client files from Application Engine servers
(UNIX)"
The Content Engine 4.0 Client installation includes a UNIX tool for uninstalling
the Client files.

Removing Content Engine 4.0 Client files from Application Engine servers (UNIX)

The Content Engine 4.0 Client installation includes a UNIX tool for uninstalling the
Client files.

To uninstall Content Engine Client 4.0. instances:

1. 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`.

2. In the Uninstall IBM FileNet Client Installer screen, click **Next**.
3. In the Select Components screen, select Complete Uninstall. Click **Next**.
4. In the Uninstall Complete screen, click **Done**.
5. Repeat these steps for each additional instance of Content Engine Client
   updater or installer programs on this server.

Installing the latest Content Engine Client files on Application Engine servers interactively

The installation wizard provides an interactive way to install the Content Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To install the Content Engine Client files interactively:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** → **Filter** → **AutoFilter** command is enabled. To view only Content Engine values, filter by **CE Client Installer** in the **Installation or Configuration Program** column.

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 installation software from the Content Engine installation software to the temporary directory. The version of the client installation software must match the version of Content Engine.

5. Expand the compressed Content Engine Client installation software file in the temporary directory.

6. Close all instances of Workplace and any other Application Engine client applications. From the application server administrative console, stop and undeploy Application Engine.

Option	Description
WebSphere Application Server	Stop and undeploy the Application Engine Web application.

7. Access the IBM FileNet Content Engine client update software in the temporary directory.

8. Start the installation program by running the following command, where *CE_version* is the version of Content Engine you intend to run, for example, 4.5.1.

Operating System	Install Program
AIX	`P8CE-CLIENT-CE_version-AIX.BIN`

9. Complete the installation screens by using the values from your installation worksheet.

10. When the installation completes, check the Content Engine Client log file for errors. The path to the log file depends on the type of operating system on the machine where you installed the Content Engine Client files:

Operating System	Path to Log File
UNIX	`/opt/IBM/FileNet/CEClient/` `ceclient_install_log_4_5_1.txt`

Related concepts

⇥ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the
worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing the Content Engine Client files silently

The command line method provides a way to silently install the Content Engine
Client files. You can use the values in your installation worksheet to edit the silent
input text file before you run the installation.

Be sure that you have available the Installation and Upgrade Worksheet that was
completed during your planning activities.

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command
   is enabled. To view only Content Engine values, filter by **CE Client Installer**
   in the **Installation or Configuration Program** column.

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 installation software from the Content Engine
   installation software to the temporary directory. The version of the client
   installation software must match the version of Content Engine.

5. Expand the compressed Content Engine Client installation software file in the
   temporary directory.

6. Close all instances of Workplace and any other Application Engine client
   applications. From the application server administrative console, stop and
   undeploy Application Engine.

Option	Description
WebSphere Application Server	Stop and undeploy the Application Engine Web application.

7. Make a backup copy of the `CEClient_silent_install.txt` input file in the
   temporary directory.

8. 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.

9. Navigate to the temporary directory path containing the Content Engine
   Client installation program.

10. Start the installation program by running the following command, where:

    CE_version is the version of Content Engine you intend to run, for example,
    4.5.1.

    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

11. When the installation completes, check the Content Engine Client log file for errors. The path to the log file depends on the type of operating system on the machine where you installed the Content Engine Client files:

Operating System	Path to Log File
UNIX	/opt/IBM/FileNet/CEClient/ ceclient_install_log_4_5_1.txt

Related concepts

⏏ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Installing the latest Process Engine client files on Application Engine servers

The Process Engine Client software enables communication between the Application Engine and the Process Engine. You must install the Process Engine Client files on all Application Engine servers.

"Upgrading the latest Process Engine Client files interactively"
The installation wizard provides an interactive way to upgrade the Process Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

"Upgrading the latest Process Engine Client files silently" on page 174
You can upgrade the Process Engine Client software on the Application Engine by using a silent input text file and running the installation program from the command line. Use the values in your upgrade worksheet to edit the silent input text file before you run the upgrade.

Upgrading the latest Process Engine Client files interactively

The installation wizard provides an interactive way to upgrade the Process Engine Client files. You can use the values you collected on your worksheet to fill in the required value for each field on the wizard screens.

To upgrade the Process Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data → Filter → AutoFilter** command is enabled. To view only Process Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.

2. 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.

3. 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.

4. Expand the compressed Process Engine Client install software file within the temporary directory. The jre.tar file that is packaged with the Process Engine Client installer must be in the same directory as the Process Engine Client installer when the installer is run.

5. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Application Engine is deployed.

   To run the program interactively, run one of the commands in the table below, and complete the screens.

Operating System	Install Program
AIX	P8PE-CLIENT-*PE_version*-AIX.BIN

6. For Process Engine upgrades from version 4.0 only, add required JAR files to the Component Manager library. 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:

Option	Description
UNIX	routercmd.sh

   a. In the Process Task Manager console, under Application Engine, **Component Manager** node, before starting up any existing router, click the **Required** libraries tab.

   b. Click **Add** and browse to /Router/lib.

   c. Select the following files to add:

      ContentExtendedOps.jar
      smtp.jar
      pe3pt.jar

Related concepts

➡ Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Upgrading the latest Process Engine Client files silently

You can upgrade the Process Engine Client software on the Application Engine by using a silent input text file and running the installation program from the command line. Use the values in your upgrade worksheet to edit the silent input text file before you run the upgrade.

To upgrade the Process Engine Client files:

1. Open your completed Installation and Upgrade Worksheet file.

   Tip: In the worksheet file, verify that the **Data** ⟶ **Filter** ⟶ **AutoFilter** command is enabled. To view only Process Engine values, filter by **PE Client Installer** in the **Installation or Configuration Program** column.

2. 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.

3. 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.

4. Expand the compressed Process Engine Client install software file within the temporary directory. The jre.tar file that is packaged with the Process Engine Client installer must be in the same directory as the Process Engine Client installer when the installer is run.

5. The expanded install software contains the Process Engine Client install program specific to the operating system on the machine where Application Engine is deployed.

6. In the expanded install software, open the file `PEClient_silent_install.txt` and change the `Variable_CheckboxAE` line to the following value:

   `-V Variable_CheckboxAE="true"`

7. Choose one of the following commands to run the silent installation:

Operating System	Install Program
AIX	`P8PE-CLIENT-PE_version-AIX.BIN -silent -options "PEClient_silent_install.txt"`

8. For Process Engine upgrades from version 4.0 only, add required JAR files to the Component Manager library. 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:

Option	Description
UNIX	`routercmd.sh`

   a. In the Process Task Manager console, under Application Engine, **Component Manager** node, before starting up any existing router, click the **Required** libraries tab.

   b. Click **Add** and browse to `/Router/lib`.

   c. Select the following files to add:

   ```
   ContentExtendedOps.jar
   smtp.jar
   pe3pt.jar
   ```

Related concepts

Installation and upgrade worksheet
For more information about the installation and upgrade worksheet, see the worksheet topics in *Plan and Prepare Your Environment for IBM FileNet P8*.

Configuring Application Engine upgrades from 3.5

After you complete the Application Engine upgrade from 3.5, you must configure Application Engine files and your application sever to enable communication between Application Engine and Content Engine. You can choose the appropriate configuration tasks for your application server type and environment.

"Configuring Application Engine upgrades from 3.5 on WebSphere Application Server" on page 176

After you upgrade the Application Engine server from 3.5, you must configure WebSphere Application Server to work with Application Engine. You can also configure changes for optional modes like SSO.

Configuring Application Engine upgrades from 3.5 on WebSphere Application Server

After you upgrade the Application Engine server from 3.5, you must configure WebSphere Application Server to work with Application Engine. You can also configure changes for optional modes like SSO.

"Editing web.xml for container-managed authentication"
You can choose to use WebSphere Application Server with container-managed authentication. To enable this optional authentication approach, edit the `web.xml` file on the WebSphere Application Server.

"Editing web.xml for SSO" on page 179
You can use SSO with a proxy server in your Application Engine environment. To enable this optional approach, edit the `web.xml` file on the WebSphere Application Server.

"Configuring Java Virtual Machine settings for JAAS login and memory" on page 183
You can use the Java Virtual Machine settings in the WebSphere Application Server to set up JAAS login information for authentication and memory settings for resource usage.

"Configuring Lightweight Third Party Authentication (LTPA)" on page 184
To set up LTPA security, configure settings to match on both the Content Engine application server and the Application Engine application server. If your Application Engine and Content Engine are on the same WebSphere Application Server, it is not necessary to configure LTPA.

"Configuring stand-alone Lightweight Directory Access Protocol (LDAP)" on page 185
To enable LDAP communication between Application Engine and Content Engine, you must configure settings on the WebSphere Application Server.

"Configuring Lightweight Directory Access Protocol (LDAP) for federated repositories" on page 187
If you have a multiple domain environment, configure LDAP settings for federated repositories on the WebSphere Application Server to enable LDAP communication between Application Engine and Content Engine.

Editing web.xml for container-managed authentication

You can choose to use WebSphere Application Server with container-managed authentication. To enable this optional authentication approach, edit the `web.xml` file on the WebSphere Application Server.

1. Make a backup copy of `web.xml`.

   `AE_install_path/Workplace/WEB-INF/web.xml`

2. Open `web.xml` for editing, search for the parameter `challengeProxyEnabled`, and set it to false.

   ```
   <param-name>challengeProxyEnabled</param-name>
   <param-value> false </param-value>
   ```

3. 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 thecontainer 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>
```

4. 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>
```

5. 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>
```

6. 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>
```

7. Search for the first instance of a `<security-constraint>` tag, and add the following `<security-constraint>` tag before that tag.

   Important: 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>
```

8. Save your changes to `web.xml` and close the file.

Editing web.xml for SSO

You can use SSO with a proxy server in your Application Engine environment. To enable this optional approach, edit the `web.xml` file on the WebSphere Application Server.

1. Make a backup copy of `web.xml`.

   AE_install_path/Workplace/WEB-INF/web.xml

2. Open `web.xml` for editing, search for the parameter `challengeProxyEnabled`, and set it to false.

```
<param-name>challengeProxyEnabled</param-name>
<param-value> false </param-value>
```

3. 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 thecontainer 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>
```

4. 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>
```

5. 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>
```

6. Search for the first instance of a <security-constraint> tag, and add the following <security-constraint> tag before that tag.

Important: 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>
```

7. 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>
-->
```

8. 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 deployment 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 values:

```
<param-name>ssoProxyContextPath</param-name>
<param-value>/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 values:

```
<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>
```

9. Save your changes to `web.xml` and close the file.

Configuring Java Virtual Machine settings for JAAS login and memory

You can use the Java Virtual Machine settings in the WebSphere Application Server to set up JAAS login information for authentication and memory settings for resource usage.

1. Log in to the WebSphere administrative console.
2. Navigate to the Java Virtual Machine settings.

Option	Description
WebSphere Application Server 6.1	**Servers → Application Servers →** *server_name* **→ Java & Process Management → Process Definition → Java Virtual Machine**
WebSphere Application Server 7.0	**Servers → Server Types → WebSphere application servers →** *server_name* **> Java & Process Management → Process Definition → Java Virtual Machine**

3. In the **Generic JVM arguments** field, add the following entry:

```
-Djava.security.auth.login.config=AE_install_path\CE_API\config\
jaas.conf.WebSphere
```

Replace *AE_install_path* in the entry above with your actual installation path, such as:

UNIX `/opt/FileNet/AE`

Your path might be slightly different based on the version of your client installations, or whether you have chosen a custom path for the installation. Verify the location of this file and specify the install path location before you enter the path.

Important: Do not copy and paste the text from this guide into the field in the console because hidden formatting can cause problems with the entry. Use a paste special option in a text editor to remove formatting first, or type the entry into the field.

4. Set the Initial heap size and Maximum heap size, and save your changes.

For example, you might set the Initial heap size to 512 and the Maximum heap size to 1024. However, these values vary significantly depending on your machine size.

Refer to your application server vendor recommendation for Initial heap size and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Performance Tuning Guide*.

5. (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 Application Server.

Related information

➦ Product documentation for IBM FileNet P8 Platform
Download the IBM FIleNet P8 Performance Tuning Guide, as well as PDF's of the rest of the IBM FileNet P8 Platform documentation.

Configuring Lightweight Third Party Authentication (LTPA)

To set up LTPA security, configure settings to match on both the Content Engine application server and the Application Engine application server. If your Application Engine and Content Engine are on the same WebSphere Application Server, it is not necessary to configure LTPA.

If you are already using LTPA with your Content Engine application server, export the existing keys and copy the key file to the Application Engine server. Check with your Content Engine administrator.

To configure LTPA, take the following steps:

1. On the Content Engine server, log in to the WebSphere administrative console.

2. Navigate to the LTPA settings page.

Option	Description
WebSphere Application Server 6.1	1. Navigate to **Security > Secure administration, applications, and infrastructure**. 2. From the right side of the panel, select **Authentication Mechanisms and expiration**.
WebSphere Application Server 7.0	1. Navigate to **Security > Global security**. 2. From the right side of the panel, select **LTPA**.

3. Enter a value for the **LTPA timeout** that is larger than the default. For example, if the timeout value is left at the default value of 120 minutes, the LTPA key expires after 2 hours. Users will not be able to log in to Workplace after being logged in for 2 hours.

4. Save your changes.

5. In the box for **Cross-cell single sign-on**, type a password to create the **LTPA password**.

For password restrictions, see the WebSphere Application Server documentation. If you have already configured Content Engine for LTPA, use the existing password in the Application Engine configuration below.

6. Enter the fully qualified path for the **Key File Name**. For example, /opt/LTPA/*ltpa_key_name*.

7. Click **Export keys**. Verify that a message similar to the following message is displayed: The keys were successfully exported to the file *ltpa_key_name*.

8. Click **OK**, then click **Save changes directly to the master configuration**.

9. Stop and restart WebSphere Application Server.

10. Copy the key file from the Content Engine server location you specified above to a directory on the Application Engine server. For example, /opt/LTPA/*ltpa_key_name*

11. On the Application Engine server, log in to the WebSphere administrative console.

12. Navigate to the LTPA settings page.

Option	Description
WebSphere Application Server 6.1	1. Navigate to **Security > Secure administration, applications, and infrastructure**. 2. From the right side of the panel, select **Authentication Mechanisms and expiration**.
WebSphere Application Server 7.0	1. Navigate to **Security > Global security**. 2. From the right side of the panel, select **LTPA**.

13. Enter a value for the **LTPA timeout** that is larger than the default. For example, if the timeout value is left at the default value of 120 minutes, the LTPA key expires after 2 hours. Users will not be able to log in to Workplace after being logged in for 2 hours.

14. Save your changes.

15. In the box for **Cross-cell single sign-on**, type and confirm the **LTPA password** you created for Content Engine.

    For password restrictions, see the WebSphere Application Server documentation. If you have already configured Content Engine for LTPA, use the existing password in the Application Engine configuration below.

16. Specify the path for the key file that you copied to the Application Engine server. For example, /opt/LTPA/*ltpa_key_name*.

17. Click **Import keys**. Verify that a message similar to the following one is displayed: The keys were successfully imported from the file *ltpa_key_name*.

18. Save your changes.

Configuring stand-alone Lightweight Directory Access Protocol (LDAP)

To enable LDAP communication between Application Engine and Content Engine, you must configure settings on the WebSphere Application Server.

To configure stand-alone Lightweight Directory Access Protocol (LDAP):

1. Open the WebSphere administrative console.

2. Navigate to **Security > Secure administration, applications, and infrastructure**.

3. Disable security using the following Security settings:
   - Clear the **Enable Administrative Security** flag.
   - Clear the **Enable application security** flag.
   - Clear the **Java 2 security** flag.

4. From the bottom of the panel, in the dropdown list called **Available realm definitions**, select **Standalone LDAP registry** and click **Configure**.

5. Configure the LDAP provider to exactly match the corresponding settings on the Content Engine application server.

   Tip: Open the WebSphere administrative console for Content Engine to the same panels to see and copy all settings.
   - Primary administrative user name
   - Select Automatically generated server identity.
   - Type
   - Host
   - Port
   - Base distinguished name (DN)
   - Bind distinguished name (DN)
   - Bind password

6. Configure the LDAP user registry settings to exactly match the corresponding settings on the Content Engine application server.

   Tip: Open the WebSphere administrative console for Content Engine to the same panel to see and copy all settings.
   - User filter
   - Group filter
   - User ID map
   - Group member ID map
   - Certificate map mode
   - Certificate filter

7. Save these settings.

8. Next to **Available realm definitions**, ensure that **Standalone LDAP registry** is still selected, and click **Set as current**.

9. Re-enable the following Security settings:
   - Select the **Enable Administrative Security** flag.
   - Select the **Enable application security** flag.
   - Clear the **Java 2 security** flag.

   The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.

10. Save your changes to the master configuration.

11. 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.

Configuring Lightweight Directory Access Protocol (LDAP) for federated repositories

If you have a multiple domain environment, configure LDAP settings for federated repositories on the WebSphere Application Server to enable LDAP communication between Application Engine and Content Engine.

Important: If you are using federated repositories, your WebSphere administrative console user cannot have the same username or ID as a user in the LDAP repository.

To configure Lightweight Directory Access Protocol (LDAP) for federated repositories:

1. Open the WebSphere administrative console.
2. Navigate to the security settings page.

Option	Description
WebSphere Application Server 6.1	Navigate to **Security > Secure administration, applications, and infrastructure**.
WebSphere Application Server 7.0	Navigate to **Security > Global security**.

3. Disable the following Security settings:
   - Clear the **Enable Administrative Security** flag.
   - Clear the **Enable application security** flag.
   - Clear the **Java 2 security** flag.
4. From the bottom of the panel, in the dropdown list called **Available realm definitions**, select **Federated Repositories** and click **Configure**.
5. Configure the LDAP provider to exactly match the corresponding **General Properties** on the Content Engine application server.

   Tip: Open the WebSphere administrative console for Content Engine to the same panels to see and copy all settings.
   - Realm name
   - Primary administrative user name
   - Select Automatically generated server identity.
   - Ignore case for authorization
   - Repositories in the realm
6. Save these settings.
7. Next to **Available realm definitions**, ensure that **Federated repositories** is still selected, and click **Set as current**.
8. Re-enable the following Security settings:
   - Select the **Enable Administrative Security** flag.
   - Select the **Enable application security** flag.
   - Clear the **Java 2 security** flag.

   The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.
9. Save your changes to the master configuration.
10. Restart the WebSphere instance.
11. Test the connection to the repository.

a. In the WebSphere administrative console, navigate to **Users and Groups** → **Manage Users**.
   b. Click **Search By User ID** , and enter a known user.
   c. Click **Search**. This should return the user from the configured LDAP repository.

Manually copying custom data

If your custom data was not retained in the upgrade, you must manually copy the data from the version 3.5 or 4.0.1 backup files (`deploy_backup`) to the upgraded directories.

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 "Preparing for Application Engine upgrade."

When you have completed the manual copy of your custom data, deploy the Workplace application.

Related tasks

"Preparing for Application Engine upgrade" on page 162
Before you upgrade your Application Engine installation, you must prepare your existing configuration for the upgrade. This preparation includes backing up existing files and undeploying your web application.

Deploying upgraded instances of Application Engine

After you upgrade and configure Application Engine, deploy your client application, Workplace, on your application server. It might be necessary to re-create the WAR or WAR and EAR files before you deploy.

"Deploying upgraded instances of Application Engine on WebSphere Application Server"
After you upgrade and configure Application Engine, deploy your client application, Workplace, on your WebSphere Application Server. It might be necessary to re-create the WAR or WAR and EAR files before you deploy.

Deploying upgraded instances of Application Engine on WebSphere Application Server

After you upgrade and configure Application Engine, deploy your client application, Workplace, on your WebSphere Application Server. It might be necessary to re-create the WAR or WAR and EAR files before you deploy.

"Re-creating the WAR or EAR file" on page 189
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 re-create the WAR or EAR file and redeploy your changes.

"Deploying Application Engine" on page 189
The method for deploying Application Engine varies depending on your WebSphere Application Server version.

Re-creating 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 re-create the WAR or EAR file and redeploy your changes.

Remember: Before re-creating the EAR file, you must first re-create the WAR file.

To re-create the WAR or EAR files:

1. Verify that all modified /Workplace directory files have been saved.
2. Recreate the app_engine.war file by running the script for your platform.

Option	Description
UNIX	*AE_install_path*/deploy/ create_app_engine_war.sh

3. (For WAR file deployment) Rename the newly recreated app_engine.war file to Workplace.war or *custom_name*.war to create the context root for your application. For 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.

   Important: You must rename the WAR file every time you regenerate it. The create_app_engine_war script will by default create a file with the name app_engine.war.

   Important: Do not rename the WAR file if you are using the WAR file to recreate the EAR file for deployment.

4. (For EAR file deployments only) Re-create the app_engine.ear file by running the script for your platform.

Option	Description
UNIX	*AE_install_path*/deploy/ create_app_engine_ear.sh

Deploying Application Engine

The method for deploying Application Engine varies depending on your WebSphere Application Server version.

"Deploying the application (WebSphere Application Server 6.1)"
To use Workplace as a web application, you must deploy it as an application in the WebSphere 6.1 administrative console. You can specify the context root for the application URL during the deployment.

"Deploying the application (WebSphere Application Server 7.0)" on page 191
To use Workplace as a web application, you must deploy it as an application in the WebSphere 7.0 administrative console. You can specify the context root for the application URL during the deployment.

Deploying the application (WebSphere Application Server 6.1):

To use Workplace as a web application, you must deploy it as an application in the WebSphere 6.1 administrative console. You can specify the context root for the application URL during the deployment.

To deploy Application Engine:

1. Log on to the WebSphere administrative console.
2. Navigate to **Applications > Enterprise Applications > Install**
3. Select the file to deploy.
   - (If the administrative console is running *locally*) Select **Local file system** and enter or browse to the location of the app_engine.war or app_engine.ear file created by the installation program. For example, *AE_install_path*/deploy/app_engine.war or *AE_install_path*/deploy/app_engine.ear. Do not enter the machine name.
   - (If the administrative console is *remote*) Select **Remote file system** and enter the fully-qualified pathname to the app_engine.war or app_engine.ear file. For example, *AE_install_path*/deploy/app_engine.war or *AE_install_path*/deploy/app_engine.ear. Do not enter the machine name.
4. If you are deploying a WAR file, enter the context root.

   Enter Workplace and click **Next** to proceed to deploying a new application.

   Tip: 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. Click **Prompt only when additional information is required**, and click **Next**. Complete the dialogs for installing a new application, using the following settings:

   For **Application name**, enter Workplace, or the name you chose to call the application.

   For **WebServer**, specify the server you are planning to use. Verify that your application name is selected and associated with the correct WebServer.

   For **Map virtual hosts for Web modules**, for virtual host, choose the default_host.

   (WAR file deployment only) For **Context root**, specify Workplace, or whatever you want to call the application.
6. Save your configuration.
7. Navigate to **Applications > Enterprise Applications > Workplace > Class loading and update detection**.
8. Set the polling interval for updated files with a number appropriate for your environment, for example, 3 seconds.
9. Change the **Class loader order** to **Classes loaded with application class loader first**.
10. Click **Apply**.
11. Click **Workplace > Manage Modules**.
12. Click **Workplace**.
13. Change the **Class loader order** to **Classes loaded with application class loader first**.

    Important: Do this only for the specific web application. Do not change the similar settings for the entire application server.
14. If you are using container-managed authentication, navigate to **Enterprise Applications > Workplace > Security role to user/group mapping**. Verify that the **All Authenticated** column is checked for the "All Authenticated" role.
15. Save all your changes to the master configuration.
16. Stop and restart the WebSphere Application Server instance where the application is installed.

To troubleshoot the deployment, check the following log:

WAS_install_path/AppServer/profiles/*profile_name*/logs/*server_name*/
SystemOut.log

17. Start Workplace (or whatever you named your application) from the administrative console.

Deploying the application (WebSphere Application Server 7.0):

To use Workplace as a web application, you must deploy it as an application in the WebSphere 7.0 administrative console. You can specify the context root for the application URL during the deployment.

To deploy Application Engine:

1. Log on to the WebSphere administrative console.
2. Navigate to **Applications > WebSphere Enterprise Applications > New Enterprise Applications > Install**
3. Select the file to deploy.
   - (If the administrative console is running *locally*) Select **Local file system** and enter or browse to the location of the app_engine.war or app_engine.ear file created by the installation program. For example, *AE_install_path*/deploy/app_engine.war or *AE_install_path*/deploy/app_engine.ear. Do not enter the machine name.
   - (If the administrative console is *remote*) Select **Remote file system** and enter the fully-qualified pathname to the app_engine.war or app_engine.ear file. For example, *AE_install_path*/deploy/app_engine.war or *AE_install_path*/deploy/app_engine.ear. Do not enter the machine name.
4. If you are deploying a WAR file, enter the context root.

   Enter Workplace and click **Next** to proceed to deploying a new application.

   Tip: The context root is the name of the application you log in to by using the web interface, such as: http://*ApplicationEngineServerName*:*port#*/*Context_Root.*
5. Click **Prompt only when additional information is required**, and click **Next**. Complete the dialog boxes for installing a new application by using the following settings:

   For **Application name**, enter Workplace, or the name you chose to call the application.

   For **WebServer**, specify the server you are planning to use. Verify that your application name is selected and associated with the correct WebServer.

   For **Map virtual hosts for Web modules**, for virtual host, choose the default_host.

   (WAR file deployment only) For **Context root**, specify Workplace, or whatever you want to call the application.
6. Save your configuration.
7. Navigate to **Applications > Enterprise Applications > Workplace > Class loading and update detection**.
8. Set the polling interval for updated files with a number appropriate for your environment, for example, 3 seconds.
9. Change the **Class loader order** to **Classes loaded with local class loader first (parent last)**.
10. Click **Apply**.

11. Click **Workplace > Manage Modules**.
12. In the **Modules** column, cllick **Workplace**.
13. Change the **Class loader order** to **Classes loaded with local class loader first (parent last)**.

    Important: Do this only for the specific web application. Do not change the similar settings for the entire application server.
14. If you are using container-managed authentication, navigate to **Enterprise Applications > Workplace > Security role to user/group mapping**.
    a. Select the **All Authenticated** role and click **Map Special Subjects**.
    b. Map the **All Authenticated** role and to **All Authenticated in Applications realm**.
15. Save all your changes to the master configuration.
16. Stop and restart the WebSphere Application Server instance where the application is installed.

    To troubleshoot the deployment, check the following log:

    `WAS_install_path`/AppServer/profiles/`profile_name`/logs/`server_name`/
    `SystemOut.log`
17. Start Workplace (or whatever you named your application) from the administrative console.

Completing post-upgrade Application Engine configuration

After you upgrade and deploy your Application Engine instances, you can complete additional configuration. Actions include modifying a configuration file, verifying Workplace sign-in, and establishing the Process Engine connection point.

To complete post-upgrade Application Engine configuration:
1. If it is not already running, start the Workplace application on your application server.
2. Sign in to Workplace to test your connection.
    a. On any computer, open a browser and type: `http://server_name:port#/Workplace`
    b. Enter a user name and password, and click **Sign in**.
3. Set the Process Engine Connection Point.

    The Process Router from version 3.5 has been replaced with the Application 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.
        1) In Workplace, click **Tasks.**
        2) Verify that the Tasks page displays.
    f. Sign out of Workplace.

Part 2. Removing software

Removing IBM FileNet P8 software can involve deleting one or more core components, expansion products, and the FileNet P8 documentation.

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*.

Removing the IBM FileNet P8 documentation

Removing the documentation for the FileNet P8 Platform and its expansion products involves a number of application server configuration possibilities. Use these examples only for reference. Your specific installation directories and application names might vary.

"Removing the FileNet P8 documentation from a WebSphere Application Server"
To remove the FileNet P8 documentation, delete the entire FileNet P8 documentation folder (for example, `ecm_help.war`) structure from the installation location.

Removing the FileNet P8 documentation from a WebSphere Application Server

To remove the FileNet P8 documentation, delete the entire FileNet P8 documentation folder (for example, `ecm_help.war`) structure from the installation location.

1. Log on to the WebSphere FileNet P8 documentation server.

Option	Description
UNIX	Log on as a user with delete access to where the FileNet P8 Platform documentation files are installed.

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.

   Important: Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

Removing Content Engine

You can uninstall an entire Content Engine installation or selected Content Engine components.

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.

"Removing an entire Content Engine installation interactively (UNIX)"
To remove an entire Content Engine installation, navigate to the *install_path*/FileNet/ContentEngine/_ceuninst directory and issue the following command: ce_uninstaller.

"Removing Content Engine silently"
In silent mode, the uninstaller removes all Content Engine components.

"Removing data associated with Content Engine"
After uninstalling Content Engine, you can remove its associated data.

Removing an entire Content Engine installation interactively (UNIX)

To remove an entire Content Engine installation, navigate to the *install_path*/FileNet/ContentEngine/_ceuninst directory and issue the following command: ce_uninstaller.

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

Removing Content Engine silently

In silent mode, the uninstaller removes all Content Engine components.

To uninstall Content Engine silently, run one of the following commands:

Option	Description
UNIX	ce_uninstaller -i silent

Removing data associated with Content Engine

After uninstalling Content Engine, you can remove its associated data.

1. Use the application server console or command lines to undeploy Content Engine.
2. Use the application server console or command lines to remove any database JDNI data sources associated with Content Engine object stores.
3. Use your database tools to drop any databases or table spaces for the object stores and the GCD.
4. Use your LDAP tools to delete users and groups you created in *Plan and Prepare Your Environment for IBM FileNet P8*.

5. Use your operating system commands to delete any directories, users, and groups that were used for installing and administering Content Engine.
6. Use your operating system commands to delete any file-storage-area directories that contain content, such as documents.
7. User your operating system commands to delete any index-area directories (K2 collections).

Removing Content Search Engine

To completely remove 4.5 Content Search Engine (Autonomy K2) and collections from your IBM FileNet P8 Platform installation, you must disable full-text indexing. You must also remove the Autonomy K2 installations from all servers associated with the Content Engine.

Important:

- If you remove or disable Autonomy K2 before you disable CBR and full-text indexing in Enterprise Manager, your system is rendered unusable and require considerable reconstruction.
- This procedure presumes you have a running installation of version 4.5 Content Search Engine and that you have existing collections.
- If you intend to remove Content Engine, skip to 4.

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** check box and click **OK**.
   d. A dialog asks if you want 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 reindex any of these objects that were previously enabled:
   - Document
   - Annotation
   - Custom Object
   - Folder

   The index job disables all full-text indexing and content-based retrieval settings and deletes any associated collections. After the index job is complete, continue with the next step.
4. As *k2_os_user*, log on to the Master Administration Server machine in your IBM FileNet P8 Content Search Engine configuration.
5. Complete the steps in the table to remove the Autonomy K2 software on the machine.

Option	Description
UNIX	1. Stop the Autonomy K2 Administration Server service and the Tomcat application server. Use the following command, according to your environment: **AIX** *verity_install_path*`/k2/` `_rs6k43/bin/k2adminstop` 2. Change directory to: *verity_install_path*`/_cseuninst` 3. Run the following command: `./CSE_uninstaller`

6. Delete *verity_install_path* on the machine.
7. Repeat 5 on page 199 and 6 for each remaining machine in your IBM FileNet P8 Content Search Engine configuration.

Removing additional locales (UNIX)

Remove the additional locales that you installed on UNIX machines.

To remove locales:
1. Change your current directory to *verity_install_path*/k2.
2. To remove all single-language locales, run the following command: _uninst_loc
3. To remove all European locales, run the following command: _uninst_eur

Removing Process Engine (UNIX)

Removal of the Process Engine on UNIX platforms requires that you stop IBM FileNet P8 software components and run specific commands.

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 `System 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, run the following command:

    `slibclean`

5. If the /fnsw/_pws_uninst2 directory exists, run the following command:

    `/fnsw/_pws_uninst2/pws_uninstall.bin`

    Otherwise run the following command:

    `/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.

Removing Application Engine (WebSphere)

Removal of the Application Engine for WebSphere on UNIX and Windows platforms requires that you log on to the application server and run the uninstall program.

To remove the Application Engine software:

1. Log on to the application server.

Option	Description
UNIX	Log on as a user with read, write, and execute access to the directory whereApplication Engine is installed.

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:

Option	Description
UNIX	uninstaller.bin

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.

   Important: 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.

Removing the Application Engine ISRA servlet

Removal of the Application Engine ISRA Servlet for Windows and UNIX environments requires that you logon to the application server, undeploy the servlet, and run the uninstall program.

Since the installed names for the ISRA Servlet are configurable on the supported application servers, the following information might 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.

Option	Description
UNIX	Log on as a user with write access to the directory where ISRA Servlet is installed.

2. Undeploy the ApplicationEngineISRAServlet application. This step is similar to that required to undeploy the Workplace application.

Option	Description
IBM WebSphere Application Server	1. Stop the ApplicationEngineISRAServlet process in the Admin console. 2. Uninstall the ApplicationEngineISRAServlet application from Enterprise Applications.

3. Navigate to the `/_uninstISRA` directory under the ISRA Servlet installation location.
4. Run the uninstall program:

Option	Description
UNIX	`uninstall.bin`

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 Application Server only) Delete the following temporary working folders for the Application Engine ISRA Servlet: *WAS_Home*`\WebSphere\ AppServer\profiles\default\installedApps\`*servername*`\ ApplicationEngineISRAServlet.ear\`

Removing Workplace Application Integration

You can remove Application Integration by using the Microsoft Windows Control Panel, or you can remove the program files silently from the command line.

Removing Workplace File Tracker

You can remove File Tracker by using the Microsoft Windows Control Panel, or you can remove the program files silently from the command line.

Part 3. Appendixes

Appendix A. 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.

You use Configuration Manager to define the following information for the Content Engine instance:

- Application server properties
- Java Database Connectivity (JDBC) data source properties for the Global Configuration Data (GCD) database
- Java Database Connectivity (JDBC) data source properties for each object store database
- Directory service (LDAP) provider properties
- Content Engine application login modules
- Content Engine bootstrap properties

"Overview of Configuration Manager" on page 216
You can use Configuration Manager to generate one or more unique Content Engine configuration profiles. A profile is a collection of information required to configure and deploy new or upgraded Content Engine instances.

"Handling passwords in Configuration Manager" on page 219
To provide the highest possible security, Configuration Manager's default settings do not save passwords from the GUI application. The password save setting is a preference setting in the Configuration Manager graphical interface. 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 a profile in Configuration Manager, the passwords are removed from memory. When you open Configuration Manager or the profile again to run a saved task, you will must reenter passwords to run the tasks.

"Accessing the Configuration Manager log files" on page 219
Configuration Manager maintains three types of log files: the Configuration Manager session log, the Configuration Managererror log, and the task status file.

"Adding an SSL signer to the Configuration Manager keystore (WebSphere)" on page 220
If you are using SSL for communication between Content Engine and WebSphere Application Server you might receive an SSL signer error when you test the connection to the application server or when you run the Deploy Application task. To resolve the issue, make sure that you have an entry for the SSL signer in the truststore that Configuration Manager uses.

"Correcting an SSL Signer Exchange Prompt error (WebSphere)" on page 221
If you are using SSL for communication between Content Engine and WebSphere Application Server, you might receive the SSL signer error SSL SIGNER EXCHANGE PROMPT when you run the task for configuring the JDBC data sources for the Global Configuration Data (GCD) or for an object store.

"Configuration Manager user interface reference" on page 222
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.

"Configuration Manager command-line reference" on page 238
Configuration Manager can be run from a command line. This section covers
the syntax for the command-line version of Configuration Manager.

Overview of Configuration Manager

You can use Configuration Manager to generate one or more unique Content
Engine configuration profiles. A profile is a collection of information required to
configure and deploy new or upgraded Content Engine instances.

"Configuration profile concepts"
The information for a profile is collected in XML files in the form of properties
and values that describe the associated configuration and deployment tasks.
You must provide values for the profile properties that are specific to each
configuration at your site, such as the application server name.

"Using the graphical and command-line user interfaces" on page 218
Both the graphical user interface (GUI) and the command-line interface (CLI)
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.

"Gathering Configuration Manager values by using the installation and
upgrade worksheet" on page 218
You can use the installation and upgrade worksheet to record the values that
you must enter in Configuration Manager.

Configuration profile concepts

The information for a profile is collected in XML files in the form of properties and
values that describe the associated configuration and deployment tasks. You must
provide values for the profile properties that are specific to each configuration at
your site, such as the application server name.

The XML files are stored in a directory that is unique to a given profile. Because
the profile name is used for both the directory name and the configuration file
name, you must provide a profile name that is a valid directory name for your
operating system. By default, the profiles are stored in the `ce_install_path`/tools/
configure/profiles directory, where `ce_install_path` is the location where Content
Engine is installed.

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 Configuration Manager to perform the following tasks that are associated with
a Content Engine configuration profile:

- **Set the application server properties**. Content Engine will be deployed as an
  application on the application server. You must 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. All profiles include the
  application server properties. By default, the application server properties are
  stored in the `ce_install_path`/tools/configure/profiles/`myprofile`/
  applicationserver.xml file, where `myprofile` is the name of your profile.

You provide the application server properties when you create a profile by using the Create New Installation Profile wizard, and you can edit the application server properties at any time as needed. See "Creating a profile for a new installation" on page 230 or "Editing the application server properties" on page 232 for detailed procedures.

- **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 must 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. By default, the JDBC properties are stored in the *ce_install_path*/tools/configure/profiles/*myprofile*/configurejdbcgcd.xml file or the the the *ce_install_path*/tools/configure/profiles/*myprofile*/configurejdbcos.xml file, where *myprofile* is the name of your profile.

  See "Editing the configure JDBC data source task" on page 233 for the procedure to set the JDBC data source properties after you have created a profile.

- **Configure the application login modules**. The login modules provide authentication information for the Content Engine application. Run this task to create the login modules on the application server.

  See "Editing the configure LDAP task" on page 233 for the procedure to create the login modules after you have created a profile.

- **Configure the directory service (LDAP) provider**. Content Engine connects to the directory service provider to authenticate users. Because the application server uses the directory service information to connect the Content Engine to the directory service provider, you cannot skip this task even if you have already configured your application server prior to installing Content Engine. 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 that you select determines some of the properties and their default values. By default, the LDAP properties are stored in the *ce_install_path*/tools/configure/profiles/*myprofile*/configureldap.xml file, where *myprofile* is the name of your profile.

  See "Editing the configure LDAP task" on page 233 for the procedure to set the LDAP properties after you have created a profile.

- **Configure the Content Engine bootstrap settings**. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine. By default, the bootstrap properties are stored in the *ce_install_path*/tools/configure/profiles/*myprofile*/configurebootstrap.xml file, where *myprofile* is the name of your profile.

  See "Editing the configure bootstrap properties settings" on page 233 for the procedure to set the bootstrap properties after you have created a profile.

- **Deploy the Content Engine application**. This action deploys the Content Engine EAR file with the JDBC, LDAP, and bootstrap settings on the application server. Any time that you update the properties for an existing deployed Content Engine instance or update the Process Engine Client files on the Content Engine, you must redeploy for the changes to take effect. By default, the deployment properties are stored in the *ce_install_path*/tools/configure/profiles/*myprofile*/deployapplication.xml file, where *myprofile* is the name of your profile.

Tip: After you deploy the Content Engine application, you use Enterprise Manager to customize Content Engine for your site's requirements.

See "Editing the deploy application task" on page 234 for the procedure to edit the deployment properties after you have created a profile.

- Check the status of a particular configuration task. Status messages are displayed when you run a task. You can also explicitly display the task status any time after you run the task. See "Checking the task status" on page 237.

- **Configure a profile for upgrading an existing Content Engine from 4**. An upgrade profile includes the Upgrade Bootstrap task and the Deploy Application task. By default, the upgrade bootstrap properties are stored in the *ce_install_path*/tools/configure/profiles/*myprofile*/upgradebootstrap.xml file, and the deployment properties are stored in the ce_install_path/tools/configure/profiles/*myprofile*/deployapplication.xml file, where *myprofile* is the name of your profile.

  See "Creating a profile for a new installation" on page 230 for a detailed procedure.

Using the graphical and command-line user interfaces

Both the graphical user interface (GUI) and the command-line interface (CLI) 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.

Configuration Manager has a GUI and a CLI. The GUI version of the tool displays the properties and default values that you must set. When you save your changes in the GUI tool, the configuration XML files are updated for you. If you use the CLI 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 run the tasks to apply the saved settings. After you have set the required values, you use either version of Configuration Manager to deploy the Content Engine application.

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 CLI version, and then use the GUI version to open the profile, edit the values, and run the configuration tasks.

Gathering Configuration Manager values by using the installation and upgrade worksheet

You can use the installation and upgrade worksheet to record the values that you must enter in Configuration Manager.

Be sure that you have available the Installation and Upgrade Worksheet that was completed during your planning activities.

To see only the properties you must specify for Configuration Manager:
1. Open your completed Installation and Upgrade Worksheet file.
2. Verify that the **Data** ⇒ **Filter** ⇒ **AutoFilter** command is enabled.
3. Filter by one of the following Configuration Manager options in the **Installation or Configuration Program** column:
   - **CM: config_mgr_user**
   - **CM: Create New Installation Profile**

- **CM: Configure GCD JDBC Data Sources**
- **CM: Configure Object Store JDBC Data Sources (object store 1)**
- **CM: Configure LDAP**
- **CM: Configure Bootstrap Properties**
- **CM: Deploy Application**
- **CM: Upgrade bootstrap**
- **CM: Upgrade Configuration Profile**

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 a preference setting in the Configuration Manager graphical interface. 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 a profile in Configuration Manager, the passwords are removed from memory. When you open Configuration Manager or the profile again to run a saved task, you will must reenter passwords to run the tasks.

The following passwords are used to run the tasks:

- The application server administrator password. Select **File → Edit Application Server Properties** to enter the password.
- The database administrator password. Edit the **Configure GCD JDBC Data Sources** task or the **Configure Object Store JDBC Data Sources** task.
- The directory service bind user password. Edit the **Configure LDAP** task.
- The bootstrap user password. Edit the **Configure Bootstrap** task.

The Configuration Manager command line passes the passwords from an XML configuration file to the required application when you run a task. You can use the **storepasswords** command to add encrypted passwords to the XML files, or you can enter plain text passwords when you edit the files. However, saving the passwords to the XML files might not be FIPS 140-2 compliant.

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.

Accessing the Configuration Manager log files

Configuration Manager maintains three types of log files: the Configuration Manager session log, the Configuration Managererror log, and the task status file.

- The Configuration Manager session log file. The session log contains brief information about the tasks that were run in the current session of the Configuration Manager graphical user interface (GUI). As you run additional tasks for the same or a different profile, new messages are added to the log. The messages in the session log are not time-stamped. When you exit Configuration Manager, the session log is cleared. You can view the session log from within the GUI, or you can open the session log file in a text editor:

ce_install_path/tools/configure/logs/P8ConfigMgr_Console.log

For more information about viewing the session log, see "Checking the task status" on page 237.

- The Configuration Manager error log. The error log contains details about any error occurring during execution of a task. The error messages are time-stamped, and include more information than is displayed in the session log. Unlike the session log, the error log is cumulative and contains data for more than one session. You can open the error log file in a text editor:

  ce_install_path/tools/configure/logs/P8Installer_common.log error.

- Task status file. When you run a task in Configuration Manager, the task execution messages are displayed in the Console pane (GUI) or on the command line (CLI), and the status messages are written to a text file in the following directory:

  ce_install_path/tools/configure/profiles/*myprofile*/status

  The file name is based on the task that was run. You can check the status on a task at any time to review the status completion messages. For more information about checking the task status from the GUI, see "Checking the task status" on page 237. For more information about checking the task status from the CLI, see "Running Configuration Manager commands" on page 238.

Adding an SSL signer to the Configuration Manager keystore (WebSphere)

If you are using SSL for communication between Content Engine and WebSphere Application Server you might receive an SSL signer error when you test the connection to the application server or when you run the Deploy Application task. To resolve the issue, make sure that you have an entry for the SSL signer in the truststore that Configuration Manager uses.

1. Identify the serial number for the SSL certificate on the web application server.
   a. From the server where Configuration Manager is installed, browse to the WebSphere administrative console address.
   b. In the Security Alert dialog box, click **View Certificate**.
   c. Click the **Details** tab.
   d. Record the value for **Serial number** for the certificate.
   e. Click **OK** to dismiss the Certificate dialog box.
   f. Click **Yes** in the Security Alert dialog box to proceed.
2. Identify the truststore location and filename.
   a. Log in to the WebSphere administrative console.
   b. Select **Security** → **SSL certificate and key management**.
   c. Select **SSL configurations**.
   d. Click the default SSL setting, **NodeDefaultSSLSettings**.
   e. Under the **Related items** link, click **Key stores and certificates**.
   f. Record the filename in the **Path** column for the truststore **NodeDefaultTrustStore**, such as trust.p12.
3. Start IBM Key Management by entering one of the following commands at a command prompt:

Option	Description
UNIX	*WAS-Home*/AppServer/bin/ikeyman.sh

4. Select **Keybase File** → **Open**.

a. For the **Key database type**, select **PKCS12**.

b. Click **Browse** to locate the filename you recorded in step 2 on page 220. For example, the File Name filed contains the filename, such as `trust.p12`. The **Location** field contains the absolute path to the truststore, such as `C:\Program Files\IBM\WebSphere\AppServer\profiles\AppSrv01\etc\` for Windows.

c. Click **OK**.

d. Enter the password and click **OK**. The default password is `WebAS`.

5. Locate the signer certificate with the serial number that matches the serial number that you recorded in step 1 on page 220.

a. Double-click a certificate name other than **default_signer** to view the serial number for the certificate.

b. Click **OK** to close the dialog box.

c. Repeat until you have located the correct signer certificate.

6. Extract the certificate.

a. Select the signer certificate with the correct serial number, and click **Extract**.

b. Provide a name and location, and then click **OK**.

7. Add the certificate that you extracted to the trust file for Content Engine.

a. Open the `DummyClientTrustFile.jks` key database file located in the WebSphere profile for Content Engine, such as `C:\Program Files\IBM\WebSphere\AppServer\profiles\AppSrv01\etc\` for Windows.

b. Add the certificate that you extracted in step 6.

8. Close IBM Key Management.

Correcting an SSL Signer Exchange Prompt error (WebSphere)

If you are using SSL for communication between Content Engine and WebSphere Application Server, you might receive the SSL signer error SSL SIGNER EXCHANGE PROMPT when you run the task for configuring the JDBC data sources for the Global Configuration Data (GCD) or for an object store.

Use this procedure if you receive the following error when you run a task for configuring the JDBC data sources, such as **Configure GCD JDBC Data Sources**:

```
Error while executing Configure GCD JDBC Data Sources
Execution failed with the following message: The data source configuration failed.
*** SSL SIGNER EXCHANGE PROMPT ***
SSL signer from target host null is not found in trust store
```

1. From a command prompt, navigate to the *WAS_HOME*/AppServer/profiles/*app_server_name*/bin directory.

2. Run the following command:

```
wsadmin.bat -conntype SOAP -port portnumber -username username
 -password userpassword
```

Where:

portnumber
    is the same value that you entered for the **Application server SOAP** field in the configure JDBC task properties. The WebSphere default is 8880.

username
    is the same value that you entered for the **Application server administrator user name** field in the configure JDBC task properties.

userpassword
> is the same value that you entered for the **Application server administrator password** field in the configure JDBC task properties.

3. At the prompt to add the signer to the trust store, enter `Yes`.

4. Close the command prompt.

5. Use Configuration Manager to run the configure JDBC task again, from the graphical user interface or from the command line.

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.

Restriction: 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 238.

"Starting Configuration Manager"
You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.

"Configuration Manager window" on page 223
The default 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 on the lower right side. You can drag, resize, and rearrange the panes to change the location if needed.

"Configuration Manager menus and commands" on page 226
You can access the Configuration Manager commands through the menu or by right-clicking an item. The following tables describe the available commands.

"Working with Configuration Manager" on page 228
You can use the Configuration Manager commands, icons, and panes to work with your profiles and configuration tasks.

"configmgr.ini parameters" on page 238
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/configmgr.ini file.

Starting Configuration Manager

You can start the graphical interface version of Configuration Manager to configure a Content Engine application instance on a web application server.

To start Configuration Manager:

1. 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:

Option	Description
UNIX	Run this command: `ce_install_path`/tools/configure/configmgr

The first time that you start Configuration Manager, the Welcome is displayed.

2. Select one of the links in the Welcome to learn more or to start working in a wizard, or close the Welcome by clicking the **X** in the tab at the upper left. You can reopen the Welcome later, as needed, from the **Help** menu.

Configuration Manager window

The default 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 on the lower right side. You can drag, resize, and rearrange the panes to change the location if needed.

The following table describes each pane:

Pane	Description
Content Engine Task View	Displays a profile and the tasks for that profile. Only one profile can be open at a time.
Task Editor	Displays the properties and values for a selected task. The Task Editor pane is empty until a specific task is selected from the Content Engine Task View pane. Each open task is displayed in a separate tab in the Task Editor pane. More than one task tab can be displayed at a time.
Console	Displays task execution messages, results from the Check Status command, or the session log.

To move from one open task tab in the Task Editor pane to another open task tab, press `Ctrl-F6`.

To move from one pane to the next, press `Ctrl-F7`.

To restore a closed pane, select **Window** → **Show View**.

"Main toolbar"
The Configuration Manager main toolbar is located just below the menu bar. The main toolbar contains multiple icons for working with your profiles.

"Profile toolbar" on page 224
The Configuration Manager profile toolbar is located at the upper right of the Content Engine Task View pane. The profile toolbar contains multiple icons for working with a selected profile.

"Task toolbar" on page 224
The Configuration Manager task toolbar is located at the upper right of the Task Editor pane when a Task tab is open.

"Console toolbar" on page 225
The Console toolbar is located at the upper right of the Console pane. The Console toolbar contains the multiple icons for working with the Console pane.

Main toolbar

The Configuration Manager main toolbar is located just below the menu bar. The main toolbar contains multiple icons for working with your profiles.

Icon	Command name and Description
	New Installation Profile
	Click this icon to create a profile for a new installation. 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.
	New Upgrade Profile
	Click this icon to create a 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. The New Configuration Profile Wizard starts.
	Open 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.
	Save
	Click this icon to save the current configuration profile settings.
	View Configuration Manager Log File
	Click this icon to view the Configuration Manager session log. The session log is cleared when you open Configuration Manager or when you close a profile.

Profile toolbar

The Configuration Manager profile toolbar is located at the upper right of the Content Engine Task View pane. The profile toolbar contains multiple icons for working with a selected profile.

Icon	Command name and Description
	Edit 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.

Task toolbar

The Configuration Manager task toolbar is located at the upper right of the Task Editor pane when a Task tab is open.

The task toolbar contains the following icons for working with the current task:

Icon or Button	Command Name and Description
	Enable the Selected Task Click this icon to enable a task. The Run All Tasks command only runs tasks that are enabled.
	Disable the Selected Task Click this icon to disable a task. The Run All Tasks command only runs tasks that are enabled.
	Save Click this icon to save the current task.
Test Database Connection	Test Database Connection Configure GCD JDBC Data Sources task or Configure Object Store JDBC Data Sources task tab only. Click this button to test the connection to the database by using the values provided.
Test LDAP Connection	Test LDAP Connection Configure LDAP task tab only. Click this button to test the connection to the directory service provider by using the values provided.

Console toolbar

The Console toolbar is located at the upper right of the Console pane. The Console toolbar contains the multiple 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 the 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.

Icon	Command name and Description
	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.

Configuration Manager menus and commands

You can access the Configuration Manager commands through the menu or by right-clicking an item. The following tables describe the available commands.

Main menu

The following table lists the menus and commands that are available in Configuration Manager.

Table 26. Main Menu Commands

Menu name	Command name	Description
File		Provides commands for creating, saving, or opening a configuration profile.
	New Installation Profile	Creates a profile for a new installation. See "Creating a profile for a new installation" on page 230 for detailed procedures to create a profile.
	New Upgrade Profile	Creates a profile for upgrading Content Engine. See "Creating a profile for a new installation" on page 230 for detailed procedures to create a profile.
	Open Profile	Opens an existing profile for viewing or editing. Keyboard shortcut: Ctrl+O
	Close Profile	Closes the current 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.
	Edit Application Server Properties	Opens the Edit Application Server Properties wizard for editing the application server values for the profile. Keyboard shortcut: Ctrl+U
	Run All	Runs all of the tasks in the profile to apply your settings. If a particular task is disabled, that task is skipped when you select Run All Tasks. Keyboard shortcut: Ctrl+R

Table 26. Main Menu Commands (continued)

Menu name	Command name	Description
	Exit	Closes Configuration Manager.
Window		Provides commands for viewing a log file, changing the view in a Configuration Manager pane, and setting preferences
	View Log File	Provides commands that give you access to logs and various views of information.
	Show View	Displays the session log in the Console pane. The session log lists results from tasks that you have run since you opened the current profile.
	Preferences	Provides choices for determining the behavior of Configuration Manager, such as the password save preference and the warning message preference.
Help		Displays help pages about Configuration Manager.
	Welcome	Provides a quick introduction to the use of Configuration Manager.
	Help Contents	Provides reference information for Configuration Manager.
	About Configuration Manager	Displays the copyright and related information about Configuration Manager.

Pop-Up menus

The pop-up menus are displayed when you right-click an item in Configuration Manager.

Table 27. Profile Icon Pop-Up Menu Commands

Command	Description
Edit Application Server Properties	Opens the Edit Application Server Properties wizard for the profile.
Add New Task > Configure JDBC Data Sources	Adds a New_Configure JDBC Data Sources task to the profile.
Add New Task > Configure LDAP	Adds a New_Configure LDAP task to the profile.
Add New Task > Configure Bootstrap	Adds a Configure Bootstrap task to the profile.
Run All	Displays a sub-menu of the tasks associated with the profile. Click the task name that you want to run.
Close Configuration Profile	Closes the configuration profile.

Table 28. Task Icon Pop-Up Menu Commands

Command	Description
Edit Selected Task	Opens the task in the task pane for editing.

Table 28. Task Icon Pop-Up Menu Commands (continued)

Command	Description
Enable Selected Task Disable Selected Task	Toggles the state of the task between enabled and disabled. Any task that is disabled will not run with either the **Run Selected Task** command or the Run All Tasks command. If the task is currently enabled, selecting the **Disable Selected Task** command prevents the task from running, changes the font for the task icon to italic, and appends "(disabled)" to the icon label. If the task is currently disabled, selecting the **Enable Selected Task** command allows the task to run, restores the original font, and removes "(disabled)" from the icon label.
Copy the Selected Task	Creates a copy of the selected task. This command is available only for the Configure JDBC Data Sources and the Configure LDAP tasks.
Add New Task > Configure JDBC Data Sources	Adds a Configure JDBC Data Sources task to the profile. You can have multiple Configure JDBC Data Sources tasks in a profile.
Add New Task > Configure Login Modules	Adds a Configure Login Modules task to the profile. You can only have one Configure Login Modules task in a profile.
Add New Task > Configure LDAP	Adds a Configure LDAP task to the profile. You can have multiple Configure LDAP tasks in a profile.
Add New Task > Configure Bootstrap Properties	Adds a Configure Bootstrap Properties task to the profile. You can have multiple Configure Bootstrap Properties tasks in a profile.
Add New Task > Deploy Application	Adds a Deploy Application task to the profile. You can only one Deploy Application task in a profile.
Reset Selected Task Status	Resets the status of the task to indicate that it has not yet been run.
Check Task Status	Displays the current status of the task.
Rename Task	Opens an edit window in which you can change the name of the label associated with the task.
Delete the Selected Task	Deletes the task from the profile. You cannot delete a task if it is open for editing in the task pane.
Move Selected Task Up	Moves the selected task up in the list of tasks.
Move Selected Task Down	Moves the selected task down in the list of tasks.

Working with Configuration Manager

You can use the Configuration Manager commands, icons, and panes to work with your profiles and configuration tasks.

"Configuring a Content Engine instance" on page 229
You can use the Configuration Manager graphical user interface to configure a Content Engine instance.

"Setting the password save preference" on page 230
By default, the Configuration Manager password save preference is set to not save passwords to a file. If your site security requirements permit you to save passwords to a file, you can change the password save preference setting.

"Creating a profile for a new installation" on page 230
You must create a new configuration profile for each Content Engine
application that you deploy.

"Creating a profile for an upgrade" on page 231
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.

"Opening and closing an existing profile or task" on page 231
You can save a profile, and open it later to edit the saved settings or to run
tasks.

"Editing the application server properties" on page 232
You initially provide the application server properties when you create a new
profile. You can open the application server properties for editing at any time,
but you cannot change the application server type for an existing profile. The
application server properties must be set before you run any tasks.

"Editing the properties for a specific task" on page 232
You must provide the required property values for each task in your profile
before you run the task.

"Applying the property settings by running a specific task" on page 234
You must run a task to apply the values that you provided.

"Adding an additional task to your profile" on page 235
A profile can contain an unlimited number of tasks. You can configure more
than one Configure Object Store JDBC Data Sources task, more than one
Configure LDAP task, or more than one Configure Bootstrap Properties task.
You can add a new task for a task type that you do not already have in the
profile.

"Deleting a task from a profile" on page 236
You can delete any unneeded tasks from your profile.

"Running all tasks at the same time" on page 237
You can run all the tasks in a profile at the same time. Disabled tasks will not
run.

"Running a single task" on page 237
You can run each configuration task individually. Disabled tasks will not run.

"Checking the task status" on page 237
Task execution messages are displayed in the console pane when you run a
task, and you can view the status of a specific task at any time.

"Viewing the session log" on page 238
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.

"Saving your changes to a task or profile" on page 238
You can save your profile and task settings at any time.

Configuring a Content Engine instance

You can use the Configuration Manager graphical user interface to configure a
Content Engine instance.

To configure a Content Engine instance:

1. Create a configuration profile. See "Creating a profile for a new installation" on
   page 230.

2. Edit the configuration tasks included in the profile. See one or more of the following topics:
    - "Editing the configure JDBC data source task" on page 233
    - "Editing the configure LDAP task" on page 233
    - "Editing the configure bootstrap properties settings" on page 233
3. Apply the configuration settings by running the tasks. See "Running all tasks at the same time" on page 237.
4. Deploy the application by running the Deploy Application task. Because deployment can take a long time, it is a best practice to run the Deploy Application task after you have completed all other configuration tasks. See "Editing the deploy application task" on page 234.

Setting the password save preference

By default, the Configuration Manager password save preference is set to not save passwords to a file. If your site security requirements permit you to save passwords to a file, you can change the password save preference setting.

When you close the profile, the passwords are erased from memory. Each time that you start Configuration Manager or open a saved profile, you must specify any passwords required by the tasks and for the application server properties before you run the tasks; otherwise, the tasks won't run successfully. If your site security requirements permit you to save passwords to a file, you can change the password save preference setting.

To change the password save preference:
1. Click **Window** → **Preferences**.
2. Complete one of the following actions:

Option	Description
To save passwords to file	Select the **Save all passwords to file when saving a task or profile** check box.
To prevent writing passwords to file	Clear the **Save all passwords to file when saving a task or profile** check box.

3. Click **OK**.

Creating a profile for a new installation

You must create a new configuration profile for each Content Engine application that you deploy.

To create a new configuration profile:
1. Start the Create New Installation Profile wizard.
    - Click the **New Installation Profile** icon in the toolbar.
    - Select **File** → **New Installation 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. If you selected **Yes** and your current profile has any unsaved changes, the Save Resource message box opens.
    b. Click **Yes** to save your changes, or click **No** to continue without saving your changes.

       The Create New Configuration Profile Wizard opens.

3. Complete the wizard screens. For details on the fields in the wizard screens, hover your mouse over the field name to view the property description.

The new profile is displayed as an icon in the left-hand panewith icons for the tasks for the tasks that you selected. The default profile contains the following tasks:

- Configure GCD JDBC Data Sources
- Configure Object Store JDBC Data Sources
- Configure Login Modules
- Configure LDAP
- Configure Bootstrap Properties
- Deploy Application (Disabled)

Creating a profile for an upgrade

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 by using one of these methods:
   - Click **Create a Profile for an Upgrade Installation** in the toolbar.
   - Select **File → New Upgrade 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. Complete the wizard screens. For details on the fields in the wizard screens, hover your mouse over the field name to view the property description.

The profile you created is displayed as an icon in the left-hand pane with the name that was presented in the confirmation screenwith icons for the **Configure Bootstrap** task and the **Deploy Application** task.

Opening and closing an existing profile or task

You can save a profile, and open it later to edit the saved settings or to run tasks.

1. Open a profile:
   - Click **Open an Existing Profile** in the toolbar.
   - Select **File → Open Profile**.
   a. If a profile is already open, the Action Required message box opens. Respond to the messages as follows:
      1) Click **Yes** to continue creating a new profile, or click **No** to cancel.
      2) 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**.
   b. Either type in the fully qualified path to the *myprofile*.cfgp profile file, or click **Browse** to locate the file.
   c. Click **OK**.
2. Open a task:
   More than one task tab can be open at a time in the Task Editor pane.

a. If the profile is collapsed in the Content Engine Task View pane, click **+** next to the profile name to expand it.

b. 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.

3. Switch between open tasks in the Task Editor pane by clicking the *tab name* in the Task Editor pane for the desired task.

4. Close a task:
   a. If the task is not the actively selected task, click the *tab name* in the Task Editor pane for the desired task.

   b. Click **Close** in the tab for the task. You will be prompted to save any changes to the task.

5. Close the profile by selecting **File** → **Close Profile**.

   You will be prompted to save any changes when you close a profile. Passwords are removed from memory when you close a profile.

Editing the application server properties

You initially provide the application server properties when you create a new profile. You can open the application server properties for editing at any time, but you cannot change the application server type for an existing profile. The application server properties must be set before you run any tasks.

1. Start the Edit Application Server Properties wizard by using one of theses methods:
   - Click **Edit Application Server Properties** in the main toolbar.
   - Select **File** → **Edit Application Server properties**.

2. Provide values for the application server properties. For details on the fields, hover your mouse over the field name to view the property description.

3. Optional: WebSphere Application Server and Oracle WebLogic Server only. In the Set Properties for Application Server window, click **Test Connection** to test the connection between Configuration Manager and the application server by using the information that you have provided. The test is optional, and you can proceed in the wizard even if the test fails. If the test fails, make sure that the application server is running and that the application server property values that you entered match the values that are defined in your application server.

4. Click **Finish**.

Editing the properties for a specific task

You must provide the required property values for each task in your profile before you run the task.

"Editing the configure JDBC data source task" on page 233
The JDBC data source information is used by Content Engine to connect to Global Configuration Data and object store databases. Configuration Manager provides two tasks for configuring the JDBC data sources: Configure GCD JDBC Data Sources and Configure Object Store JDBC Data Sources.

"Editing the configure LDAP task" on page 233
The LDAP information is used to connect Content Engine to the directory service provider to authenticate users.

"Editing the configure bootstrap properties settings" on page 233
The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine.

"Editing the deploy application task" on page 234
You can edit the deployment property values without applying the settings.
When you apply the settings, the Content Engine is deployed as an application
on the application server. Because deploying an application can take time, we
recommend that you do not deploy the application until after you have
installed any dependent files, such as Process Engine client files, Content Search
Engine, or customized applications for Content Engine.

Editing the configure JDBC data source task:

The JDBC data source information is used by Content Engine to connect to Global
Configuration Data and object store databases. Configuration Manager provides
two tasks for configuring the JDBC data sources: Configure GCD JDBC Data
Sources and Configure Object Store JDBC Data Sources.

The procedure for editing the Configure GCD JDBC Data Sources task and the
Configure Object Store JDBC Data Sources task is the same. To complete the
Content Engine configuration, you must configure the GCD data sources and the
data sources for each object store.

1. Open either the GCD data source task or the object store data source task for
   editing.
   - Double-click **Configure GCD JDBC Data Sources**.
   - Double-click **Configure Object Store JDBC Data Sources**.
2. Provide the property values for your database. Place your mouse on a field
   name to view the property description.
3. Optional: WebSphere Application Server and Oracle WebLogic Server only.
   Click **Test Database Connection** to test the connection to the database by using
   the database user name, database server name, database name, port number,
   and password that you provided. The test does not create the data sources.
4. Select **File** ‣ **Save** to save your changes. Saving your changes to disk does not
   apply the settings to the application server.

Editing the configure LDAP task:

The LDAP information is used to connect Content Engine to the directory service
provider to authenticate users.

1. Double-click **Configure LDAP** in the Content Engine Task View pane to open
   the task for editing.
2. Provide the property values for your LDAP provider. For details on the fields,
   hover your mouse over the field name to view the property description.
3. Optional: WebSphere Application Server and Oracle WebLogic Server only.
   Click **Test LDAP Connection** to test the connection to the directory service
   provider by using the directory service bind user name, host name, port
   number, and password that you provided.
4. Select **File** ‣ **Save** to save your changes.

Editing the configure bootstrap properties settings:

The bootstrap information is needed for creating the Global Configuration Data
and for starting Content Engine.

1. Double-click **Configure Bootstrap Properties** in the Content Engine Task View
   pane to open the task for editing.
2. Select a value for the **Bootstrap operation** field:

Option	Description
To provide bootstrap information for a new installation profile	Select **Create New**.
To modify an existing deployed EAR file	Select **Modify Existing**.
To upgrade an existing deployed EAR file	Select **Upgrade**.

3. Provide the bootstrap property values. For details on the fields, hover your mouse over the field name to view the property description.
4. Optional: If you selected **Modify Existing** or **Upgrade**, then click **Verify Bootstrap Information** to display the bootstrap information from the modified EAR file. You can verify the information in the modified EAR file for a new installation after you run the task.
5. Select **File** ‣ **Save** to save your changes.

Editing the deploy application task:

You can edit the deployment property values without applying the settings. When you apply the settings, the Content Engine is deployed as an application on the application server. Because deploying an application can take time, we recommend that you do not deploy the application until after you have installed any dependent files, such as Process Engine client files, Content Search Engine, or customized applications for Content Engine.

To edit the Deploy Application task:
1. Right-click the **Deploy Application** task in the profile pane (left pane), and select **Edit Selected Task**.
2. Provide the property values for your deployment. For details on the fields, hover your mouse over the field name to view the property description..
3. Select **File** ‣ **Save**.

Applying the property settings by running a specific task
You must run a task to apply the values that you provided.

"Applying the JDBC data source settings"
Your JDBC data source settings are stored when you save the task, but the settings are not applied to the application server until you run the task.

"Applying the LDAP settings" on page 235
Your LDAP settings are stored when you save the task, but the settings are not applied to the application server until you run the task.

"Applying the bootstrap properties settings" on page 235
The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine. Your bootstrap settings are stored when you save the task, but the settings are not applied to the application server until you run the task.

"Deploying the application" on page 235
Your Deploy Application task settings are stored when you save the task, but the Content Engine EAR file is not deployed to the application server until you run the task.

Applying the JDBC data source settings:

Your JDBC data source settings are stored when you save the task, but the settings are not applied to the application server until you run the task.

1. Select the configure data sources task that you want to run.
   - Right-click **Configure GCD JDBC Data Sources** in the Content Engine Task View pane, and select **Run Task**.
   - Right-click **Configure Object Store JDBC Data Sources** in the Content Engine Task View pane, and select **Run Task**.

   Running the configuration task can take a few minutes.
2. Close the Configure JDBC Data Sources task pane.

The task execution status messages are displayed in the Console pane below the JDBC data source properties.

Applying the LDAP settings:

Your LDAP settings are stored when you save the task, but the settings are not applied to the application server until you run the task.
1. Right-click **Configure LDAP** in the Content Engine Task View pane, and then select **Run Task**. Running the configuration task can take a few minutes.
2. Close the Configure LDAP task pane.

The task execution status messages are displayed in the Console pane below the LDAP properties.

Applying the bootstrap properties settings:

The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine. Your bootstrap settings are stored when you save the task, but the settings are not applied to the application server until you run the task.
1. Right-click **Configure Bootstrap Properties** in the Content Engine Task View pane, and then select **Run Task**. Running the configuration task can take a few minutes.
2. Optional: If you selected **Create New**, then click **Verify Bootstrap Information** to display the bootstrap information from the modified EAR file.
3. Close the Configure Bootstrap properties task pane.

The task execution status messages are displayed in the Console pane below the bootstrap properties.

Deploying the application:

Your Deploy Application task settings are stored when you save the task, but the Content Engine EAR file is not deployed to the application server until you run the task.
1. Right-click the **Deploy Application** task in the left pane, and then select **Run Task**. Running the deployment task will take a few minutes.
2. Close the Deploy Application properties task pane.

The task execution status messages are displayed in the Console pane below the Deploy Application task properties.

Adding an additional task to your profile

A profile can contain an unlimited number of tasks. You can configure more than one Configure Object Store JDBC Data Sources task, more than one Configure

LDAP task, or more than one Configure Bootstrap Properties task. You can add a
new task for a task type that you do not already have in the profile.

The default configuration profile contains the following tasks:
- Configure GCD JDBC Data Sources
- Configure Object Store JDBC Data Sources
- Configure LDAP
- Configure Bootstrap Properties
- Deploy Application

You can add a new task whenever the following conditions apply:
- You do not already have an existing task of that type in the profile. You can
  have only one Configure GCD JDBC Data Sources task or one Deploy
  Application task.
- You have more than one object store in your site.
- You have multiple LDAP realms in your site.
- You must modify or upgrade you bootstrap properties.

To add a new task to your profile:
1. If your configuration profile is not open in Configuration Manager, open the
   profile.
2. Right-click any *task name* in the profile pane, and select **Add New Task**.
3. In the pop-up menu, select the *task name* that you want to create. The new
   task is added to the profile. If a task of the same type exists, the new task name
   begins with "New_."
4. Optional: Rename the task to provide a more useful name.
   a. Right-click the new *task name*, and select **Rename Task**.
   b. Enter a useful name for the task, such as `Configure Object Store 2 JDBC
      Data Sources` or `Configure Bootstrap Properties for Upgrade`.
   c. Click **OK**.

Deleting a task from a profile
You can delete any unneeded tasks from your profile.

The default configuration profile contains the following tasks:
- Configure GCD JDBC Data Sources
- Configure Object Store JDBC Data Sources
- Configure LDAP
- Configure Bootstrap Properties
- Deploy Application

When you delete a task, all the property values are removed. If you later need the
task, you will have to add a new task and re-enter all your values.

To delete a task from your profile:
1. If your configuration profile is not open in Configuration Manager, open the
   profile.
2. Right-click the *task name* that you want to remove, and select **Delete Selected
   Task**.
3. In the confirmation message box, click **OK**.

Running all tasks at the same time

You can run all the tasks in a profile at the same time. Disabled tasks will not run.

To run all tasks in a profile:

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 All Tasks command:
   - Click **Run All Tasks** in the Content Engine Task View pane toolbar.
   - Select **File → Run All**.

   Important: Tasks will not complete if you have not entered passwords correctly during the current Configuration Manager session. See "Handling passwords in Configuration Manager" on page 219.

The Console pane displays the task execution messages.

Running a single task

You can run each configuration task individually. Disabled tasks will not run.

To run a selected task:

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.

   Important: Tasks will not complete if you have not entered passwords correctly during the current Configuration Manager session. See "Handling passwords in Configuration Manager" on page 219.

The Console pane displays the task execution messages.

Checking the task status

Task execution messages are displayed in the console pane when you run a task, and you can view the status of a specific task at any time.

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 the *task name* in the Content Engine Task View pane, and then 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.

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.

Status Result	Description
INCOMPLETE	The task is incomplete.
NO STATUS AVAILABLE	The task has not been run.
FAILED	The task failed to complete. Additional information about the failure is displayed.

Viewing 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. Subsequently running a task will replace the displayed session log with the current task status. To redisplay the session log, repeat this procedure as needed.

Saving your changes to a task or profile

You can save your profile and task settings at any time.

Use one of these methods to save changes to the currently open task:
- Click **Save** in the toolbar.
- Select **File** › **Save**.

configmgr.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/configmgr.ini file.

UNIX example

```
-vm
/opt/IBM/FileNet/ContentEngine/_cejvm/bin
```

Configuration Manager command-line reference

Configuration Manager can be run from a command line. This section covers the syntax for the command-line version of Configuration Manager.

"Running Configuration Manager commands"
From the command line, you can generate the configuration XML files, run a configuration task to apply the settings in the configuration XML file, and check the status of a task. Use the following syntax to enter the Configuration Manager commands.

Running Configuration Manager commands

From the command line, you can generate the configuration XML files, run a configuration task to apply the settings in the configuration XML file, and check the status of a task. Use the following syntax to enter the Configuration Manager commands.

Use one of these commands to run the Configuration Manager command line:

Operating System	File Name
UNIX	`ce_install_path`/tools/configure/`configmgr_cl`

"How to read the syntax diagrams" on page 240
Syntax diagrams describe how you must enter commands and what options are available.

"**checkstatus** command" on page 240
The **checkstatus** command checks the status of the specified configuration task. The command name is not case sensitive.

"**execute** command" on page 244
The **execute** command applies the settings from a configuration XML file for the specified configuration task. The command name is not case sensitive.

"**generateconfig** command" on page 248
The **generateconfig** command generates the configuration XML file for the specified configuration task. The command name is not case sensitive.

"**generateupgrade** command (with Servers.xml file)" on page 253
The **generateupgrade** command generates the configuration XML files for upgrading an existing Content Engine. The **generateupgrade** command has two syntax formats, depending on whether you have a valid copy of the Servers.xml file in the directory where Content Engine is installed. The command name is not case sensitive.

"**generateupgrade** command (without Servers.xml file)" on page 255
The **generateupgrade** command generates the configuration XML files for upgrading an existing Content Engine. The **generateupgrade** command has two syntax formats, depending on whether you have a valid copy of the Servers.xml file in the directory where Content Engine is installed. The command name is not case sensitive.

"**gui** command" on page 257
The **gui** command opens the Configuration Manager graphical user interface. The command is not case sensitive.

"**listservers** command" on page 257
The **listservers** command displays a list of the currently deployed Content Engine servers. Only servers that are included in the `ce_install_path\`Servers.xml file are listed, where `ce_install_path` is the location where Content Engine is installed. You must supply the server names when you create an upgrade profile with the **generateUpgrade** command.

"**listtasks** command" on page 257
The **listtasks** command displays a list of the tasks and the task files in the configuration profile. The command name is not case sensitive.

"**movetask** command" on page 260
The **movetask** command moves a task to a different position in the list of tasks. The task position determines the order that the tasks are run when you run all the tasks at the same time. You use the **listtask** command to display the task order. The command names are not case sensitive.

"**removetask** command" on page 263
The **removetask** command removes the specified task from the configuration profile. When you remove the task, the configuration XML file is deleted from the profile directory. The command name is not case sensitive.

"**storepasswords** command" on page 266
The **storepasswords** command prompts for passwords that are blank in a

profile and stores the encrypted passwords in the file. Storing encrypted passwords might not be FIPS 140-2 compliant. You can run the command for a single task or for all tasks in the profile. The command name is not case sensitive.

How to read the syntax diagrams

Syntax diagrams describe how you must enter commands and what options are available.

The syntax topics uses several conventions to indicate variable, parameters, required items, and optional items. Enter all commands on a single line, even if the command or syntax examples wrap to the next line.

```
command_name -option_1 variable_1 [-option_2 variable_2]
 [-option_3 variable_3 | -option_4]
```

Where:

command_name
> The **command_name** is required.

-option_1
> The **-option_1** parameter is a required parameter.

variable_1
> The *variable_1* value is a required variable for the **-option_1** parameter.

[-option_2 *variable_2***]**
> Square braces [] indicate optional items. The **-option_2** parameter with its value is optional.

-option_3 *variable_3* **| -option_4**
> A vertical bar indicates a choice of parameters. Use the **-option_3** parameter with its value, or use the **-option_4** parameter. In this example, both items are optional, and you can use only one or the other.

checkstatus command

The **checkstatus** command checks the status of the specified configuration task. The command name is not case sensitive.

checkstatus command syntax

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl checkstatus [-task task_type | -taskfile task_file_name]
 -profile myprofile [-help]
```

checkstatus command parameters

-task *task_type*
> The **-task** *task_type* parameter specifies which task to use for the status check. You can omit the **-task** *task_type* parameter if you want to check all of the tasks or if you specify the **-taskfile** *task_file_name* parameter. The *task_type* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 29. task_type values

Option	Configuration file	Description

omitted	`configurebootstrap.xml`
	
	 `configurebootstrap.n.xml`
	
	 `configurejdbcgcd.xml`
	
	 `configurejdbcos.xml`
	
	 `configurejdbcos.n.xml`
	
	 `configureldap.xml`
	
	 `configureldap.n.xml`
	
	 `configureloginmodules.xml`
	
	 `deployapplication.xml`
	
	 Where *n* is an integer starting with 2
		If you omit the **-task** *task_type* parameter and the **-taskfile** *task_file_name* parameter, then the status for all the configuration files in the profile is displayed.

configurebootstrap	`configurebootstrap.xml` `configurebootstrap.n.xml`, where *n* is an integer starting with 2	Checks status for applying the settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named `configurebootstrap.2.xml`. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. For example, you might must check the status for the initial bootstrap task for the new profile and to check the status for modifying the EAR file to apply a fix pack. If your profile contains more than one `configurebootstrap` task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to check.
configurejdbcgcd	`configurejdbcgcd.xml`	Checks the status for configuring the JDBC connections to the Global Configuration Data (GCD) database used by Content Engine.

Table 29. task_type values (continued)

Option	Configuration file	Description
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Checks the status for configuring JDBC connections to a single object store database used by Content Engine. You need to generate, edit, and apply the configurejdbcos task settings for each object store in your database. When you generate a second object store JDBC configuration file, it is named configurejdbcos.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. If your profile contains more than one configurejdbcos task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to check.
configureLDAP	configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Checks status for configuring the connection to the directory server for authenticating Content Engine users. (WebSphere Application Server or Oracle WebLogic Server only.) If you generate a second LDAP configuration file, then it is named configureldap.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. If your profile contains more than one ConfigureLDAP task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to check.

Table 29. task_type values (continued)

Option	Configuration file	Description
configureloginmodules	configureloginmodules.xml	(WebSphere Application Server or JBoss Application Server only.) Checks status for creating the login modules for the Content Engine application.
deployapplication	deployapplication.xml	Checks status for deploying the Content Engine application on the application server.

-taskfile *task_file_name*

The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.

If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.

If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*

The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
- The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
- The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-help

The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

checkstatus command examples

The following examples include line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

Check the status of a profile with only one file for the configurejdbcos task.

The following command checks the status of the configurejdbcos task in the *ce_install_path*/tools/configure/profiles/wstdb2jdbc directory:

```
configmgr_cl checkstatus -task configurejbcos -profile wstdb2jdbc
```

Check the status of a profile with several files for the `configurejdbcos` task.
>The following command checks the status of the configurejdbcos.2.xml
>task file in the *ce_install_path*/tools/configure/profiles/wstdb2jdbc
>directory:
>
>```
>configmgr_cl checkstatus -taskfile configurejdbcos.2.xml
>-profile wstdb2jdbc
>```

Check the status for creating the login modules.
>The following command checks the status of the configureloginmodules
>task in the *ce_install_path*/tools/configure/profiles/wstdb2jdbc
>directory:
>
>```
>configmgr_cl checkstatus -task configureloginmodules
>-profile wstdb2jdbc
>```

Display the help for the checkstatus command.
>The following command displays the help for the **checkstatus** command:
>
>```
>configmgr_cl checkstatus -help
>```

execute command

The **execute** command applies the settings from a configuration XML file for the specified configuration task. The command name is not case sensitive.

execute command syntax

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl execute [-task task_type | -taskfile task_file_name]
 -profile myprofile [-silent] [-force] [-help]
```

execute command parameters

-task *task_type*
>The **-task** *task_type* parameter is indicates which task to run. You can omit the
>**-task** *task_type* parameter if you want to run all of the tasks or if you specify
>the **-taskfile** *task_file_name* parameter. The *task_type* value is not case sensitive.
>The following table lists the valid task names, the associated configuration
>XML file, and a description of the Content Engine settings affected by the task.

Table 30. task_name values

Option	Configuration file to execute	Description
omitted	configurebootstrap.xml	If you omit the **-task** *task_type* parameter and the **-taskfile** *task_file_name* parameter, then all the configuration files in the path are run. Any configuration XML file that has the **enabled** attribute value in the **<configuration>** tag set to false is skipped.
	configurebootstrap.*n*.xml	
	configurejdbcgcd.xml	
	configurejdbcos.xml	
	configurejdbcos.*n*.xml	
	configureldap.xml	
	configureldap.*n*.xml	
	configureloginmodules.xml	
	deployapplication.xml	
	where *n* is an integer starting with 2	

Table 30. task_name values (continued)

Option	Configuration file to execute	Description
configurebootstrap	configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Applies the settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named configurebootstrap.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. For example, you might must run the initial bootstrap task for the new profile and then later run another task for modifying the EAR file to apply a fix pack. If your profile contains more than one configurebootstrap task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to run.
configurejdbcgcd	configurejdbcgcd.xml	Applies the settings for configuring the JDBC connections to the Global Configuration Data (GCD) database used by Content Engine.

Table 30. task_name values (continued)

Option	Configuration file to execute	Description
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Applies the settings for configuring JDBC connections to a single object store database used by Content Engine. You need to generate, edit, and apply the configurejdbcos task settings for each object store in your database. When you generate a second object store JDBC configuration file, it is named configurejdbcos.2.xml. The filename increments for each new file you generate. You can edit the file name as needed. If your profile contains more than one configurejdbcos task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to run.
configureldap	configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Configures the connection to the directory server for authenticating Content Engine users. WebSphere Application Server or Oracle WebLogic Server only.) If you generate a second LDAP configuration file, then it is named configureldap.2.xml. The filename increments for each new file you generate. You can rename the file as needed. If your profile contains more than one configureldap task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to run.
configureloginmodules	configureloginmodules.xml	(WebSphere Application Server or JBoss Application Server only.) Creates the login modules for the Content Engine application.
deployapplication	deployapplication.xml	Deploys the Content Engine application on the application server.

-taskfile *task_file_name*

The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.

If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.

If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*

The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the `ce_install_path`/tools/configure/profiles directory, where `ce_install_path` is the location where the Content Engine software is installed.
- The absolute path to the profile directory, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2"` or `opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2`.
- The absolute path to the profile input file, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"` or `opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg`.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg"`.

-silent

The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-force

The **-force** parameter is optional and only applies when the **-silent** parameter is used. When **-force** is specified, then the task is run without pausing for required responses to validation error messages, such as messages about missing passwords or invalid port numbers.

-help

The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

execute command examples

The following examples include line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

Run all the tasks in a profile.

The following command runs all the tasks in the wstdb2jdbc profile, which

is located in the *ce_install_path*/tools/configure/profiles/wstdb2jdbc
directory. If you include the **-silent** parameter in the command, you must
also include the **-force** parameter.

```
configmgr_cl execute -profile wstdb2jdbc
```

Run the `configurejdbcos` task in a profile with one `configurejdbcos` task.
The following command runs the configurejdbc task in the
ce_install_path/tools/configure/profiles/wstdb2jdbc directory:

```
configmgr_cl execute -task configurejdbcos -profile wstdb2jdbc
```

**Run a single `configurejdbcos` task in a profile with multiple `configurejdbcos`
tasks:** The following command runs the configurejdbcos.2.xml task file in the
ce_install_path/tools/configure/profiles/wstdb2jdbc directory:

```
configmgr_cl execute -taskfile configurejdbc.2.xml -profile wstdb2jdbc
```

Display the help for the execute command.
The following command displays the help for the **execute** command:

```
configmgr_cl execute -help
```

generateconfig command

The **generateconfig** command generates the configuration XML file for the
specified configuration task. The command name is not case sensitive.

generateconfig command syntax

The following syntax includes line breaks to format the command for reading.
Enter the command and options on a single line, without any line breaks.

```
configmgr_cl generateconfig -appserver app_server_type
-repositorytype ldap_repository_type
 -db database_type -ldap ldap_type -bootstrap bootstrap_operation
 -deploy deploy_type -task task_type -taskname display_name
-profile myprofile [-silent] [-force] [-help]
```

generateconfig command parameters

-appserver *appserver_name*
   The **-appserver** *appserver_type* specifies the type of application server and must
   be WebSphere, WebLogic, or JBoss.

-repositorytype *ldap_repository_type*
   WebSphere Application Server only. The **-repositorytype** *ldap_repository_type*
   parameter is required only when you are generating the XML files. This
   parameter specifies the type of LDAP repository to use and must be
   standalone or federated.

-db *database_type*
   The **-db** *database_type* parameter is required only when you are generating all
   the files at the same time or when you are generating a single file by using the
   configurejdbcgcd or configurejdbcos option. This parameter specifies the type
   of database to be used by Content Engine and must be mssql, oracle,
   oracle_rac, db2, or db2zos.

-ldap *ldap_type*
   The **-ldap** *ldap_type* parameter is required only when you are generating all the
   files at the same time or when you are generating a single file by using the
   configureldap option. This parameter specifies the type of directory service
   repository that Content Engine uses for authenticating users and must be
   activedirectory, adam, edirectory, oid, sunjava, or tivoli. The adam option
   applies to both Microsoft ADAM and AD LDS.

-bootstrap *bootstrap_operation*

The **-bootstrap** *bootstrap_operation* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the `configurebootstrap` option. This parameter specifies the bootstrap operation for the profile and must be `new`, `modify`, or `upgrade`.

-deploy *deploy_type*

The **-deploy** *deploy_type* parameter is required only when you are generating all the files at the same time or when you are generating a single file by using the `deployapplication` option. This parameter specifies the type of Content Engine deployment. The value must be `standard`, `cluster`, or `netdeploy` (network deployment).

Specify `standard` if you are deploying Content Engine to a stand-alone (that is, a server that is not managed or clustered) WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server.

Specify `cluster` if you are deploying Content Engine to a WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server cluster.

Specify `netdeploy` if you are deploying Content Engine to a managed WebSphere Application Server instance. That is, you are using Network Deployment to manage individual servers that are not necessarily in a cluster.

-task *task_type*

The **-task** *task_type* parameter indicates which task to generate. You can omit the **-task** *task_type* parameter if you want to generate all the tasks or if you specify the **-taskfile** *task_file_name* parameter. The *task_type* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 31. task_type values

Option	Configuration file to generate	Description
omitted	`applicationserver.xml`	If you omit the `-task` *task_type* parameter and the **-taskfile** *task_file_name* parameter, then all the default configuration files for a profile are created.
	`configurebootstrap.xml`	
	`configurejdbcgcd.xml`	
	`configurejdbcos.xml`	
	`configureldap.xml`	
	`configureloginmodules.xml`	
	`deployapplication.xml`	

Table 31. task_type values (continued)

Option	Configuration file to generate	Description
configurebootstrap	applicationserver.xml configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Generates the file for the application server properties and the file with the settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named configurebootstrap.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. For example, you might must generate the initial bootstrap task for the new profile and then generate a second bootstrap task for modifying the EAR file to apply a fix pack. If the profile already contains an applicationserver.xml file, the existing file is retained.
configurejdbcgcd	applicationserver.xml configurejdbcgcd.xml	Generates the file for the application server properties and the file with the settings for configuring the JDBC connections to the Global Configuration Data (CGD) database used by Content Engine. If the profile already contains an applicationserver.xml file, the existing file is retained.

Table 31. task_type values (continued)

Option	Configuration file to generate	Description
configurejdbcos	applicationserver.xml configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Generates the file for the application server properties and the file with the settings for configuring JDBC connections to a single object store database used by Content Engine. You need to generate, edit, and apply the configurejdbcos task settings for each object store in your database. When you generate a second object store JDBC configuration file, it is named configurejdbcos.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. If the profile already contains an applicationserver.xml file, the existing file is retained.
configureldap	applicationserver.xml configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Generates the file for the application server properties and the file for configuring the connection to the directory server for authenticating Content Engine users. (WebSphere Application Server or Oracle WebLogic Server only.) When you generate a second LDAP configuration file, it is named configureldap.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task. If the profile already contains an applicationserver.xml file, the existing file is retained.

Table 31. task_type values (continued)

Option	Configuration file to generate	Description
configureloginmodules	applicationserver.xml configureloginmodules.xml	(WebSphere Application Server or JBoss Application Server only.) Generates the file for the application server properties and the files for the login modules for the Content Engine application. If the profile already contains an applicationserver.xml file, the existing file is retained.
deployapplication	applicationserver.xml deployapplication.xml	Generates the file for the application server properties and the file with settings for deploying the Content Engine application on the application server. If the profile already contains an applicationserver.xml file, the existing file is retained.

-taskname *display_name*
> The **-taskname** *display_name* parameter is optional and is valid only for generating the files. This parameter specifies the value for the displayName attribute in the configuration XML file. If the display name includes spaces, put the entire name inside quotation marks. The display name is used in the graphical user interface to identify the task.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
> - The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.
>
> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-silent
> The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are

written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-force
> The **-force** parameter is optional and only applies when the **-silent** parameter is used. When **-force** is specified, then the task is run without pausing for required responses to validation error messages, such as messages about missing passwords or invalid port numbers.

-help
> The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

generateconfig command examples

The following examples include line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

Generate all configuration files at the same time.
> The following command generates all the configuration XML files for a new installation profile for a standard deployment on WebSphere with IBM Tivoli Directory Server that uses a stand-alone LDAP repository and DB2 in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:

```
configmgr_cl generateconfig -appserver websphere -repositorytype standalone
 -db db2 -ldap tivoli -bootstrap new -deploy standard -profile wstdb2
```

Generate only the `configurejdbcos` task file for an object store.
> The following command generates only the configurejdbcos.*n*.xml file for a new installation profile for deployment on WebSphere that uses a stand-alone LDAP repository in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:

```
configmgr_cl generateconfig -appserver websphere -repositorytype standalone
 -task configurejdbcos -profile wstdb2jdbc
```

Generate only the `configurejdbcos` task file for an object store and provide a display name for the task.
> The following command generates only the configurejdbcos.*n*.xml file for a new installation profile for deployment on WebSphere that uses a stand-alone LDAP repository in the *ce_install_path*/tools/configure/profiles/wstdb2 directory and uses a display name of Configure Object Store OS23 Data Sources:

```
configmgr_cl generateconfig -appserver websphere -repositorytype standalone
 -task configurejdbcos -taskname "Configure Object Store OS23 Data Sources"
 -profile wstdb2jdbc
```

Display the help for the generateconfig command.
> The following command displays the help for the **generateconfig** command:

```
configmgr_cl generateconfig -help
```

generateupgrade command (with Servers.xml file)

The **generateupgrade** command generates the configuration XML files for upgrading an existing Content Engine. The **generateupgrade** command has two syntax formats, depending on whether you have a valid copy of the Servers.xml file in the directory where Content Engine is installed. The command name is not case sensitive.

generateupgrade command syntax

Use the following syntax for the **generateupgrade** command if you have an existing `Servers.XML` file with valid server information in the directory where Content Engine is installed (`ce_install_path`).

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl generateupgrade  -server server_name -deploy deploy_type
-profile myprofile [-silent] [-force] [-help]
```

generateupgrade command parameters

-server *server_name*
> The **-server** *server_name* parameter indicates the server to create the upgrade profile for. Use the **listservers** command to display a list of valid *server_name* values before you run the **generateupgrade** command.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
> - The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.
>
> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-silent
> The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-help
> The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

generateupgrade command examples

Generate the upgrade configuration files when `Servers.xml` file exists.
> The following command generates all the upgrade configuration XML files for a server named ceserver1 from an existing Content Engine for the profile UpgradeServer1 in the default profile directory:
>
> ```
> configmgr_cl generateupgrade -server ceserver1 -profile UpgradeServer1
> ```

Generate the upgrade configuration files when `Servers.xml` file exists, but does not have the deployment type.

> The following command generates all the upgrade configuration XML files for a server named ceserver1 from an existing Content Engine for the profile UpgradeServer1 in the default profile directory:
>
> ```
> configmgr_cl generateupgrade -server ceserver1 -deploy cluster
> -profile UpgradeServer1
> ```

Display the help for the generateupgrade command.

> The following command displays the help for the **generateupgrade** command:
>
> ```
> configmgr_cl generateupgrade -help
> ```

generateupgrade command (without Servers.xml file)

The **generateupgrade** command generates the configuration XML files for upgrading an existing Content Engine. The **generateupgrade** command has two syntax formats, depending on whether you have a valid copy of the `Servers.xml` file in the directory where Content Engine is installed. The command name is not case sensitive.

generateupgrade command syntax

Use the following syntax for the **generateupgrade** command if you do not have an existing `Servers.XML` file with valid server information in the directory where Content Engine is installed (`ce_install_path`).

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl generateupgrade -appserver app_server_type -server server_name
-deploy deploy_type -profile myprofile [-silent]
[-force] [-help]
```

generateupgrade command parameters

-appserver *appserver_type*

> The **-appserver** *appserver_type* parameter is required only if the `ce_install_path`/Servers.xml file does not exist or does not contain any servers. This parameter specifies the type of application server and must be one of the following values: WebSphere, WebLogic, or JBoss.

-server *server_name*

> The **-server** *server_name* parameter indicates the server to create the upgrade profile for. Use the listservers command to display a list of valid *server_name* values before you run the generateupgrade command.

-deploy *deploy_type*

> The **-deploy** *deploy_type* parameter is required only if the `ce_install_path`/Servers.xml file does not exist or does not contain the deployment type for the selected server. This parameter specifies the type of Content Engine deployment. If a value for the deployment type exists in the Servers.xml file for the selected server, then the value from the Servers.xml file is used instead the value you specify on the command line.
>
> Valid values are standard, cluster, or netdeploy (network deployment).
>
> Specify standard if you are deploying Content Engine to a stand-alone (that is, a server that is not managed or clustered) WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server.

Specify `cluster` if you are deploying Content Engine to a WebSphere Application Server, Oracle WebLogic Server, or JBoss Application Server cluster.

Specify `netdeploy` if you are deploying Content Engine to a managedWebSphere Application Server instance. That is, you are using Network Deployment to manage individual servers that are not necessarily in a cluster.

-profile *myprofile*
>   The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
>
>   - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
>   - The absolute path to the profile directory, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2"` or `opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2`.
>   - The absolute path to the profile input file, such as `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg"` or `opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg`.
>
>   **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter `"C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg"`.

-silent
>   The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-help
>   The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

generateupgrade command examples

Generate the upgrade configuration files when `Servers.xml` file does not exists.
>   The following command generates the upgrade configuration XML files for a server named `ceserver1` from an existing Content Engine cluster deployment that uses the profile `UpgradeServer1` in the default profile directory:
>
>   `configmgr_cl generateupgrade -deploy cluster -profile UpgradeServer1`

Display the help for the generateupgrade command.
>   The following command displays the help for the **generateupgrade** command:
>
>   `configmgr_cl generateupgrade -help`

gui command

The **gui** command opens the Configuration Manager graphical user interface. The command is not case sensitive.

gui command syntax

```
configmgr_cl gui
```

gui command Example

The following command starts the Configuration Manager graphical user interface:

```
 configmgr_cl gui
```

listservers command

The **listservers** command displays a list of the currently deployed Content Engine servers. Only servers that are included in the *ce_install_path*\Servers.xml file are listed, where *ce_install_path* is the location where Content Engine is installed. You must supply the server names when you create an upgrade profile with the **generateUpgrade** command.

listservers command syntax

```
configmgr_cl listservers [-help]
```

listservers command examples

List the deployed Content Engine servers.

The following command lists the deployed Content Engine servers.

```
 configmgr_cl listservers
```

A message like the following message is displayed:

```
These are the currently deployed Content Engine servers:

ce_server01Node01Cell01
    Application Server: WebSphere
    Ear file: C:\Program Files\IBM\FileNet\ContentEngine\lib\deployed\
Engine-ws.ear
    No profile is set.  Use the generateUpgrade command to generate a
profile to upgrade this server.
    Deployment Type: Network
```

Display the help for the listservers command.

The following command displays the help for the **listservers** command:

```
configmgr_cl listservers -help
```

listtasks command

The **listtasks** command displays a list of the tasks and the task files in the configuration profile. The command name is not case sensitive.

listtasks command syntax

```
configmgr_cl listtasks [-task task_type]  -profile myprofile
 [-help]
```

listtasks command parameters

-task *task_type*

The **-task** *task_type* parameter is optional and indicates which task type to list. The *task_type* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 32. task_type values

Option	Configuration files	Description
omitted	configurebootstrap.xml configurebootstrap.*n*.xml configurejdbcgcd.xml configurejdbcos.xml configurejdbcos.*n*.xml configureldap.xml configureldap.*n*.xml configureloginmodules.xml deployapplication.xml Where *n* is an integer starting with 2	If you omit the -task *task_type* parameter, then all the configuration tasks and the associated task files for the profile are listed.
configurebootstrap	configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Lists the task for configuring the settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named configurebootstrap.2.xml. The filename increments for each new file you generate.
configurejdbcgcd	configurejdbcgcd.xml	Lists the task for configuring the settings for the JDBC connections to the Global Configuration Data (CGD) database used by Content Engine.
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Lists the task for configuring the settings for the JDBC connections to a single object store database used by Content Engine. You must generate, edit, and apply the configurejdbcos task settings for each object store in your database. When you generate a second object store JDBC configuration file, it is named configurejdbcos.2.xml. The filename increments for each new file you generate.

Table 32. task_type values (continued)

Option	Configuration files	Description
configureldap	configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Lists the tasks for configuring the connection to the directory server for authenticating Content Engine users. (WebSphere Application Server or Oracle WebLogic Server only) When you generate a second LDAP configuration file, it is named configureldap.2.xml. The filename increments for each new file you generate. When you generate a second LDAP configuration file, it is named configureldap.2.xml. The filename increments for each new file you generate. You cannot change the filename, but you can edit the value in the file for the name of the task.
configureloginmodules	configureloginmodules.xml	(WebSphere Application Server or JBoss Application Server only.) Lists the task for configuring the login modules for the Content Engine application.
deployapplication	deployapplication.xml	Lists the task for deploying the Content Engine application on the application server.

-profile *myprofile*

The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

- The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.

- The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.

- The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-help
> The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

listtasks command examples

List all the configuration tasks in the profile.

> The following command lists all the configuration tasks and the associated configuration XML files in the wstdb2 profile located in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:

```
configmgr_cl listtasks -profile wstdb2
```

> A message like the following example is displayed:

```
All tasks in profile wstdb2.  The tasks will be executed in the
positions indicated when executing all tasks:
Task 1:
        Name: Configure GCD JDBC Data Sources
        Type: configurejdbcgcd
        File: configurejdbcgcd.xml
Task 2:
        Name: Configure Object Store JDBC Data Sources
        Type: configurejdbcos
        File: configurejdbcos.xml
Task 3:
        Name: Configure Login Modules
        Type: configureloginmodules
        File: configureloginmodules.xml
Task 4:
        Name: Configure LDAP
        Type: configureldap
        File: configureldap.xml
Task 5:
        Name: Configure Bootstrap Properties
        Type: configurebootstrap
        File: configurebootstrap.xml
Task 6:
        Name: Deploy Application
        Type: deployapplication
        File: deployapplication.xml
```

List all the `configurejdbcos` tasks in the profile.

> The following command lists all the configure JDBC settings for the object store tasks and the associated configuration XML files in the wstdb2 profile located in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:

```
configmgr_cl listtasks -task configurejdbcos -profile wstdb2
```

> A message like the following example is displayed:

```
Tasks in profile wstdb2 of the task type configurejdbcos:
Task name: Configure Object Store JDBC Data Sources
        File: configurejdbcos.xml
```

Display the help for the listtasks command.
> The following command displays the help for the **listtasks** command:

```
configmgr_cl listtasks -help
```

movetask command

The **movetask** command moves a task to a different position in the list of tasks. The task position determines the order that the tasks are run when you run all the tasks at the same time. You use the **listtask** command to display the task order. The command names are not case sensitive.

movetask command syntax

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl movetask -task task_type | -taskfile task_file_name
-position new_position -profile myprofile [-silent][-force][-help]
```

movetask command parameters

-task *task_type*

The **-task** *task_type* parameter indicates which task to move. This parameter must be included if the **-taskfile** *task_file_name* parameter is omitted. The *task_type* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 33. task_type values

Option	Configuration files	Description
configurebootstrap	configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Moves the task for configuring the settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named configurebootstrap.2.xml. The filename increments for each new file you generate.
configurejdbcgcd	configurejdbcgcd.xml	Moves the task for configuring the settings for the JDBC connections to the Global Configuration Data (CGD) database used by Content Engine.
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Moves the task for configuring the settings for the JDBC connections to a single object store database used by Content Engine. When you generate a second object store JDBC configuration file, it is named configurejdbcos.2.xml. The filename increments for each new file you generate.

Table 33. task_type values (continued)

Option	Configuration files	Description
configureldap	configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Moves the task for configuring the connection to the directory server for authenticating Content Engine users. (WebSphere Application Server or Oracle WebLogic Server only) When you generate a second LDAP configuration file, it is named configureldap.2.xml. The filename increments for each new file you generate. When you generate a second LDAP configuration file, it is named configureldap.2.xml. The filename increments for each new file you generate.
configureloginmodules	configureloginmodules.xml	Moves the task for configuring the login modules for the Content Engine application.
deployapplication	deployapplication.xml	Moves the task for deploying the Content Engine application on the application server.

-taskfile *task_file_name*
> The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.
>
> If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.
>
> If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-position*new_position*
> The **-position***new_position* parameter specifies the new position in the list for the item. You can run the **listtasks** to view the list of tasks in the profile and their position before you run the **movetask** command.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> - The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> - The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.

- The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-silent
> The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-force
> The **-force** parameter is optional and only applies when the **-silent** parameter is used. When **-force** is specified, then the task is run without pausing for required responses to validation error messages, such as messages about missing passwords or invalid port numbers.

-help
> The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

movetask command examples

Move the task for the `configurejdbcos.3.xml` file in a profile with more than one `configurejdbcos` task.
> The following command moves the task for the configurejdbcos.3.xml file for the wstdb2 profile located in the *ce_install_path*/tools/configure/ profiles/wstdb2 directory to position 1:
>
> ```
> configmgr_cl movetask -taskfile configurejdbcos.3.xml -position 1
> -profile wstdb2
> ```

Move the `configureloginmodules` task.
> The following command moves the configureloginmodules task from the default position of task 3 to task 5 in the wstdb2 profile located in the *ce_install_path*/tools/configure/profiles/wstdb2 directory:
>
> ```
> configmgr_cl movetask -task configureoginmodules -position 5
> -profile wstdb2
> ```

Display the help for the movetask command.
> The following command displays the help for the **movetask** command:
>
> ```
> configmgr_cl movetask -help
> ```

removetask command

The **removetask** command removes the specified task from the configuration profile. When you remove the task, the configuration XML file is deleted from the profile directory. The command name is not case sensitive.

removetask command syntax

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl removetask -task task_type | -taskfile task_file_name
 -profile myprofile [-silent][-force][-help]
```

removetask command parameters

-task *task_type*

The **-task** *task_type* parameter is optional and indicates which task to remove. You can omit the **-task** *task_type* parameter if you specify the **-taskfile** *task_file_name* parameter. The *task_type* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 34. task_type values

Option	Configuration file to remove	Description
configurebootstrap	configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Settings for the Content Engine bootstrap on the application server. The bootstrap information is needed for creating the Global Configuration Data and for starting Content Engine When you generate a second bootstrap configuration file, it is named configurebootstrap.2.xml. The filename increments for each new file you generate.
configurejdbcgcd	configurejdbcgcd.xml	Settings for configuring the JDBC connections to the Global Configuration Data (CGD) database used by Content Engine.
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Settings for configuring JDBC connections to a single object store database used by Content Engine. You need to generate, edit, and apply the configurejdbcos task settings for each object store in your database. If you generate a second object store JDBC configuration file, then it is named configurejdbcos.2.xml. The filename increments for each new file you generate.

Table 34. task_type values (continued)

Option	Configuration file to remove	Description
configureldap	configureldap.xml configureldap.*n*.xml, where *n* is an integer starting with 2	Configures the connection to the directory server for authenticating Content Engine users. (WebSphere Application Server or Oracle WebLogic Server only) If you generate a second LDAP configuration file, then it is named configureldap.2.xml. The filename increments for each new file you generate. If your profile contains more than one configureldap task, then you must specify the **-taskfile** *task_file_name* parameter to identify the task to remove.
configureloginmodules	configureloginmodules.xml	(WebSphere Application Server or JBoss Application Server only) Creates the login modules for the Content Engine application.
deployapplication	deployapplication.xml	Deploys the Content Engine application on the application server.

-taskfile *task_file_name*
> The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.
>
> If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.
>
> If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:
> * The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.
> * The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.
> * The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

Remember: If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

-silent
The **-silent** parameter is optional. When **-silent** is specified, then no prompts or informational messages are displayed in the console, but the errors are written to the log. Failure messages and validation error messages are displayed as needed, such as messages about missing passwords or invalid port numbers. If you run the execute command to run all the tasks in a profile, and you specify the **-silent** parameter, you must also specify the **-force** parameter.

-force
The **-force** parameter is optional and only applies when the **-silent** parameter is used. When **-force** is specified, then the task is run without pausing for required responses to validation error messages, such as messages about missing passwords or invalid port numbers.

-help
The **-help** parameter is optional and displays a brief message on the command syntax instead of running the command.

removetask command examples

Remove the `configurejdbcos` task from a profile with only one `configurejdbcos` task.   The following command removes the `configurejdbcos` task and the `configurejdbcos.xml` file from the profile named wstdb2jdbc_one.

```
configmgr_cl removetask -task configurejdbcos -profile wstdb2jdbc_one
```

Remove the `configurejdbcos` task from a profile with several `configurejdbcos` tasks.   The following command removes the `configurejdbcos` task and the `configurejdbcos.2.xml` file from the profile named wstdb2jdbc_many.

```
configmgr_cl removetask -taskfile configurejdbcos.2.xml
 -profile wstdb2jdbc_many
```

Remove the `configurejdbcos` task from a profile with several `configurejdbcos` tasks by using an absolute path to the profile directory.
The following command removes the `configurejdbcos` task and the `configurejdbcos.2.xml` file from the profile named wstdb2jdbc_many that is located in the `c:\temp\myprofiles\wstdb2jdbc_many` directory.

```
configmgr_cl removetask -taskfile configurejdbcos.2.xml
-profile  c:\temp\myprofiles\wstdb2jdbc_many
```

Display the help for the removetask command.
The following command displays the help for the **removetask** command:

```
configmgr_cl removetask -help
```

storepasswords command

The **storepasswords** command prompts for passwords that are blank in a profile and stores the encrypted passwords in the file. Storing encrypted passwords might not be FIPS 140-2 compliant. You can run the command for a single task or for all tasks in the profile. The command name is not case sensitive.

storepasswords command syntax

The following syntax includes line breaks to format the command for reading. Enter the command and options on a single line, without any line breaks.

```
configmgr_cl storepasswords [-task task_type | -taskfile task_file_name]
  -profile myprofile [-help]
```

storepasswords command parameters

-task *task_type*

The **-task** *task_type* parameter specifies a specific task to encrypt passwords for. You can omit the **-task** *task_type* parameter if you want to store passwords for all the tasks or if you specify the **-taskfile** *task_file_name* parameter. The *task_name* value is not case sensitive. The following table lists the valid task names, the associated configuration XML file, and a description of the Content Engine settings affected by the task.

Table 35. task_type values

Option	Configuration file	Description
omitted	applicationserver.xml configurebootstrap.xml configurebootstrap.*n*.xml configurejdbcgcd.xml configurejdbcos.xml configurejdbcos.*n*.xml configureldap.xml configureldap.*n*.xml deployapplication.xml Where *n* is an integer starting with 2	If you omit the **-task** *task_type* parameter and the **-taskfile** *task_file_name* parameter, then you are prompted to enter the passwords for each configuration XML file in the profile. Each password is encrypted before it is added to the XML file.
configurebootstrap	configurebootstrap.xml configurebootstrap.*n*.xml, where *n* is an integer starting with 2	Encrypts the password for the BootstrapPassword property that is used to create the Global Configuration Data and to start Content Engine. If you have more than one configurebootstrap.*n*.xml file, run the command for each task file.
configurejdbcgcd	configurejdbcgcd.xml	Encrypts the password for the DatabasePassword property that Content Engine uses to access the GCD database.

Table 35. task_type values (continued)

Option	Configuration file	Description
configurejdbcos	configurejdbcos.xml configurejdbcos.*n*.xml, where *n* is an integer starting with 2	Encrypts the password for the `DatabasePassword` property that Content Engine uses to access the object store database. If you have more than one configurejdbcos.*n*.xml file, run the command for each task file.

-taskfile *task_file_name*
> The **-taskfile** *task_file_name* parameter specifies the configuration XML file to use.

> If only one task file exists for the *task_type*, then the **-taskfile** *task_file_name* parameter is optional.

> If more than one task file for the *task_type* exists, then you must include the **-taskfile** *task_file_name* parameter. You can omit the **-task** *task_type* parameter when you specify the **-taskfile** *task_file_name* parameter.

-profile *myprofile*
> The **-profile** *myprofile* parameter specifies the profile to use. The *myprofile* value can be one of the following items:

> • The name of the profile, such as ce_was_tiv_db2. The profile must be located in the *ce_install_path*/tools/configure/profiles directory, where *ce_install_path* is the location where the Content Engine software is installed.

> • The absolute path to the profile directory, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2" or opt/IBM/FileNet/ContentEngine/tools/configure/ profiles/ce_was_tiv_db2.

> • The absolute path to the profile input file, such as "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ ce_was_tiv_db2\ce_was_tiv_db2.cfg" or opt/IBM/FileNet/ContentEngine/ tools/configure/profiles/ce_was_tiv_db2/ce_was_tiv_db2.cfg.

> **Remember:** If the path includes a directory name with spaces, enclose the entire path in quotation marks. For example, enter "C:\Program Files\IBM\FileNet\ContentEngine\tools\configure\profiles\ce_was_tiv_db2\ ce_was_tiv_db2.cfg".

storepasswords command examples

Encrypt and save all passwords for a profile.
> The following command encrypts and saves passwords for any blank entries in the profile named wstdb2jdbc_one.
> ```
> configmgr_cl storepasswords -profile wstdb2jdbc_one
> ```

Encrypt and save the passwords for a specific configurejdbcos task in a profile with several configurejdbcos tasks.
> The following command encrypts and saves passwords for any blank entries in the configurejdbcos.2.xml file from the profile named wstdb2jdbc_many.

```
configmgr_cl storepasswords -taskfile configurejdbcos.2.xml
-profile wstdb2jdbc_many
```

Display the help for the storepasswords command.

The following command displays the help for the **storepasswords** command:

```
configmgr_cl storepasswords -help
```

Appendix B. Troubleshooting installation and upgrade

This section provides troubleshooting solutions for problems you might encounter when installing or upgrading IBM FileNet P8 software.

Notices

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

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

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

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

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

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

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

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

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

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

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.

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

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

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

COPYRIGHT LICENSE:

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

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

Trademarks

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

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

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

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

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

Index

IBM ®

Program Number: 5724-R76
                 5724-R81