


## Installation and Upgrade Guide





## Installation and Upgrade Guide

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 677.

This edition applies to version 4.0.1 of IBM FileNet Content Manager (product number 5724-R81), version 4.0.2 of IBM FileNet Business Process Manager (product number 5724-R76), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2001, 2008. All rights reserved.**  
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Typographical Conventions

This document uses the conventions in the following table to distinguish elements of text.

Convention	Usage
UPPERCASE	Environment variables, status codes, utility names.
<b>Bold</b>	Program names and selected terms such as command parameters or environment variables that require emphasis.
<b>Bold Gray</b>	Clickable user-interface elements (such as buttons).
<b>Bold Olive</b>	Paths and file names.
<i>Italic</i>	User-supplied variables and new terms introduced in text, names of additional documents (such as <i>IBM FileNet® P8 Platform Installation and Upgrade Guide</i> ).
<italic>	User-supplied variables that replace everything between and including the angle bracket delimiters (< and >).
Monospace	Code samples, examples, display text, and error messages.

**NOTE** Some path names in this document that are identical (except for the directory-separator character) on both UNIX® and Windows® platforms are specified in UNIX syntax only (that is, with forward-slash directory separators).

**WARNING** This document contains examples of text to be typed on a command line. Be sure to manually type the command, rather than copying and pasting it from this document. Otherwise, your command line may contain unrecognized characters and will not execute properly.

## Revision Log

The following table identifies changes made to this document since the IBM FileNet P8 4.0.0 release.

Date	Revision
05/08	Corrected executable name to run Process Engine Setup on Windows in <a href="#">“Install Process Engine (Windows)”</a> on page 266.
03/08	In <a href="#">“Installation Planning Considerations”</a> on page 24, added a statement in the “Database Considerations” subtopic about the effect of a database upgrade on the application server.
	Corrected the caution note in step 2 of the “To specify the WebSphere environment variables” procedure in <a href="#">“Configure an Application Server for Content Engine (WebSphere)”</a> on page 106.
	Fixed errors in procedure for encrypting LDAP password in <a href="#">“Configure an Application Server for Content Engine (JBoss)”</a> on page 118.
	Added a step to configure Autonomy K2 Dashboard to use SSL in <a href="#">“To configure Autonomy K2 for Content-Based Retrieval”</a> on page 153.
	<p>In <a href="#">“Install and Deploy Content Engine”</a> on page 160:</p> <ul style="list-style-type: none"> <li>• Revised descriptions of the following Content Engine Setup screens: <ul style="list-style-type: none"> <li>– Choose Application Server</li> <li>– Review Installation Summary</li> <li>– Configure JDBC</li> <li>– Specify Authentication Provider</li> <li>– Verify Installation</li> </ul> </li> <li>• Clarified steps 2 and 3 in the “To enable the Log4j logging API (optional)” procedure.</li> </ul>
	<p>In <a href="#">“Configure Content Engine Application Server Database Connectivity (JBoss 4.0.x)”</a> on page 247, made the following revisions:</p> <ul style="list-style-type: none"> <li>• Fixed XML syntax errors.</li> <li>• Fixed errors in step 1 in the following procedures: <ul style="list-style-type: none"> <li>– “To configure JBoss 4.0.x database connectivity (MS SQL Server)”</li> <li>– “To configure JBoss 4.0.x database connectivity (DB2)”</li> </ul> </li> <li>• Clarified step 2 in the “To encrypt data source passwords” procedure.</li> </ul>
<p>In <a href="#">“Configure Application Engine (JBoss)”</a> on page 381, added quotation marks to the “%JAVA%” %JAVA_OPTS% command example for Windows to accommodate the word spacing in the “Program Files” directory.</p>	

Date	Revision
	<p>Corrected batch file names for creating EAR and WAR files in <a href="#">“Deploy Application Engine (WebLogic)” on page 391</a>.</p> <p>Removed requirement that an Oracle upgrade can occur only after upgrading IBM FileNet P8 components in <a href="#">“Upgrade Planning Considerations” on page 461</a>.</p> <p>In <a href="#">“Upgrade Planning Considerations” on page 461</a>, added a statement in the “Database Considerations” subtopic about the effect of a database upgrade on the application server.</p> <p>In <a href="#">“Upgrade Content Engine Data” on page 498</a>, removed step 3 from the “To create the XML upgrade status file” procedure, as it duplicates steps 2 through 5 of the “To check Upgrader Tool prerequisites” procedure.</p> <p>Clarified the Application Engine install-path location for an upgrade in <a href="#">“Upgrade Application Engine” on page 551</a>.</p>
12/07	Clarified configuration requirements for supported J2EE application servers in <a href="#">“J2EE Application Server Considerations” on page 36</a> .
11/07	Updated links to online documentation for the new IBM FileNet FTP site.
09/07	<p>Added notes and topics throughout the guide for new support of JBoss application servers for Content Engine, Application Engine and FileNet P8 documentation servers. For example, added or updated the following topics:</p> <ul style="list-style-type: none"> <li>• <a href="#">“Install IBM FileNet P8 Platform Documentation (JBoss)” on page 141</a></li> <li>• <a href="#">“Upgrade IBM FileNet P8 Documentation” on page 478</a></li> <li>• <a href="#">“Configure an Application Server for Content Engine (JBoss)” on page 118</a></li> <li>• <a href="#">“Configure an Application Server for Application Engine (JBoss)” on page 126</a></li> <li>• <a href="#">“Configure Application Engine (JBoss)” on page 381</a></li> <li>• <a href="#">“Deploy Application Engine (JBoss)” on page 394</a></li> </ul> <p>Added notes and topics throughout the guide for new support of Microsoft® Active Directory Application Mode (ADAM) directory service.</p> <p>Added notes throughout the guide for new support of Microsoft SQL Server 2005 databases.</p> <p>Added notes throughout the guide for new support of Oracle 10g.</p> <p>Added information in <a href="#">“Install Process Engine (HP-UX)” on page 320</a> and <a href="#">“Install Process Engine (Solaris)” on page 285</a> to support DB2® databases.</p> <p>Added notes throughout the guide for new support of IBM® WebSphere® 6.1 application servers for Content Engine and Application Engine.</p> <p>Modified the topic <a href="#">“Create Additional File Storage Areas” on page 428</a>, including its title, to change the focus to the second and all subsequent file storage areas.</p>

Date	Revision
	Modified the title and approach of the topic <a href="#">“Deploy Multiple Content Engine Instances” on page 436</a> . It now clarifies that each Content Engine instance needs its own application server configuration.
	Added new topic <a href="#">“Complete Post-Upgrade Content Engine Configuration” on page 513</a> to cover final steps required for upgrading and recreating Content Engine data, including content search collections.
	Made numerous changes to instructions related to upgrading Content Search Engine. For example, see the new <a href="#">“Complete Content Search Engine Upgrade and Create Collections”</a> procedure in the topic <a href="#">“Complete Post-Upgrade Content Engine Configuration” on page 513</a> .
	Applied significant updates and improvements to the Application Engine installation, configuration, and upgrade topics.
	In the topic <a href="#">“Complete Post-Upgrade Process Engine Configuration” on page 539</a> , removed steps to clean up database views. Process Engine post-upgrade steps generally simplified and the process made more user friendly.
	Removed the silent installation and upgrade sample response files from the guide and renamed the associated appendix topic to be <a href="#">“Encrypt Passwords for Silent Installations and Upgrades” on page 642</a> . This was done because the actual sample response files are included with the software distributions.
06/07	Replaced sections <a href="#">“Installation Roadmaps”</a> and <a href="#">“Upgrade Roadmaps”</a> with multiple topics to clarify process for installation, upgrade, and software updates.
	Added silent install and uninstall information to the topics <a href="#">“Install Application Integration” on page 429</a> and <a href="#">“Install File Tracker” on page 433</a> .
	In the <a href="#">“Security Considerations”</a> section of the topic <a href="#">“Upgrade Planning Considerations” on page 461</a> , added a section describing the upgrade of 3.5.x authentication parameters.
	In the <a href="#">“Security Considerations”</a> section of the topic <a href="#">“Upgrade Planning Considerations” on page 461</a> , added information about reconciling the Process Engine user security information.



Date	Revision
03/07	<p data-bbox="399 352 1268 415">Numerous topics added or updated to cover new IBM DB2 database certification:</p> <ul data-bbox="399 436 1263 1119" style="list-style-type: none"> <li data-bbox="399 436 1027 468">• <a href="#">“Installation Planning Considerations” on page 24</a></li> <li data-bbox="399 489 873 520">• <a href="#">“Sample Configurations” on page 37</a></li> <li data-bbox="399 541 862 573">• <a href="#">“Installation Checklists” on page 42</a></li> <li data-bbox="399 594 995 625">• <a href="#">“Specify IBM FileNet P8 Accounts” on page 60</a></li> <li data-bbox="399 646 1227 678">• <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a></li> <li data-bbox="399 699 1232 730">• <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a></li> <li data-bbox="399 751 1016 783">• <a href="#">“Install Process Engine (Windows)” on page 266</a></li> <li data-bbox="399 804 951 835">• <a href="#">“Install Process Engine (AIX)” on page 304</a></li> <li data-bbox="399 856 1011 888">• <a href="#">“Configure Process Task Manager” on page 340</a></li> <li data-bbox="399 909 1003 940">• <a href="#">“Upgrade Process Engine (UNIX)” on page 517</a></li> <li data-bbox="399 961 1049 993">• <a href="#">“Upgrade Process Engine (Windows)” on page 527</a></li> <li data-bbox="399 1014 1260 1045">• <a href="#">“Complete Post-Upgrade Process Engine Configuration” on page 539</a></li> <li data-bbox="399 1066 1102 1098">• <a href="#">“IBM FileNet P8 Database Character Sets” on page 653</a></li> <li data-bbox="399 1119 971 1150">• <a href="#">“IBM FileNet P8 Port Numbers” on page 643</a></li> </ul> <p data-bbox="399 1150 1398 1245">In <a href="#">“Specify IBM FileNet P8 Accounts” on page 60</a>, added links to the topics where specific users or groups are referenced, and added information about whether the user or group is different from 3.5.x.</p> <p data-bbox="399 1266 1369 1329">Added installation and upgrade planning topics <a href="#">“Installation Roadmaps”</a> and <a href="#">“Upgrade Roadmaps”</a>.</p> <p data-bbox="399 1350 1414 1478">In the topic <a href="#">“Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81</a>, changed a reference to direct users to the FileNet P8 Hardware and Software Requirements document for the JDBC version information, replacing a link to a Microsoft website.</p>

Date	Revision
	<p>The topic “Create an Object Store and Verify the Content Engine Installation” has been split into the following topics:</p> <ul style="list-style-type: none"> <li>• <a href="#">“Configure Content Engine Application Server Database Connectivity (WebSphere 5.1.x)”</a> on page 193</li> <li>• <a href="#">“Configure Content Engine Application Server Database Connectivity (WebSphere 6.0.x)”</a> on page 206</li> <li>• <a href="#">“Configure Content Engine Application Server Database Connectivity (WebLogic 8.1.x)”</a> on page 236</li> <li>• <a href="#">“Configure Content Engine Application Server Database Connectivity (WebLogic 9.2.x)”</a> on page 241</li> <li>• <a href="#">“Create Object Stores”</a> on page 258</li> <li>• <a href="#">“Verify the Content Engine Installation”</a> on page 263</li> </ul> <p>In the topic <a href="#">“Complete Post-Install Process Engine Configuration”</a> on page 342, added a step to restart Process Engine after creating a new key in the registry.</p> <p>In the topics <a href="#">“Upgrade Process Engine (UNIX)”</a> on page 517 and <a href="#">“Upgrade Process Engine (Windows)”</a> on page 527, added a note to install the 4.0.1 Service Pack before proceeding to post-upgrade steps.</p> <p>In the topic <a href="#">“Upgrade Process Engine (UNIX)”</a> on page 517, added information on editing the ims_start file if on HP with memory greater than 2GB.</p> <p>In the topic <a href="#">“Complete Post-Upgrade Process Engine Configuration”</a> on page 539, added information on view cleanup, a step that is only necessary if stopping at the FCS release. There is no need to drop and recreate database views if proceeding straight to the GA release.</p>
12/06	Initial release of the <i>FileNet P8 Platform 4.0.0 Installation and Upgrade Guide</i> .

# Table of Contents

<b>Typographical Conventions</b> .....	<b>5</b>
<b>Revision Log</b> .....	<b>6</b>
<b>About this Document.</b> .....	<b>20</b>
Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs. ....	21
<b>Installation Planning and Procedures</b> .....	<b>22</b>
Plan the Installation .....	23
Installation Planning Considerations .....	24
General Requirements for all IBM FileNet P8 Systems .....	24
Operating System Considerations .....	25
General .....	25
Content Engine .....	26
Process Engine .....	26
Application Engine .....	28
Network Considerations .....	28
Security Considerations .....	28
Proxy server or firewall configurations .....	28
Authentication and Authorization .....	28
Microsoft Windows Active Directory .....	30
Microsoft Active Directory Application Mode (ADAM) .....	30
Novell eDirectory .....	30
Sun Java System Directory Server .....	31
IBM Tivoli Directory Server .....	32
Database Considerations .....	32
All platforms .....	32
Microsoft SQL Server .....	32
Oracle .....	33
DB2 .....	35
J2EE Application Server Considerations .....	36
Sample Configurations .....	37
Baseline Configuration .....	38
Baseline Configuration With Optional Components .....	39
Developer Configuration .....	40
Demo Configuration .....	41
Installation Checklists .....	42
IT Administrator .....	42
P8 Platform Environment .....	42
Content Engine .....	42
Process Engine .....	42
Application Engine .....	43

Security Administrator	43
User Accounts	43
Content Engine	44
Application Engine	45
Database Administrator	46
Content Engine	46
Process Engine	47
Installation Administrator	49
Environment Considerations	50
Content Engine	50
Content Search Engine	51
Process Engine	51
Application Engine	51
<b>Prerequisite Tasks</b>	<b>53</b>
Task 1a: Configure Windows Active Directory	54
Task 1b: Configure Active Directory Application Mode (ADAM)	55
Task 1c: Configure Sun Java System Directory Server	56
Configure Sun Java System Directory Server (v 5.1 SP2)	56
Configure Sun Java System Directory Server (v 5.2)	57
Task 1d: Configure Novell eDirectory	58
Task 1e: Configure IBM Tivoli Directory Server	59
Task 2: Specify IBM FileNet P8 Accounts	60
Task 3a: Verify that Microsoft SQL Server Is Installed for IBM FileNet P8	81
Task 3b: Verify that Oracle Server Is Installed for IBM FileNet P8	85
Task 3c: Verify that DB2 Server Is Installed for IBM FileNet P8	92
Task 4a: Verify that Microsoft SQL Server Client Is Installed for IBM FileNet P8	98
Task 4b: Verify that Oracle Client Is Installed for IBM FileNet P8	100
Task 4c: Verify that DB2 Client Is Installed for IBM FileNet P8	103
Task 5a: Configure an Application Server for Content Engine (WebSphere)	106
Task 5b: Configure an Application Server for Content Engine (WebLogic)	110
Task 5c: Configure an Application Server for Content Engine (JBoss)	118
Manually Configure JBoss Data Sources for the GCD (Optional)	119
Encrypt LDAP Password (Optional)	121
Encrypt GCD Data Source Passwords (Optional)	122
Task 6a: Configure an Application Server for Application Engine (WebSphere)	123
Task 6b: Configure an Application Server for Application Engine (WebLogic)	124
Task 6c: Configure an Application Server for Application Engine (JBoss)	126
<b>Installation Tasks</b>	<b>128</b>
Task 7a: Install IBM FileNet P8 Platform Documentation (WebSphere)	130
Overview of Procedures	130
Install IBM FileNet P8 Platform Documentation	131
Install Expansion Product Documentation	131
Update Documentation Search Index	133
Complete and Verify the Documentation Installation	134
Task 7b: Install IBM FileNet P8 Platform Documentation (WebLogic)	135
Overview of Procedures	135
Install the IBM FileNet P8 Platform Documentation	136

Install Expansion Product Documentation . . . . .	137
Update the Documentation Search Index. . . . .	138
Complete and Verify the Documentation Installation . . . . .	140
Task 7c: Install IBM FileNet P8 Platform Documentation (JBoss). . . . .	141
Overview of procedures . . . . .	141
Install the IBM FileNet P8 Platform documentation . . . . .	142
Install functional expansion documentation . . . . .	142
Update the documentation search index . . . . .	143
Complete and verify the documentation installation. . . . .	145
Task 8: Install and Configure Content Search Engine . . . . .	147
Task 9: Install and Deploy Content Engine . . . . .	160
Task 10: Install Content Engine Software Updates . . . . .	177
Task 11: Complete Post-Install Content Engine Configuration . . . . .	178
Install Centera Shared Libraries . . . . .	178
Verify the Data Sources and Connection Pools . . . . .	180
Verify That Content Engine Has Deployed. . . . .	181
Enable Log4J Logging API (Optional) . . . . .	182
Task 12: Install Enterprise Manager . . . . .	183
Task 13: Establish the FileNet P8 Domain and Global Configuration Data (GCD). . . . .	187
Task 14a: Configure Content Engine Application Server Database Connectivity (WebSphere 5.1.x) . . . . .	193
Task 14b: Configure Content Engine Application Server Database Connectivity (WebSphere 6.0.x) . . . . .	206
Task 14c: Configure Content Engine Application Server Database Connectivity (WebSphere 6.1.x) . . . . .	220
Task 14d: Configure Content Engine Application Server Database Connectivity (WebLogic 8.1.x) . . . . .	236
Task 14e: Configure Content Engine Application Server Database Connectivity (WebLogic 9.2.x) . . . . .	241
Task 14f: Configure Content Engine Application Server Database Connectivity (JBoss 4.0.x) . . . . .	247
Task 15: Prepare Storage Areas for Object Stores . . . . .	252
Configure File Servers for File Storage Areas . . . . .	253
Remote File Access Protocols . . . . .	253
Users and Groups . . . . .	254
Configure the Remote Access Protocol on the Client Machine . . . . .	256
Task 16: Create Object Stores . . . . .	258
Task 17: Verify the Content Engine Installation . . . . .	263
Task 18: Install Content Search Engine Software Updates. . . . .	265
Task 19a: Install Process Engine (Windows) . . . . .	266
To Complete Additional Configuration . . . . .	281
Task 19b: Install Process Engine (Solaris) . . . . .	285
Task 19c: Install Process Engine (AIX) . . . . .	304
Task 19d: Install Process Engine (HP-UX) . . . . .	320
Task 20: Install Process Engine Software Updates . . . . .	338
Task 21: Install the Latest Content Engine Client Files on Process Engine Servers . . . . .	339
Task 22: Configure Process Task Manager . . . . .	340
Task 23: Complete Post-Install Process Engine Configuration. . . . .	342
Task 24: Install Application Engine . . . . .	344
Task 25a: Configure Application Engine (WebSphere) . . . . .	356
Task 25b: Configure Application Engine (WebLogic) . . . . .	373
Task 25c: Configure Application Engine (JBoss) . . . . .	381
Task 26: Install Application Engine Software Updates. . . . .	383
Task 27: Install the Latest Content Engine and Process Engine Client Files on Application Engine Servers . . . . .	384
Task 28a: Deploy Application Engine (WebSphere) . . . . .	385

Task 28b: Deploy Application Engine (WebLogic) .....	391
Task 28c: Deploy Application Engine (JBoss) .....	394
<b>Configuration/Startup Tasks .....</b>	<b>396</b>
Task 29: Configure Content Engine for Content-Based Retrieval .....	397
Task 30: Set Application Engine Bootstrap Preferences .....	403
Bootstrap Preferences .....	403
Enhanced Timezone Detection .....	403
Task 31: Create a Process Engine Isolated Region .....	408
Task 32: Create a Process Engine Connection Point .....	409
Task 33: Configure the Process Engine Connection Point for Application Engine .....	410
Task 34: Set Up Content Engine and Client Transport SSL Security .....	412
Task 35: Set Up Application Engine SSL Security .....	416
Additional procedure for WebSphere 5.1 .....	419
Using Java Applets in an SSL Environment .....	419
Task 36: Perform Additional Configuration Tasks .....	420
<b>Optional Installation Tasks .....</b>	<b>422</b>
Task 37: Install and Configure IBM FileNet Publishing .....	423
Task 38: Enable the Process Engine Component Integrator .....	424
Task 39: Install an Additional Instance of Enterprise Manager .....	427
Task 40: Create Additional File Storage Areas .....	428
Task 41: Install Application Integration .....	429
Task 42: Install File Tracker .....	433
Task 43: Deploy Multiple Content Engine Instances .....	436
Task 44: Deploy Content Engine to Other Application Servers .....	437
Deploy Content Engine Into a Non-Managed Server Environment .....	437
Deploy Content Engine Into a Managed Server Environment .....	439
Task 45: Deploy Multiple Application Engine Instances .....	443
Task 46: Install Additional Content Search Engine Servers .....	447
Task 47: Enable Application Engine to Use ISRA .....	451
ISRA SSL support .....	451
To install and deploy the Application Engine ISRA Servlet .....	452
To configure the Workplace site preferences .....	454
To set the ISRA Interface Servlet URL .....	454
To log on to Image Services via LDAP .....	455
To access IS library documents .....	455
Task 48: Install and Configure IBM FileNet System Manager .....	456
Task 49: Install Software Updates for Optional Components .....	457
<b>Upgrade Planning and Procedures .....</b>	<b>458</b>
<b>Plan the Upgrade .....</b>	<b>459</b>
Upgrade Overview .....	460
Upgrade Planning Considerations .....	461
General Requirements for all IBM FileNet P8 Systems .....	461
Expansion Product Considerations .....	462
eForms .....	462

Records Manager	463
Third-party Software Considerations	463
Fixed Content Devices	463
Operating System Considerations	463
General	463
Content Engine	464
Process Engine	465
Security Considerations	466
Network Considerations	467
Database Considerations	468
General	468
Oracle	468
Microsoft SQL Server	469
Application Server Considerations	469
Upgrade Checklists	470
IT Administrator	470
Before the Upgrade	470
Content Engine	470
Process Engine	471
Application Engine	471
Security Administrator	472
General	472
Content Engine	472
Application Engine	473
Database Administrator	473
General	474
Content Engine	474
Process Engine	474
Upgrade Administrator	474
Before the Upgrade	474
Environment Considerations	475
Content Engine	475
Process Engine	475
Application Engine	476
Upgrade Core Components	477
Task 1: Upgrade IBM FileNet P8 Documentation	478
Overview	478
Refresh IBM FileNet P8 4.0.0 Documentation Without Uninstalling	479
Update IBM FileNet P8 4.0.0 Documentation by Uninstalling and Reinstalling	481
Update Help Search Index	484
Deploy and Verify IBM FileNet P8 Documentation Web Site	485
Task 2: Verify Current Release Level	486
Task 3: Upgrade Content Engine Software	487
Task 4: Install Content Engine Software Updates	497
Task 5: Upgrade Content Engine Data	498
Graphical User Interface to Upgrader Tool	500
Prepare Items for Upgrade	503
Perform the Upgrade	506

Command Line Interface to Upgrader Tool . . . . .	507
Task 6: Install Content Search Engine Software Updates. . . . .	512
Task 7: Complete Post-Upgrade Content Engine Configuration . . . . .	513
Check Completion Status of Asynchronous Upgrade Events . . . . .	513
Create CodeModules Folders for Upgraded Object Stores . . . . .	513
Clear the Read-Only Attribute for NTFS File Storage Areas . . . . .	514
Move File Storage Areas from Windows to UNIX. . . . .	514
Uninstall Version 3.5.x of Content Engine . . . . .	515
Complete Content Search Engine Upgrade and Create Collections . . . . .	516
Task 8a: Upgrade Process Engine (UNIX). . . . .	517
Task 8b: Upgrade Process Engine (Windows) . . . . .	527
Task 9: Install Process Engine Software Updates . . . . .	537
Task 10: Install the Latest Content Engine Client Files on Process Engine Servers . . . . .	538
Task 11: Complete Post-Upgrade Process Engine Configuration . . . . .	539
Task 12: Upgrade Application Engine . . . . .	551
Task 13: Complete Post-Upgrade Application Engine Configuration. . . . .	556
Upgrade Add-On Components . . . . .	558
Task 14: Upgrade Enterprise Manager (Standalone) . . . . .	559
Task 15: Upgrade Application Integration and File Tracker. . . . .	560
Upgrade considerations . . . . .	560
Task 16: Upgrade IBM FileNet Publishing Components. . . . .	562
Task 17: Upgrade Custom Applications . . . . .	563
Content Engine COM API. . . . .	563
New Content Engine 4.0 Object Types . . . . .	563
COM API Object Types Not Supported . . . . .	564
New Content Engine 4.0 Properties on Supported Object Types . . . . .	564
COM API Properties Not Supported . . . . .	564
Content Engine 4.0 Methods Not Supported . . . . .	565
Validation of Property Values . . . . .	566
Default Method Short-Cutting . . . . .	566
VB Error Handling . . . . .	567
Property Side-effects . . . . .	568
Managed Client Applications Using Catclient . . . . .	568
Error Reporting . . . . .	568
Security . . . . .	568
OLEDB/ADO Data Provider . . . . .	570
Microsoft IIS (Internet Information Services) WebDAV Provider . . . . .	570
Publishing . . . . .	570
Events and Subscriptions . . . . .	570
Query . . . . .	572
Storage . . . . .	573
Auditing . . . . .	574
Document Lifecycles . . . . .	574
Document Classification . . . . .	575
Content Type . . . . .	576
VersionSeries Security . . . . .	576
Addons . . . . .	576
Exception When Attempting to Modify Root Folder Properties . . . . .	576



Exception When Attempting to Update the List of PropertyDefinition Objects for a ClassDefinition . . . . .	576
GUID Strings . . . . .	576
Reusing Dependent Objects . . . . .	577
Retrieving an ObjectStore's ClassDescription . . . . .	577
Retrieving Realm Users and Groups . . . . .	577
Operating on Deleted Objects . . . . .	577
Type Property on Domain Class Not Supported . . . . .	577
GCD Changes . . . . .	577
PropertyDescription DateTime Values . . . . .	578
Content Engine Web Service . . . . .	578
Classes Not Supported . . . . .	578
Properties Not Supported . . . . .	579
Security . . . . .	579
Events and Subscriptions . . . . .	580
Query . . . . .	582
Storage . . . . .	583
Retrieving an ObjectStore's ClassDescription . . . . .	584
ObjectStore DatabaseType property . . . . .	584
Content Type . . . . .	584
Creating an Addon . . . . .	584
IsPersistent Property . . . . .	585
Document Lifecycles . . . . .	585
Document Classification . . . . .	585
GUID Strings . . . . .	586
Type Property on Domain Class Not Supported . . . . .	587
GCD Changes . . . . .	587
Content Java API . . . . .	587
JDK . . . . .	587
Removed from the API . . . . .	587
Deprecated . . . . .	588
Error Reporting . . . . .	589
Security . . . . .	590
Transactions . . . . .	593
Events and Subscriptions . . . . .	595
Query . . . . .	596
Auditing . . . . .	597
Document Lifecycle and Classification Actions . . . . .	598
Feature AddOns . . . . .	598
ServerEnvironment Class . . . . .	598
MimeType . . . . .	598
VersionSeries Security . . . . .	599
PropertyDescription DateTime Values . . . . .	599
Java SecurityManager . . . . .	599
Exception When Attempting to Modify Root Folder Properties . . . . .	599
Retrieving an ObjectStore's ClassDescription . . . . .	599
Publishing . . . . .	599
GCD Changes . . . . .	601
Process Java API . . . . .	601

Process Router .....	601
Directory Services and Authentication .....	602
Content Engine API Files .....	602
Web Application Toolkit .....	604
Upgrading Custom Applications to Web Application Toolkit Version 4.0 .....	604
Upgrading Custom Applications to Web Application Toolkit Version 3.x and Later .....	604
Upgrading Custom Applications to Access Workflows on a P8 System 4.0 .....	605
Obsolete .....	606
Workplace Application Integration Toolkit .....	609
Obsolete .....	609
Task 18: Upgrade Server-Side Scripts and COM Objects .....	610
File Document Handler .....	610
Log Event Handler .....	613
Send eMail Handler .....	615
Task 19: Upgrade ISRA Servlet .....	617
ISRA SSL Support .....	617
Task 20: Install Service Packs for Add-On Components .....	622
<b>Remove Software .....</b>	<b>623</b>
Remove the IBM FileNet P8 Documentation .....	624
Remove Content Search Engine .....	626
Remove Content Engine .....	628
Remove Process Engine (Windows) .....	631
Remove Process Engine (UNIX) .....	633
Remove Application Engine (WebSphere) .....	634
Remove Application Engine (WebLogic) .....	636
Remove Application Engine (JBoss) .....	637
Remove the Application Engine ISRA Servlet .....	639
<b>Appendixes .....</b>	<b>641</b>
Encrypt Passwords for Silent Installations and Upgrades .....	642
IBM FileNet P8 Port Numbers .....	643
Process Engine SQL Scripts .....	647
SQL Scripts for SQL Server .....	647
SQL Scripts for Oracle .....	648
Changes made by pe_filenet_site.sql and pe_ora_users_defaults.sql .....	649
Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql .....	650
IBM FileNet P8 Database Character Sets .....	653
IBM FileNet P8 Database components character sets .....	653
Additional Oracle character set information .....	653
Additional DB2 character set information .....	655
New Content Engine Classes and Properties .....	656

Content Engine Classes .....	656
Content Engine Properties .....	658
Index .....	669
<b>Notices .....</b>	<b>677</b>
Trademarks .....	679

## About this Document

Installation by an IBM FileNet Certified Professional (FCP) recommended. For more information on the FCP program, contact your IBM service representative.

Every effort has been made to provide you with complete installation instructions. If information becomes available after the documentation release from which you accessed this guide, we will provide an updated version of the guide on the Information Management support page ([www.ibm.com/software/data/support/](http://www.ibm.com/software/data/support/)). As a general rule, you should refer to the web site to obtain the current version of this guide.

A new version of the guide will be distributed as of the release of any core IBM FileNet P8 Platform component (Content Engine, Process Engine, or Application Engine). Component fix packs are often released independently of one another and can also require updates of the guide. Check the edition number applied to this common guide to determine the product versions covered.

Send your comments by e-mail to [comments@us.ibm.com](mailto:comments@us.ibm.com). Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a chapter and section title, a table number, a page number, or a help topic title).

### NOTES

- This guide provides instructions for installing and/or upgrading the core IBM FileNet P8 Platform components to their current released version. Be aware that each release of an IBM FileNet P8 component or expansion product can have multiple software updates available in the form of service packs, fix packs, and/or interim fixes, each with potentially different dependencies and installation requirements. Therefore, before you attempt to install or upgrade IBM FileNet P8 Platform, review the list of releases and their associated dependencies in the IBM FileNet Compatibility Matrix on the Information Management support page ([www.ibm.com/software/data/support/](http://www.ibm.com/software/data/support/)). See “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on page 21.
- If you do not intend to upgrade all IBM FileNet P8 components and expansion products to the most current released version, review the dependencies for the version(s) you intend to install. These dependencies will be contained in a document and location on the IBM Customer Service and Support web site for FileNet customers (equivalent to those noted in the previous bullet).
- This guide assumes that to perform your installations you are using the 4.0.0b update of the generally released Content Engine software, which among other things adds support for JBoss as an application server.
- This guide contains links to the IBM FileNet P8 Platform online help system, where you will find important additional information that can help you make decisions about your installation. The ideal way to use this guide is as an online PDF installed within the help system. To take advantage of these links while using this guide, perform the documentation install tasks ([Task 7a on page 130](#), [Task 7b on page 135](#), or [Task 7c on page 141](#)) before beginning your full IBM FileNet P8 Platform installation.

# Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs

## To access documentation for IBM FileNet products

---

1. Navigate to the Information Management support page ([www.ibm.com/software/data/support/](http://www.ibm.com/software/data/support/)).
2. Select the appropriate IBM FileNet product from the Select a category list. For IBM FileNet P8 documentation, select FileNet Content Manager or FileNet Business Process Manager.
3. From the Product Support page, under Learn, click **Product documentation**.
4. From the Product Documentation page:
  - a. If necessary, click the **Doc. Link** for the appropriate component product to display the document list.
  - b. Click the icon in the appropriate release column to access the document you need.

**NOTE** For a complete list of direct links to documentation relevant to this guide, see “[General Requirements for all IBM FileNet P8 Systems](#)” on page 24.

## To access compatibility matrices and fix packs for IBM FileNet products

---

1. Navigate to the Information Management support page ([www.ibm.com/software/data/support/](http://www.ibm.com/software/data/support/)).
2. Select the appropriate IBM FileNet product from the Select a category list. For IBM FileNet P8 documentation, select FileNet Content Manager or FileNet Business Process Manager.
3. From the Product Support page, under Download, click **Fix Packs by version**.
4. From the Fix Pack page:
  - a. Click **Matrix** to access the compatibility matrix.
  - b. Click the release number to access the fix pack you need.

## *Installation Planning and Procedures*

This installation section contains the following major topics:

- [“Plan the Installation” on page 23](#)
- [“Prerequisite Tasks” on page 53](#)
- [“Installation Tasks” on page 128](#)
- [“Configuration/Startup Tasks” on page 396](#)
- [“Optional Installation Tasks” on page 422](#)

# Plan the Installation

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

**NOTE** Although many of the bullet items below apply to upgrade as well as new installations, also see [“Plan the Upgrade” on page 459](#) for specific upgrade details.

- Retrieve updates to IBM FileNet P8 documentation and software from the [IBM Information Management support page on www.ibm.com](#).
- Review the [“Installation Planning Considerations” on page 24](#) for a list of auxiliary documentation you should gather and a list of tasks you should perform before installing IBM FileNet P8 Platform software.
- Use the items in [“Installation Checklists” on page 42](#) to track installation tasks and to record information that will be needed in later tasks.
- Review the [“Sample Configurations” on page 37](#).
- Read [“General Requirements for all IBM FileNet P8 Systems” on page 24](#), [“Installation Tasks” on page 128](#), [“Configuration/Startup Tasks” on page 396](#), and [“Optional Installation Tasks” on page 422](#) to become familiar with the tasks you will perform when setting up your IBM FileNet P8 Platform software.

**CAUTION** Follow the planning, installation, and configuration tasks in the order presented in these topics. If you do otherwise, or you install IBM FileNet P8 components on machines containing components for existing IBM FileNet P8 systems, you may encounter unforeseen problems or be required to perform additional setup steps. See the *IBM FileNet P8 Platform Troubleshooting Guide* for details. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

## Installation Planning Considerations

This section lists details that will help you prepare your environment for the installation of an IBM FileNet P8 system. In many cases, the items you see listed will be links to more detailed information that will help you plan a system rollout. Review this section thoroughly before you start to set up IBM FileNet P8 components or required third-party software.

When you have completed the considerations in this section, see the [“Installation Checklists” on page 42](#) for a convenient way to track your progress and collect necessary information throughout the prerequisite and installation stages.

### General Requirements for all IBM FileNet P8 Systems

- **Gather auxiliary documentation.**

**NOTE** To download this and other IBM FileNet product documentation from the IBM web site, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

- *IBM FileNet P8 Hardware and Software Requirements*. This document provides details for all IBM FileNet P8 system components, as well as the minimum supported levels of third-party software components. The information throughout the *IBM FileNet P8 Platform Installation and Upgrade Guide* assumes you have met all applicable requirements listed in that document.
- IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups](#). This help topic provides a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain an IBM FileNet P8 system.
- *IBM FileNet P8 Release Notes*. This document provides details on new features, known issues, and resolved problems.
- IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Shut-down and Startup](#). This help topic describes how to shut down and restart IBM FileNet P8 Platform components and some expansion products. Manual, command line, and some sample batch file procedures are provided.
- *IBM FileNet P8 Platform Installing Non-English Environments Technical Notice*. This document will help you set up the product if your environment is not English-language based.
- *IBM FileNet P8 Platform High Availability Technical Notice*. This document provides details on how to set up your IBM FileNet P8 system using clusters, farms, and other high-availability software and hardware.
- *IBM FileNet P8 Platform Planning and Deployment Guide*. This document provides details on how to deploy an IBM FileNet P8 system from a staging environment into a full production environment.
- *IBM FileNet P8 Platform Troubleshooting Guide*. This document provides troubleshooting information on all aspects of the product.
- *IBM FileNet P8 Platform Performance Tuning Guide*. This document provides performance tuning information on all aspects of the product.



- *IBM FileNet Rendition Engine Installation and Upgrade Guide*. This document provides details on how to install and upgrade Rendition Engine.
- For all optional components, see the IBM FileNet P8 4.0.0 version of the appropriate documentation, including:
  - *IBM FileNet P8 Content Federation Services for Image Services Guidelines*
  - *IBM FileNet P8 Process Analyzer Installation and Upgrade Guide*
  - *IBM FileNet P8 Process Simulator Installation and Upgrade Guide*
- **Determine your IBM FileNet P8 user application requirements.** Install Application Engine if you will be using Workplace or creating a custom user application. Alternatively, you can install and use Workplace XT. If you plan to use Workplace XT, installing Application Engine is not required.
- **Consider the Java™ requirements for the IBM FileNet P8 Documentation server.** To enable the help Search functionality, you must install the IBM FileNet P8 Documentation onto a Java-enabled application server (for example, any J2EE web application server supported by Content Engine or Application Engine).

## Operating System Considerations

### General

- **Ensure hosts file contents.** On each UNIX-based IBM FileNet P8 server that does not use DNS (Domain Name Service) or NIS (Network Information Service), the `/etc/hosts` file must contain the name and Internet Protocol (IP) address of all servers it will communicate with, including the remote database server, if applicable.
- **Synchronize the time and date on all servers.** System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

The Process Engine database server (the machine that hosts the database used by Process Engine) is considered the master time keeper; the UTC time of that machine is considered the correct time. The server hosting the Process Engine API and the server hosting Content Engine must have the UTC time set to match the UTC time on the Process Engine database server, plus or minus 15 minutes.
- **(HP-UX only) Adjust number of threads per process.** On each HP-UX 11 or HP-UX 11i machine where you will install a JVM-based IBM FileNet P8 component (such as Content Engine), or where an associated third-party JVM-based component (such as WebLogic or WebSphere) will run, increase the values of the kernel parameters `max_thread_proc` (maximum number of threads per process) and `nkthread` (maximum number of kernel threads in the system) beyond their default values, which are too small for IBM FileNet P8 applications.

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

## Content Engine

- If you intend to enable content-based retrieval (CBR), you must run your Content Engine servers and associated Content Search Engine (Autonomy K2) servers on the same operating system.
- This release of IBM FileNet P8 Platform 4.0.0 is certified to support UNIX-based Content Engine servers accessing Windows-based file storage areas and index areas. However, running Windows-based Content Engine servers to access UNIX-based file storage areas and index areas is not yet certified for support.
- If you plan to configure a remote file storage area, install a UPS power supply backup system on the remote connect file system so that the systems never go down or stay up long enough to shut down gracefully. Failure to gracefully shut down a server on which a remote file storage area resides can result in data loss or corruption.

## Process Engine

### (UNIX only) Volume Manager

- Process Engine requires the presence of several partitions. Before installing PE verify that your Operating System is set up with a correctly configured volume manager. You can use the volume manager provided with the operating systems or an equivalent Veritas volume manager.

## AIX

- **Ensure minimum /tmp size.** The **/tmp** directory must have 510 MB free.
- **Set Kernel.** The Kernel must be set to 64-bit mode.
- **Ensure minimum swap space.** The swap space must be set to 1.5 - 2 times RAM.
- **Set *Maximum Number of Processes allowed per user*.** The *Maximum Number of Processes allowed per user* must be set to at least 400.
- **Set *Maximum Kbytes of real memory allowed for MBUFFS*.** The *Maximum Kbytes of real memory allowed for MBUFFS* must be set to 0. Setting the MBUFFS parameter to 0 causes the system to use the default amount of available memory. This default amount is approximately 1/8 to 1/4 the amount of real memory.
- **Set *Maximum Number of FIXED licenses (Num)*.** The *Maximum Number of FIXED licenses (Num)* must be set to a minimum of 16.
- **Install required filesets.** The following filesets must be installed and committed:
  - bos.adt.libm
  - bos.adt.lib
  - bos.adt.base
  - bos.perf.perfstat

- bos.perf.libperfstat
- bos.adt.debug

### HP-UX PA-RISC

- **Ensure minimum /tmp size.** The **/tmp** directory must have 510 MB free.
- **Ensure minimum physical memory.** The physical memory must be at least 512 MB.
- **Set Kernel.** The Kernel must be set to 64-bit mode. See the topic [“Install Process Engine \(HP-UX\)” on page 320](#) for a complete list of minimum Kernel values.

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

- **Ensure minimum swap space.** The swap space must be set as follows:
  - Two times RAM if RAM < 1GB
  - 1.5 times RAM if RAM >= 1GB

### HP-UX Integrity

- **Ensure minimum /tmp size.** The **/tmp** directory must have 510 MB free.
- **Ensure minimum physical memory.** The physical memory must be at least 512 MB.
- **Set Kernel.** The Kernel must be set to 64-bit mode. See the topic [“Install Process Engine \(HP-UX\)” on page 320](#) for a complete list of minimum Kernel values.

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

- **Ensure minimum swap space.** The swap space must be set as follows:
  - Two times RAM if RAM < 1GB
  - 1.5 times RAM if RAM >= 1GB

### Solaris

- **Ensure minimum /tmp size.** The **/tmp** directory must have 510 MB free.
- **Set Kernel.** See the topic [“Install Process Engine \(Solaris\)” on page 285](#) for a complete list of Kernel values.

## Application Engine

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

## Network Considerations

- **Ensure static IP addresses.** Assign all IBM FileNet P8 servers a static IP address.
- **Ensure TCP/IP settings.** Verify TCP/IP configuration settings on all UNIX and Windows servers and Enterprise Manager clients.
- **Ensure availability of required port numbers.** Several port numbers are required by the various IBM FileNet P8 components. For a composite list, see [“IBM FileNet P8 Port Numbers” on page 643](#).

## Security Considerations

The security information in this section is provided to assist in the security planning process but is not a complete description of any security feature or level of support. For complete information about IBM FileNet P8 security, consult the Security Guide, IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security](#). The Security Guide sections that are especially applicable to installation are Authentication, Directory Service Providers, and Users and Groups.

## Proxy server or firewall configurations

- **Proxy/firewall configurations require full HTTP/1.1 compliance.** If you plan to use proxy server or firewall configurations with your IBM FileNet P8 Platform system, the proxy/firewall configurations must be fully HTTP/1.1 compliant.

## Authentication and Authorization

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

The Content Engine Setup program can capture configuration information to create your application server authentication provider; or you can use an authentication provider that already exists on the application server. Immediately following Content Engine installation, you will be required to use Enterprise Manager to configure a directory service provider (creating what is called a Directory Configuration).

- **Understand that all logins as of IBM FileNet P8 4.0.0 are done through JAAS.** IBM FileNet P8 leverages Java Authentication and Authorization Service (JAAS) for authentication only, which is a process that occurs between a J2EE client application, a J2EE application server, and one or more JAAS login modules. This process does not involve any FileNet code.

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

- **Determine your single sign-on (SSO) requirements.** Content Engine's ability to leverage JAAS-based authentication means that if a single sign-on (SSO) provider writes a JAAS LoginModule for a supported application server, then clients of FileNet P8 applications hosted in that application server can leverage that SSO solution.
- **Determine Kerberos applicability.** You can use Kerberos for SSO authentication between FileNet Enterprise Manager and Content Engine, provided you use Windows Active Directory as the directory server. See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authentication > Kerberos for Content Engine](#).

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

- **Determine how many authentication realms you require.** At least one authentication realm is required, which you establish via your directory service or authentication provider and specify during the procedure for creating an object store. For an explanation of how to create and configure multiple realms, for example, multiple Windows domains, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms](#).
- **Ensure that you have a directory service provider in place.** Authorization in IBM FileNet P8 Platform is provided by one of the following supported directory servers:
  - Microsoft Windows Active Directory
  - Microsoft Active Directory Application Mode (ADAM)
  - Novell eDirectory
  - Sun Java System Directory Server
  - IBM Tivoli® Directory Server

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

- **Note that any WebLogic authentication provider should be dedicated to IBM FileNet P8.** For performance reasons, no authentication provider used by WebLogic for deployed IBM FileNet P8 components should be shared with applications used for other purposes.
- **Understand the users and groups required for IBM FileNet P8.** All general administrative users and groups needing access to IBM FileNet P8-based applications must reside in one of the supported directory servers. This guide provides instructions for creating the administrative accounts required for installation and initial configuration. You can find additional detailed reference information for all users, groups, and administrative roles in the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups](#).

- **Note that Process Engine delegates authentication to Content Engine.** As of the IBM FileNet P8 Platform 4.0.0 release, Process Engine no longer has a direct connection to a directory server for authentication purposes, as it did in earlier releases. Instead, it delegates authentication tasks to Content Engine. Content Engine, in turn, runs under a J2EE application server, and relies on that server's JAAS-based facilities to authenticate users and groups against the chosen directory server.

## Microsoft Windows Active Directory

This subsection applies only to installations using Microsoft Windows Active Directory as a directory server.

**NOTE** See [“Configure Windows Active Directory” on page 54](#) for details on related restrictions and requirements.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Windows Active Directory](#) for complete reference information.

- IBM FileNet P8 does not support Windows Built-in groups, roles, directory aliases, or follow referrals in Active Directory searches.
- IBM FileNet P8 does not support Windows NT® domains (Windows versions 4.0 and earlier).

## Microsoft Active Directory Application Mode (ADAM)

This subsection applies only to installations using Microsoft Windows Active Directory Application Mode (ADAM) as a directory server.

**NOTE** See [“Configure Active Directory Application Mode \(ADAM\)” on page 55](#) for details on related restrictions and requirements.

You can use ADAM as a stand-alone directory service, or ADAM can be synchronized with Active Directory, using Microsoft's builtin tools. Synchronization is invisible to IBM FileNet P8 applications and authentication.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Windows Active Directory Application Mode \(ADAM\)](#) for complete reference information.

## Novell eDirectory

This subsection applies only to installations using Novell eDirectory as a directory server.

**NOTE** [“Configure Novell eDirectory” on page 58](#) for details on related restrictions and requirements.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Novell eDirectory](#) for complete reference information.

- IBM FileNet P8 does not support dynamic groups, roles, directory aliases, or follow referrals in searches.

- Access control settings in IBM FileNet P8 require that all users have Browse access on the directory server. If you do not want to set Browse access at the individual user level, IBM recommends you establish a Public trustee for the realm.
- IBM FileNet P8 supports cross-realm group memberships. This means that IBM FileNet P8 supports a configuration in which a group is in one realm while some or all of its users are in another.
- IBM FileNet P8 does not support the Transport Layer Security (TLS) cryptographic protocol.

## Sun Java System Directory Server

This subsection applies only to installations using Sun Java System Directory Server as a directory server.

**NOTE** See [“Configure Sun Java System Directory Server” on page 56](#) for details on related restrictions and requirements.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Sun Java System Directory Server](#) for complete reference information.

- Access control settings in IBM FileNet P8 require that all users have Read access on the directory server. If you do not want to set Read access at the individual user level, you can establish this configuration by one of the following:
  - Setting up anonymous access privileges
  - Adding all users to an authenticated users group having Read access.
- If you enable one-way SSL, use the fully qualified domain name of the Sun Java System Directory Server when requesting the certificate.
- If multiple users within a realm have the same UID value, these users cannot log on via IBM FileNet P8 application sign-in screens using their short name (uid). Instead, they must configure the application server to use the fully qualified distinguished name (FQDN).

For example, consider two separate users in two different organizational units (OUs):

```
uid=jdoe,ou=Sales,dc=filenet,dc=com
```

```
uid=jdoe,ou=Marketing,dc=filenet,dc=com
```

In this case neither user can log on with a short name (uid) of jdoe. Each user must use his fully qualified name. For example, jdoe in Marketing must enter all of the information necessary to create the fully qualified name uid=jdoe,ou=Marketing,dc=filenet,dc=com. (If a custom logon application is used, some of this information can be defaulted or hidden from the user.)

- No user can have a blank password. Any user attempting to log on (for example, to Workplace or Enterprise Manager) or to configure the Process Engine LDAP connection will receive an error message.
- IBM FileNet P8 does not support dynamic groups, universal groups, roles, directory aliases, or follow referrals in searches.



## IBM Tivoli Directory Server

This subsection applies only to installations using IBM Tivoli Directory Server as a directory server.

**NOTE** See “[Configure IBM Tivoli Directory Server](#)” on page 59 for details on related restrictions and requirements.

See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > IBM Tivoli Directory Server](#) for complete reference information.

- IBM FileNet P8 does not support dynamic groups, roles, directory aliases, follow referrals for searches.

## Database Considerations

### All platforms

- **Determine your local vs. remote database engine preference.** The database engine is *local* if it is installed on a server where you will also be installing Content Engine, Process Engine and/or Rendition Engine software. The database is *remote* if it is on a separate server from the component using that database. The choice is yours with the exception of DB2, which must be remote from Process Engine. An Oracle database must be remote if it is installed on a Linux server.
- **Determine what type(s) of database engine you want to use.** See the *IBM FileNet P8 Hardware and Software Requirements* for support information related to IBM FileNet P8 components and database engines. To download this guide from the IBM support page, see “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on page 21.
- **Determine your NLS and character-set requirements.** See “[IBM FileNet P8 Database Character Sets](#)” on page 653 for information.
- **Decide when to upgrade the database.** IBM FileNet P8 requires no re-configuration of the application server when the database server is undergoing an upgrade as long as connection information in the data sources does not change and no data is lost when performing the database upgrade procedures. Be sure to update to the JDBC driver associated with the new database version as specified in *IBM FileNet P8 Hardware and Software Requirements*.

### Microsoft SQL Server

- **Determine whether you want to use a dedicated or shared database.** In this regard:
  - Content Engine, Process Engine, and Rendition Engine can share a database engine, or they can each have a dedicated (unique) database engine. The Process Analyzer expansion product must have its own database engine.
  - Content Engine, Process Engine, and Rendition Engine components can each have a dedicated SQL Server database instance, or they can share a database instance with one another or with non-IBM FileNet applications. You can use the default instance or a named instance of your choosing.



- Content Engine, Process Engine, and Rendition Engine must each have their own databases, even if they share a database engine.
- An MS SQL Server database name cannot have a numeric value as the first character.
- **Create a SQL Login for Rendition Engine purposes.** If you plan to install Rendition Engine you must create a SQL Login for the RE database. For detailed information, see IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).
- **Verify that a SQL Login (Database User) exists for Content Engine purposes.** Under MS SQL Server, you must ensure that a SQL Login account exists before you create your first Content Engine object store. For this purpose, you can create a specific user with access only to the global configuration data (GCD) database.
- **Be aware of database client software requirements.** For Process Engine, if the database is remote, you might need to install database client software on the Process Engine server. See [“Verify that Microsoft SQL Server Client Is Installed for IBM FileNet P8” on page 98](#) for details.
- **SQL Server scripts must be executed for Process Engine.** A number of SQL scripts must be executed that create a number of stored procedures. These stored procedures will then be used to create database users and other stored procedures required for PE production.

These scripts can be executed manually, before starting Process Engine installation, or executed from Process Engine setup. If the scripts are executed manually before Process Engine Setup, they will also create the PE database users. See [“Process Engine SQL Scripts” on page 647](#) for information on execution modes and associated security requirements as well as details about the scripts.

Process Engine Setup will complete only if all these scripts run successfully.

- **Determine the maximum size of the content elements your users store.** This affects setting up database storage areas or file storage areas. When you create an object store, a database storage area is provided by default, allowing you to store content as database BLOBs. You can also create one or more file storage areas to store content on local or remote file systems. If your users store large individual documents or other content elements, and you create the associated object stores on Microsoft SQL Server or IBM DB2 databases, use only file storage areas. Otherwise, users can encounter memory-related errors when retrieving or indexing the large content.

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

## Oracle

### General

- **Oracle 9i support.** Oracle 9i is nearing end-of-life and is thus supported only for upgrades from IBM FileNet P8 3.5.x. If you are planning a new Oracle-based installation of IBM FileNet P8 4.0.x, use Oracle 10g.

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

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

- **Verify that an Oracle alias (Database User) exists for Content Engine purposes.** Under Oracle, you must ensure that an Oracle account exists before you create your first Content Engine object store. For this purpose, you can create a specific user with access only to the global configuration data (GCD) tablespace.
- **Plan to use locally managed tablespaces.** For performance reasons, IBM recommends that you create locally managed, rather than dictionary managed, tablespaces for Process Engine and Content Engine. (The tablespaces you create via Oracle Database assistant (dbca) are locally managed by default.)
- **Be aware of database client software requirements.** For Process Engine, if the database is remote, you must install database client software on the Process Engine server.

## Process Engine

- **Process Engine does not support Oracle Password Complexity Verification during the installation process.** During installation this Oracle feature must be turned off for the Process Engine run-time user (f\_sw or alias). After installation is complete and the f\_sw (or alias) user password is changed via the set\_f\_maint\_pw utility, Oracle's password verification can be turned back on.
- **Oracle SQL scripts must be executed.** A number of SQL scripts must be executed that:
  - Create an Oracle database account for IBM FileNet PE use.
  - Create a number of stored procedures.
  - Grant access levels to the default tablespace and temporary tablespace specified in Process Engine Setup.

These scripts can be executed manually, before starting Process Engine installation, or executed from Process Engine setup. If the scripts are executed manually before Process Engine Setup, they will also create the PE database users. See [“Process Engine SQL Scripts” on page 647](#) for information on execution modes and associated security requirements as well as details about the scripts.

Process Engine Setup will complete only if all these scripts run successfully.

## DB2

- **DB2 server for Process Engine must be remote.** The DB2 database server must be remote from Process Engine. See the *IBM FileNet P8 Hardware and Software Requirements* for specifics on supported operating systems. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)
- **DB2 server instances must be 64-bit on Unix, 32-bit on Windows.** Instances created on a remote Unix database server must be 64-bit. On Windows database servers, instances must be 32-bit.
- **DB2 client instances must be 32-bit.** Instances created on the Process Engine server must be 32-bit.
- **Determine whether you want to use a dedicated or shared database.** In this regard:
  - Content Engine and Process Engine can share a database engine, or they can each have a dedicated (i.e., unique) database engine.
  - Content Engine and Process Engine can each have a dedicated DB2 instance, or they can share an instance with one another or with non-IBM FileNet applications.
  - Each Content Engine and Process Engine must have their own databases and tablespaces.
  - The Content Engine GCD and each object store must have their own databases.
- **Plan to use Database Managed User(DMS) for tablespaces.** For performance reasons, IBM recommends that you create database managed user and user temp tablespaces rather than system managed tablespaces for Process Engine and Content Engine.
- **Be aware of database client software requirements.** For Process Engine on UNIX, if the remote database server software is version 9, the DB2 client software on the Process Engine server must be version 8.
- **Use DB2 default collation.** For both Content Engine and Process Engine, use the default collation setting.
- **Predefine several operating system users and groups.** For Process Engine, an instance owner and fenced user, and groups must be created. Process Engine also requires two specific users be created. For Content Engine, there are no specific users or groups to be created, but an operating system user must be identified for access to the Content Engine database. This could be an existing user. Database permissions will need to be granted to this user. See [“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92](#) for details.
- **IBM recommends SERVER authentication.** SERVER\_ENCRYPT and CLIENT authentication are also supported.
- **Be aware of recommended pagesizes.** IBM recommends 8KB minimum pagesizes for Process Engine databases and requires 32KB minimum pagesizes for Content Engine databases.
- **Determine the maximum size of the content elements your users store.** This affects setting up database storage areas or file storage areas. When you create an object store, a database storage area is provided by default, allowing you to store content as database BLOBs. You can also create one or more file storage areas to store content on local or remote file systems. If your users store large individual documents or other content elements, and you create the associated object stores on Microsoft SQL Server or IBM DB2 databases, use only file storage

areas. Otherwise, users can encounter memory-related errors when retrieving or indexing the large content.

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

## *J2EE Application Server Considerations*

**Requirement for J2EE application servers.** Content Engine and Application Engine are J2EE application server-based applications. (Process Engine is not.) You must install Content Engine and Application Engine in a homogeneous J2EE environment in which all of your application servers (IBM WebSphere, BEA WebLogic, or JBoss) and their version numbers are identical for both components. Also, the applications must use Enterprise Java Bean (EJB) transport.

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

**The Java Virtual Machine determines the maximum number of object stores.** If the application server where Content Engine will be deployed is running on a 32-bit JVM, you should not create more than 75 Content Engine object stores. On a 64-bit JVM, you should not create more than 150 Content Engine object stores.

**Deploying Content Engine and other applications on the same machine.** Content Engine 4.0.x is a resource-intensive enterprise application. Running Content Engine and other J2EE applications on the same machine is possible but not recommended. Other J2EE applications will compete with Content Engine for the same CPU, memory, and disk I/O resources, and increase the complexity of the installation and the risk of the deployment, as configurations will not match what has been qualified by IBM FileNet Engineering.

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

## Sample Configurations

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

This topic includes the following sample configurations:

- [“Baseline Configuration” on page 38](#)
- [“Baseline Configuration With Optional Components” on page 39](#)
- [“Developer Configuration” on page 40](#)
- [“Demo Configuration” on page 41](#)

### **In all the sample configurations, note that:**

- None of the samples shows a database engine, but at least one is required. You can collocate either Microsoft SQL Server or Oracle database engines with IBM FileNet P8 Platform components on any of the servers shown, or you can install them on separate database servers. A DB2 database engine must be remote from Process Engines. Likewise, an Oracle database engine running on a Linux machine must be remote from all IBM FileNet P8 components.
- None of the samples shows the required directory service provider.
- You can scale out the components, but the following graphics do not attempt to show this.
- If you choose to run Content Engine on a UNIX server, as of release 4.0.0 you will also need a Windows administrative client for your Enterprise Manager installation.
- To administer Application Engine, you can run IBM FileNet P8 Workplace clients either from the computers shown or from one or more browser clients.
- You can also run Workplace as a user client to create and access stored content and processes. Optionally, you can configure these clients to integrate with Microsoft Outlook and Office applications, or to work in conjunction with expansion products, for example, IBM FileNet P8 eForms, IBM FileNet P8 Portlets, or IBM FileNet Records Manager, as they come available on a given release. You can also use WebDAV clients.
- If you want to apply business rules to Process Designer workflows, you can install a rules engine of your choice, such as ILOG JRules, which is not shown in the samples.
- Optional components not shown in the samples include: IBM FileNet P8 Portlets and Image Services Resource Adapter (ISRA). Check with your service representative for availability of other expansion products.
- You must set up Content Engine, Application Engine, and the P8 Platform Documentation on application servers. You can collocate the documentation with Content Engine or Application Engine, or deploy it on a dedicated server, as shown.

## Baseline Configuration

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

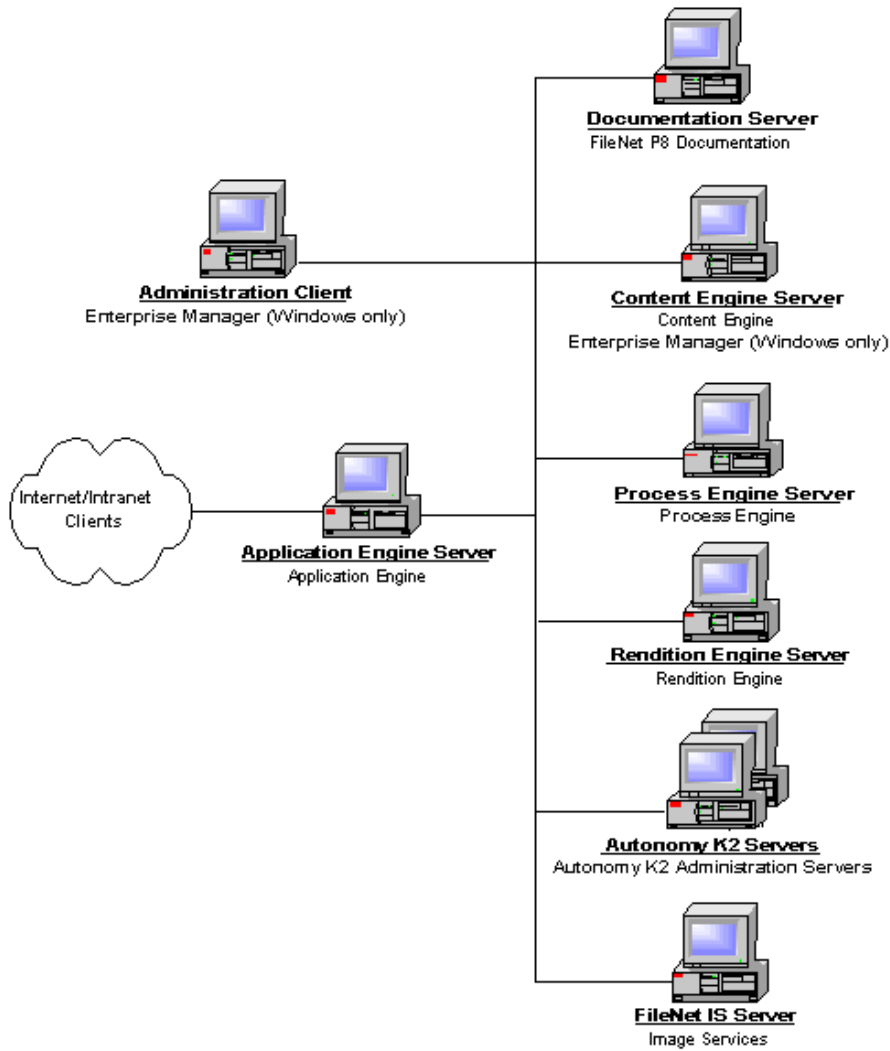


Figure 1: Baseline Configuration

## Baseline Configuration With Optional Components

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

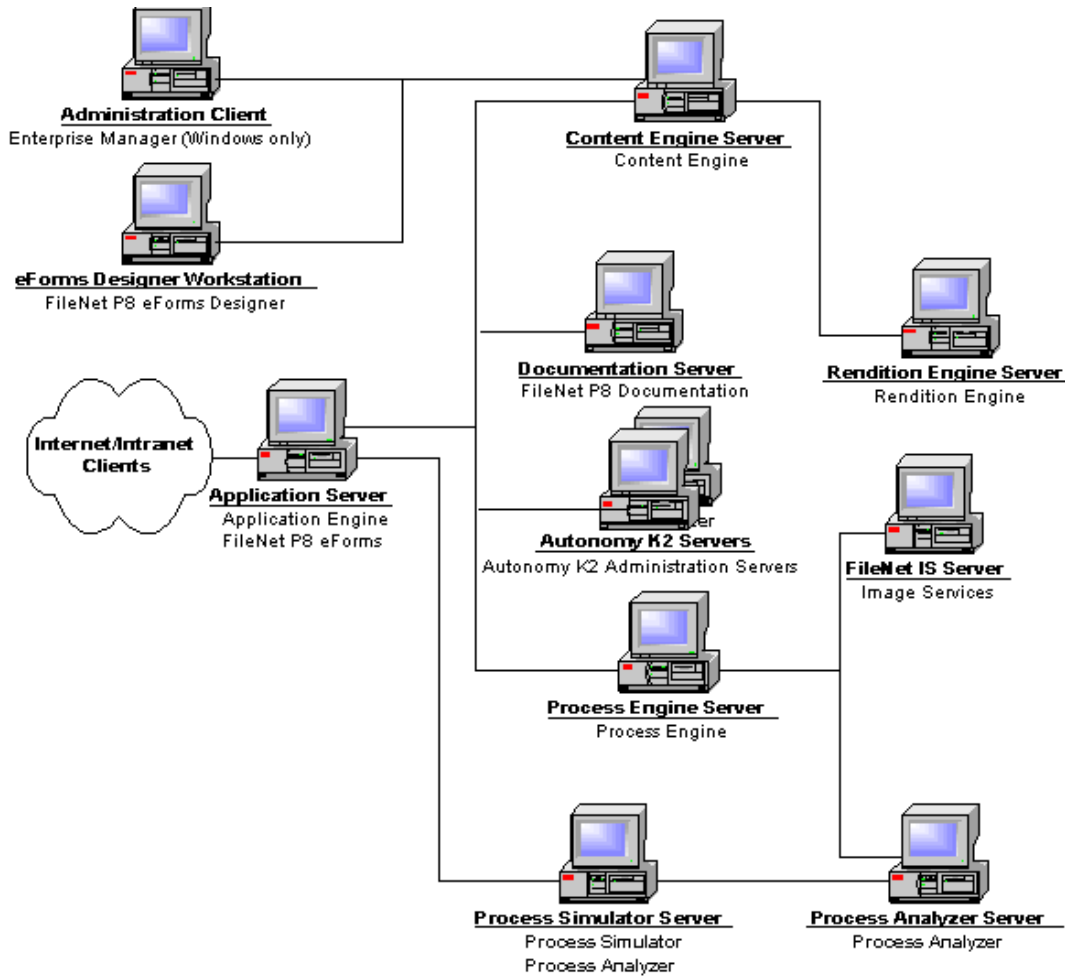


Figure 2: Baseline Configuration With Optional Components

## Developer Configuration

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

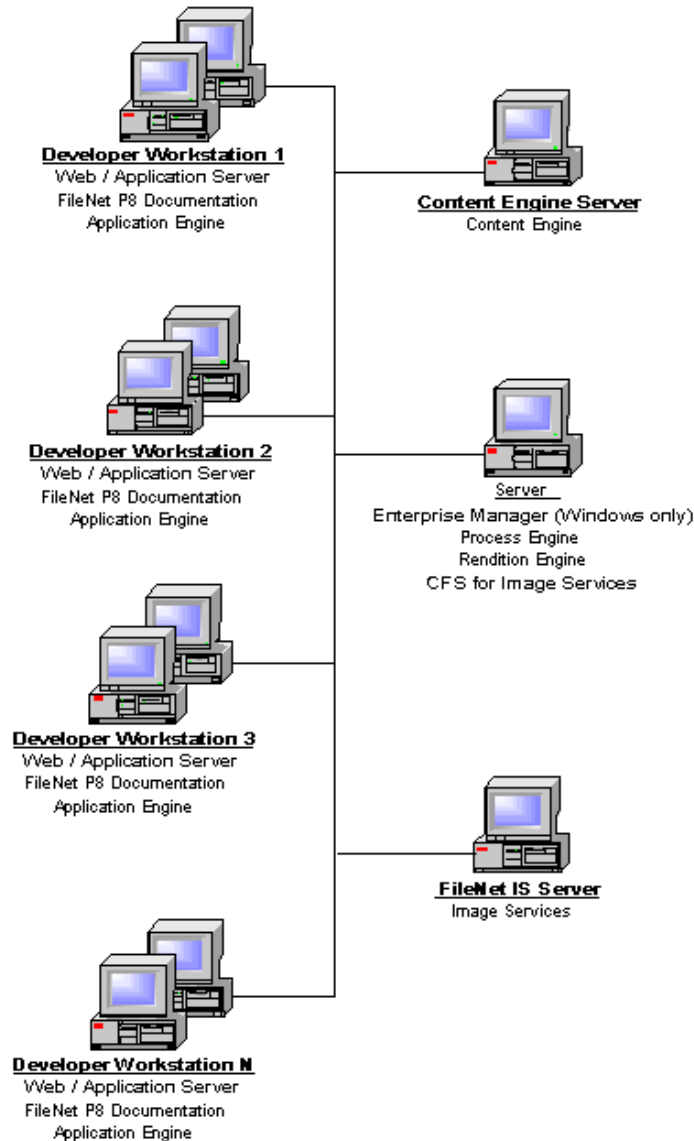


Figure 3: Developer Configuration

### NOTES

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



- This configuration represents a single IBM FileNet P8 domain.
- Refer to the IBM FileNet P8 Developer Help topic [IBM FileNet P8 Documentation > Developer Help > Developer Roadmap > Introduction](#) for information on setting up your development environment and installing the IBM FileNet P8 API toolkits.

## Demo Configuration

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



**Figure 4: Demo Configuration**

### NOTE

- Avoid collocating IBM FileNet P8 components. See the *IBM FileNet P8 Hardware and Software Requirements* for details. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

## Installation Checklists

The following topics provide checklists of items that must be completed for a successful installation or upgrade of the IBM FileNet P8 Platform. The lists are divided by user or role. Note that your organization may have different administrator roles, and that some of the responsibilities of listed roles may vary.

You can print these lists and record the information as you perform the prerequisite and installation tasks. Where applicable, the list items indicate which role will need the information you provide. Coordinate communication among the various roles to facilitate an easier installation process.

### *IT Administrator*

This role administers hardware and operating systems for your environment, as well as helping to make decisions about machine usage and configurations such as clustering and farming. The information confirmed and collected in this set of checklists should be provided to the Installation Administrator for each component.

#### P8 Platform Environment

- Consult with the application server, database, and P8 administrators to determine port requirements for all the servers in your install environment. For details, see ["IBM FileNet P8 Port Numbers" on page 643](#).

#### Content Engine

- Provide the host name of the application server machine and the directory where Content Engine is to be installed:

**Host name:** \_\_\_\_\_

**Directory:** \_\_\_\_\_

- (Windows installations only) Verify the version of Microsoft .NET Framework (2.0) and Web Services Enhancements (WSE) (3.0). If these are the versions you plan to use, and they are not installed, install them prior to the Content Engine installation.

#### Process Engine

- Decide whether to perform a standalone installation, a farmed installation, or a clustered installation.

**Standalone** \_\_\_\_\_

**Clustered** \_\_\_\_\_

**Farmed** \_\_\_\_\_

- Determine a location for the following files that will be added to the machine during installation.

**NOTE** In a clustered environment, you must place configuration and data files on a shared drive.

**Common files:** \_\_\_\_\_  
**Program files** \_\_\_\_\_  
**Configuration and data files:** \_\_\_\_\_

- (All UNIX platforms) Save the following files for the fnsw (or alias) and root users. The Process Engine installer will make modifications to them and you will be instructed to restore your custom changes after installing Process Engine software.

**.Xdefaults** \_\_\_\_\_  
**.Xresources** \_\_\_\_\_  
**.dbxinit** \_\_\_\_\_  
**.dtprofile** \_\_\_\_\_  
**.env** \_\_\_\_\_  
**.login** \_\_\_\_\_  
**.mwmrc** \_\_\_\_\_  
**.profile** \_\_\_\_\_  
**.cshrc** \_\_\_\_\_  
**.xinitrc** \_\_\_\_\_

### Application Engine

- (WebLogic) Determine the recommended MaxPermSize value for MEM\_ARGS:

**MaxPermSize:** \_\_\_\_\_

For information refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM FileNet specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs"](#) on page 21.

## Security Administrator

This role administers authentication, users and groups, passwords, encryption, and general network access considerations for the installation and eventual use of the P8 Platform software. The information confirmed and collected in the following checklist items should be provided to the Installation Administrator for each component, as well as to the P8 administrators (for example, the GCD administrator).

### User Accounts

- Create necessary P8 Platform accounts.

The task ["Specify IBM FileNet P8 Accounts"](#) on page 60 provides details about the accounts that are required to install and configure the P8 Platform, and directs you to the install or upgrade tasks that require the accounts. The topic provides a worksheet format for recording the

accounts you create or specify so that you can have them to refer to throughout the install or upgrade process.

## Content Engine

Configure the directory service provider. See one of the following:

- “Configure Windows Active Directory” on page 54
- “Configure Active Directory Application Mode (ADAM)” on page 55
- “Configure Sun Java System Directory Server” on page 56
- “Configure Novell eDirectory” on page 58
- “Configure IBM Tivoli Directory Server” on page 59

Record the following information about the directory service (authentication) provider relative to the application server.

For details on these parameters, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory Service Providers](#), and navigate to the section about your directory service provider.

**Host:** \_\_\_\_\_

**Port:** \_\_\_\_\_

**Distinguished name of user (DN) that the app server will use to connect to the authentication provider:** \_\_\_\_\_

**User Base DN:** \_\_\_\_\_

**User Name Attribute:** \_\_\_\_\_

**User From Name Filter:** \_\_\_\_\_

**Group Base DN:** \_\_\_\_\_

**Group From Name Filter:** \_\_\_\_\_

**Static Group Name Attribute:** \_\_\_\_\_

**NOTE** This information must also be available during the Application Engine install.

Designate a user account in the directory service to serve as WebSphere administrator:

**WebSphere administrator:** \_\_\_\_\_

(CE on WebSphere for Windows, using Windows Active Directory) Provide the name or IP address of the domain controller:

**Domain controller name or IP address:** \_\_\_\_\_

## Application Engine

- Specify the users and groups to add to the Application Engine Administrators access role. For details on this role, see [“Specify IBM FileNet P8 Accounts” on page 60](#).

Users or groups

- Specify the users and groups who will be allowed to create subscriptions to add to the PWDesigner access role.

**NOTE** These are not specifically required to complete an install, but required to create subscriptions later in Workplace, and can also be added later:

Users or groups

- For Single Sign On (SSO), specify the following:

**SSO proxy host URL:** \_\_\_\_\_  
**SSO proxy host server name:** \_\_\_\_\_  
**HTTP port on the SSO proxy host:** \_\_\_\_\_  
**HTTPS port on the SSO proxy host:** \_\_\_\_\_

- For SSL, specify the following:

**SSL host name and port number:** \_\_\_\_\_  
**Java Server HTTP port:** \_\_\_\_\_

- (WebLogic) If you are using container-managed authentication, configure a new password or credential in the WebLogic interface. See [“To configure Application Engine \(WebLogic 8.1.x\)” on page 376](#) or [“To configure Application Engine \(WebLogic 9.2\)” on page 378](#).

## Database Administrator

This role oversees database creation and administration. The decisions, task confirmations, and information gathered in the following checklist items should be provided to the Installation Administrator for each component.

### Content Engine

- Set up the database servers for Content Engine. Refer to one of the following:
  - “Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81
  - “Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85
  - “Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92
- For SQL, record the following information for the GCD and each object store:

**JDBC connection pool name (available from the application server administrator):** \_\_\_\_\_  
**Type of authentication (Windows or database engine):** \_\_\_\_\_  
**Database name:** \_\_\_\_\_  
**Database host name or IP address:** \_\_\_\_\_  
**Database port:** \_\_\_\_\_  
**Database user name:** (For details, see [Task 2 on page 60.](#)) \_\_\_\_\_  
**Password:** \_\_\_\_\_

- For Oracle, record the following information for the GCD and each object store:

**JDBC connection pool name (available from the application server administrator):** \_\_\_\_\_  
**Net service name:** \_\_\_\_\_  
**Database name:** \_\_\_\_\_  
**Database host name or IP address:** \_\_\_\_\_  
**Database port:** \_\_\_\_\_  
**Tablespace user names for the GCD and each object store** (See [Task 2 on page 60.](#)) \_\_\_\_\_  

- **Permanent tablespace:** \_\_\_\_\_
- **Temporary tablespace:** \_\_\_\_\_

**Password:** \_\_\_\_\_

- For DB2, record the following information for the GCD and each object store:

**JDBC connection pool name (available from the application server administrator):** \_\_\_\_\_  
**Tablespace name:** \_\_\_\_\_  
**Database host name or IP address:** \_\_\_\_\_  
**Database port:** \_\_\_\_\_  
**Database user name:** (See [Task 2 on page 60.](#)) \_\_\_\_\_  
**Password:** \_\_\_\_\_

## Process Engine

- Set up the database server for Process Engine. Refer to one of the following:
  - “Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81
  - “Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85
  - “Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92
- If using a remote database, install the database client software for Process Engine. Refer to one of the following:
  - “Verify that Microsoft SQL Server Client Is Installed for IBM FileNet P8” on page 98
  - “Verify that Oracle Client Is Installed for IBM FileNet P8” on page 100
  - “Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103
- Turn off Oracle password complexity before beginning the Process Engine install.
- Is the database local or remote?
  - Local** \_\_\_\_\_
  - Remote** \_\_\_\_\_
- (SQL) Specify the SQL sa password (needed to create the ODBC data source and to run SQL scripts from PE Setup):
- (SQL) For SQL, provide the following information
  - SQL ODBC data source name:** \_\_\_\_\_
  - Database name:** \_\_\_\_\_
  - Filegroup name:** \_\_\_\_\_
  - SQL version:** \_\_\_\_\_
- Decide whether to run SQL scripts for a SQL Server or Oracle database before running the Process Engine installer, from the installer and prompting the user for a password, or silently using operating system authentication.
 

<b>Before running setup</b>	<b>Prompt for a password from setup</b>	<b>Silently, using operating system authentication</b>
_____	_____	_____

- ❑ If running SQL scripts manually for a SQL Server database before installing Process Engine, collect the following information for inclusion in the scripts.

Parameter	Value
SQL Server database name for Process Engine	
DSN (ODBC data source name)	
runtime user	f_sw (default) or alias to be assigned in Process setup (See <a href="#">Task 2 on page 60.</a> )
maintenance user	f_maint (default) or alias to be assigned in Process setup (See <a href="#">Task 2 on page 60.</a> )

- ❑ If running SQL scripts manually for an Oracle database before installing Process Engine, collect the following information for inclusion in the scripts.

Parameter	Value
Oracle data tablespace name	
Oracle index tablespace name	
Oracle temporary tablespace name	
runtime user	f_sw (default) or alias to be assigned in Process setup (See <a href="#">Task 2 on page 60.</a> )
maintenance user	f_maint (default) or alias to be assigned in Process setup (See <a href="#">Task 2 on page 60.</a> )

- ❑ (Oracle Windows) For Oracle on Windows, provide the following information as applicable:

**Oracle SYS password:** (only if running sql scripts from installer)

**Oracle Home directory**

- Remote:

- Local:

**Global database name (remote):**

**Temporary tablespace name**

- Remote:

- Local:

**Data tablespace name**

- Remote:

- Local:



**Index tablespace name (optional)**

- Remote: \_\_\_\_\_
  - Local: \_\_\_\_\_
- Oracle SID (local): \_\_\_\_\_
- Oracle version: \_\_\_\_\_

- (Oracle UNIX) For Oracle on UNIX, provide the following information

**Oracle SYS password:** (only if running sql scripts from installer)

**Oracle Home directory** \_\_\_\_\_

- Remote: \_\_\_\_\_
- Local: \_\_\_\_\_

**Global database name (remote):** \_\_\_\_\_

**Temporary tablespace name**

- Remote: \_\_\_\_\_
- Local: \_\_\_\_\_

**Data tablespace name**

- Remote: \_\_\_\_\_
- Local: \_\_\_\_\_

**Index tablespace name (optional)**

- Remote: \_\_\_\_\_
- Local: \_\_\_\_\_

**Oracle SID (local):** \_\_\_\_\_

**Oracle user name** (See [Task 2 on page 60.](#))

- Remote: \_\_\_\_\_
- Local: \_\_\_\_\_

**Oracle DBA group** (See [Task 2 on page 60.](#))

- Remote: \_\_\_\_\_
  - Local: \_\_\_\_\_
- Oracle version: \_\_\_\_\_

- (DB2) For DB2, provide the following information

**Database alias name:** \_\_\_\_\_

**Tablespace name:** \_\_\_\_\_

**Instance owner name (UNIX only):** \_\_\_\_\_

**f\_sw (or alias) password:** \_\_\_\_\_

**f\_maint (or alias) password** \_\_\_\_\_

## ***Installation Administrator***

This role will actually perform installation tasks for your P8 Platform software. This role may also perform initial configuration, setup, and startup tasks, and is typically filled by one or more operating system administrators on the associated computers.

## Environment Considerations

- Install the IBM FileNet P8 documentation per [Task 7a on page 130](#) (WebSphere), [Task 7b on page 135](#) (WebLogic), or [Task 7c on page 141](#). Record the URL:

**Documentation Server URL:** \_\_\_\_\_

## Content Engine

- Install WebLogic or WebSphere.
- For WebSphere, record the following:

**WebSphere installation:** \_\_\_\_\_  
**WebSphere profile:** \_\_\_\_\_  
**WebSphere cell:** \_\_\_\_\_  
**Node (machine on which WebSphere is managed):** \_\_\_\_\_  
**Server (name of WebSphere instance):** \_\_\_\_\_  
**SOAP port:** \_\_\_\_\_  
**HTTP port:** \_\_\_\_\_  
**WebSphere administrator:** (See [Task 2 on page 60.](#))  
**Password:** \_\_\_\_\_

- For WebLogic, record the following:

**WebLogic domain:** \_\_\_\_\_  
**WebLogic Root Directory:** \_\_\_\_\_  
**WebLogic Configuration Tool Directory:** \_\_\_\_\_  
**WebLogic Domains Directory:** \_\_\_\_\_  
**Administration Server:** \_\_\_\_\_  
**WebLogic Administrator:** (See [Task 2 on page 60.](#))  
**Password:** \_\_\_\_\_

- Configure the application server for Content Engine. See [“Configure an Application Server for Content Engine \(WebSphere\)” on page 106](#) and [“Configure an Application Server for Content Engine \(WebLogic\)” on page 110](#).
- Install and deploy the Content Engine software. See [“Install and Deploy Content Engine” on page 160](#). You will also create a P8 domain as part of that task.
- Set up an object store and test the Content Engine installation. See [“Create Object Stores” on page 258](#) and [“Verify the Content Engine Installation” on page 263](#).
- After installation, record the Content Engine client software URL (Content Engine server name and port number), which will be entered during Process Engine and Application Engine setup:

**Content Engine Client software URL:** \_\_\_\_\_

- After installation, record the URL for the Component Manager on Content Engine, which will be entered during Application Engine setup:

**Content Engine Component Manager URL:** \_\_\_\_\_

- After installation, record the Content Engine URLs (Content Engine server name and port number) for downloading and uploading document content, which will be entered during Application Engine setup:

**Content Engine download URL:** \_\_\_\_\_  
**Content Engine upload URL:** \_\_\_\_\_

## Content Search Engine

- Install the Autonomy K2 Master Administration Server. See [“Install and Configure Content Search Engine” on page 147](#).
- (Optional) Install additional K2 Administration Servers. See [“Install Additional Content Search Engine Servers” on page 447](#).
- Configure Content Engine for content-based retrieval. See [“Configure Content Engine for Content-Based Retrieval” on page 397](#).

## Process Engine

- Before the Process Engine installation, determine the J2EE application server type (WebLogic, WebSphere, or JBoss) and version used for the Content Engine installation:

**Content Engine Application Server Type:** \_\_\_\_\_  
**Content Engine Application Server Version:** \_\_\_\_\_

- Install Process Engine per the procedure for your platform:
  - [“Install Process Engine \(Windows\)” on page 266](#)
  - [“Install Process Engine \(Solaris\)” on page 285](#)
  - [“Install Process Engine \(AIX\)” on page 304](#)
  - [“Install Process Engine \(HP-UX\)” on page 320](#)
- Configure Process Task Manager per [Task 22 on page 340](#).
- Complete post-install Process Engine configuration per [Task 23 on page 342](#).

## Application Engine

- Verify that at least one Content Engine object store has been created. See [“Create Object Stores” on page 258](#).
- Configure the application server for Application Engine, per [Task 6a on page 123](#) for WebSphere, [Task 6b on page 124](#) for WebLogic, or [Task 6c on page 126](#) for JBoss. Provide the application server type and version (this must match the type and version used for Content Engine):

**Application Engine Application Server Type:** \_\_\_\_\_  
**Application Engine Application Server Version:** \_\_\_\_\_

- Decide on an application name; Workplace is the default (relevant for custom applications):

**Application name:** \_\_\_\_\_

- Install Application Engine. See [“Install Application Engine” on page 344](#).
- Set up Application Engine per the procedure for your platform:
  - [“Configure Application Engine \(WebSphere\)” on page 356](#)
  - [“Configure Application Engine \(WebLogic\)” on page 373](#)
  - [“Configure Application Engine \(JBoss\)” on page 381](#)

# Prerequisite Tasks

## To set up and configure prerequisite software for FileNet P8 Platform components

---

1. Configure FileNet P8 Platform authentication. Do one of the following:
  - [Task 1a on page 54](#) (Active Directory)
  - [Task 1b on page 55](#) (Active Directory Application Mode - ADAM)
  - [Task 1c on page 56](#) (Sun Java System Directory Server)
  - [Task 1d on page 58](#) (Novell eDirectory)
  - [Task 1e on page 59](#) (Tivoli Directory Server)
2. Create FileNet P8 Platform groups and users. Do [Task 2 on page 60](#).
3. Verify that a database server has been installed for Content Engine and Process Engine.
  - a. For Content Engine, do one of the following:
    - [Task 3a on page 81](#) (Microsoft SQL Server)
    - [Task 3b on page 85](#) (Oracle)
    - [Task 3c on page 92](#) (DB2)
  - b. For Process Engine, do one of the following:
    - [Task 3a on page 81](#) (Microsoft SQL Server)
    - [Task 3b on page 85](#) (Oracle)
    - [Task 3c on page 92](#) (DB2)
4. Verify that database client software has been installed on any computer that needs to directly access a database engine. Do one of the following:
  - [Task 4a on page 98](#) (Microsoft SQL Server)
  - [Task 4b on page 100](#) (Oracle)
  - [Task 4c on page 103](#) (DB2)
5. Configure your application server for Content Engine. Do one of the following:
  - [Task 5a on page 106](#) (WebSphere)
  - [Task 5b on page 110](#) (WebLogic)
  - [Task 5c on page 118](#) (JBoss)
6. Configure your application server for Application Engine. Do one of the following:
  - [Task 6a on page 123](#) (WebSphere)
  - [Task 6b on page 124](#) (WebLogic)
  - [Task 6c on page 126](#) (JBoss)

## Task 1a: Configure Windows Active Directory

You will use this task to configure Windows Active Directory on a Windows server.

For a complete list of IBM FileNet P8-supported Windows Active Directory features, refer to [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers > Windows Active Directory](#).

**NOTE** IBM FileNet P8 does not support Active Directory built-in user groups.

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

---

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

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

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

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

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

## Task 1b: Configure Active Directory Application Mode (ADAM)

No special IBM FileNet P8-specific settings are required for Microsoft ADAM on a Windows server.

However, if you intend to provision some or all of your user accounts in Microsoft Active Directory you must use ADAM's userProxyFull class of objects (and not the userProxy object) to represent those user accounts. Consult your ADAM documentation for full information, including how to use the Active Directory to ADAM synchronization tool. IBM recommends that you establish the Active Directory to ADAM connection before installing IBM FileNet P8.

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

## Task 1c: Configure Sun Java System Directory Server

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

This task assumes that Sun Java System Directory Server is installed properly. Be sure to read all pertinent documentation from the third-party software vendor.

### NOTES

This step is performed on the Windows or supported UNIX computer that will be configured as the IBM FileNet P8 Platform authentication server using a Sun Java System Directory Server (v 5.1 SP2 or 5.2).

- When configuring the Sun Java System Directory Server software, you may find references in the user interface to iPlanet or Sun ONE. These terms have been replaced by Sun Java System Directory Server.
- On Windows servers, Sun Java System Directory Server should be installed on an NTFS hard drive partition.
- If there are more than 2,000 users in the Directory Server, you must increase the resource limits to correctly display users in IBM FileNet P8. IBM recommends setting this limit to -1 (unlimited). You can either set this limit for the entire LDAP server or for the individual IBM FileNet P8 users. Instructions for either procedure appear below.
- You must have at least one group and/or user account to create Content Engine object stores (see [“Specify IBM FileNet P8 Accounts” on page 60](#)).

### Configure Sun Java System Directory Server (v 5.1 SP2)

#### To set the resource limits for the entire directory server

---

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

1. From the server where Sun Java System Directory Server is installed, log on with an account that has rights to modify the Sun Java System Directory Server environment.
2. Run the Sun Java System Directory Server console and login.
3. Expand the *Domain > Server Group* containers and select your **Directory Server**. Then right-click and select **Open**.
4. Select the Directory tab and expand config > plugins > ldbm database.
5. Double-click the **config** folder.
6. From the Property Editor sheet for cn=config,cn=ldbm database,cn=plugins,cn=config, change nsslapd-lookthroughlimit value to -1 and click **OK**.
7. Select the Configuration tab.



8. Select the Performance tab, change the *Size limit* to -1 and click **Save**.
9. Select the **Tasks** tab and click **Restart** to restart the Directory Server.

**To set the resource limits for individual IBM FileNet P8 users**

---

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

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

**Configure Sun Java System Directory Server (v 5.2)**

**To set the resource limits for the entire 5.2 Directory Server**

---

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

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

## Task 1d: Configure Novell eDirectory

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

### NOTES

- This task assumes that Novell eDirectory Server and Client components are properly installed.
- IBM FileNet P8 Platform does not support global referrals because eDirectory does not provide a referral error to chase for LDAP binds.
- The Windows server where Novell eDirectory Server is installed, must have an NTFS hard drive partition.
- You must have at least one group and/or user account to create Content Engine object stores (see [“Specify IBM FileNet P8 Accounts” on page 60](#)).
- The Novell eDirectory administrator may have to create an index if the sorting attribute is not in the list of default attributes shipped by eDirectory.
- Novell eDirectory Server does not support the TLS protocol for authorization (authentication is the responsibility of the application server where Content Engine is deployed).

## Task 1e: Configure IBM Tivoli Directory Server

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

## Task 2: Specify IBM FileNet P8 Accounts

This task assumes that you have a directory server that is properly installed and configured.

The following procedures create or designate the accounts needed to install and configure IBM FileNet P8. Some IBM FileNet P8 roles are not fully described here because the information is not essential for completing the installation procedures.

### NOTES

- For a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain an IBM FileNet P8 system, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups](#).
- In a multi-domain Microsoft Active Directory environment, a logon will fail for any account whose user name and password in a parent/child domain match those in a child/parent domain.

### To create Content Search Engine accounts

---

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

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
K2 Security Group	Autonomy K2 security group used to secure K2 collections. You will specify this group in the User Group field in the Verity Domain Configuration when you configure CBR in Enterprise Manager.	Install: <a href="#">Task 8</a> Install: <a href="#">Task 29</a>	New for 4.0.0	

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
K2 Security User	Autonomy K2 security user account, used when logging on to perform Content-Based Retrieval (CBR). You will specify this account in the Verity Username field in the Verity Domain Configuration when you configure CBR in Enterprise Manager. This user must be a member of the K2 Security Group and must be defined as an authorized K2 administrator in the K2 dashboard.	Install: <a href="#">Task 8</a>  Install: <a href="#">Task 29</a>	New for 4.0.0	_____
K2 Operating System User	Autonomy K2 services will run as this user. This user must be an operating system administrator on the machine where the Autonomy K2 Master Administration server is installed. Additionally, this user must have access to the file system that contains the file storage areas and the full text index collections, because the Autonomy K2 software needs to read the file storage areas and write the full text index collections as part of the full text indexing operation.	Install: <a href="#">Task 8</a>  Install: <a href="#">Task 29</a>  Install: <a href="#">Task 40</a>  Install: <a href="#">Task 46</a>	New for 4.0.0	_____

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

**To create Content Engine accounts**

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

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
<p>&lt;SQL_Server_login&gt;</p> <p><b>NOTE</b> This account applies only if the Content Engine database is SQL Server.</p>	<p>SQL Server account for creating databases and unique database users for object stores and the GCD. The account must have at least the following server roles:</p> <ul style="list-style-type: none"> <li>• System Administrators</li> <li>• Security Administrators</li> <li>• Disk Administrators</li> <li>• Database Creators</li> </ul> <p>and at least the following database access permissions:</p> <ul style="list-style-type: none"> <li>• public</li> <li>• db_owner</li> </ul> <p>Also add this account to SQL Server's master database and grant the SqlJDBCXAUser role and the public role.</p>	<p>Install: <a href="#">Task 9</a></p> <p>Install: <a href="#">Task 14a</a> (Web Sphere 5.1.x)</p> <p>Install: <a href="#">Task 14b</a> (Web Sphere 6.0.x)</p> <p>Install: <a href="#">Task 14c</a> (Web Sphere 6.1.x)</p> <p>Install: <a href="#">Task 14d</a> (Web Logic 8.1.x)</p> <p>Install: <a href="#">Task 14e</a> (Web Logic 9.2.x)</p> <p>Install: <a href="#">Task 16</a></p> <p>Upgrade: <a href="#">Task 3</a></p>	<p>New requirement for access to the master database.</p>	

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
<p>&lt;Oracle_alias&gt;</p> <p><b>NOTE</b> This account applies only if the Content Engine database is Oracle.</p>	<p>The account you will use to create tablespaces and unique tablespace users for object stores and the GCD. Give this account at least the following permissions:</p> <ul style="list-style-type: none"> <li>• CREATE SESSION</li> <li>• CREATE TABLE</li> <li>• CREATE SEQUENCE</li> </ul>	<p>Install: <a href="#">Task 9</a></p> <p>Install: <a href="#">Task 14a</a> (Web Sphere 5.1.x)</p> <p>Install: <a href="#">Task 14b</a> (Web Sphere 6.0.x)</p> <p>Install: <a href="#">Task 14c</a> (Web Sphere 6.1.x)</p> <p>Install: <a href="#">Task 14d</a> (Web Logic 8.1.x)</p> <p>Install: <a href="#">Task 14e</a> (Web Logic 9.2.x)</p> <p>Install: <a href="#">Task 16</a></p> <p>Upgrade: <a href="#">Task 3</a></p>		<hr/>

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
<p>&lt;DB2 user&gt;</p> <p><b>NOTE</b> Required if the Content Engine database is DB2.</p>	<p>(DB2 only) An operating system user on the database server with the following DB2 database permissions. In the case of a remote database, no equivalent users are needed on the Content Engine server.</p> <ul style="list-style-type: none"> <li>• Connect to the database</li> <li>• Create tables in the tablespace (CREATETAB)</li> <li>• Use the tablespace (USE OF) for User and User Temp tablespaces</li> </ul>	<p>Install: <a href="#">Task 3c</a></p> <p>Install: <a href="#">Task 9</a></p> <p>Install: <a href="#">Task 14a</a> (Web Sphere 5.1.x)</p> <p>Install: <a href="#">Task 14b</a> (Web Sphere 6.0.x)</p> <p>Install: <a href="#">Task 14c</a> (Web Sphere 6.1.x)</p> <p>Install: <a href="#">Task 14d</a> (Web Logic 8.1.x)</p> <p>Install: <a href="#">Task 14e</a> (Web Logic 9.2.x)</p> <p>Install: <a href="#">Task 16</a></p> <p>Upgrade: <a href="#">Task 3</a></p>	<p>No change</p>	<hr/>



User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
<p>Accounts required for DB2 database setup and administration:</p> <ul style="list-style-type: none"> <li>• Groups: instance owner primary group</li> <li>• Users: instance owner</li> </ul>	<p>(DB2 only) Operating system user and group that must exist on the database server.</p>	<p>Install: <a href="#">Task 3c</a></p> <p>Install: <a href="#">Task 16</a></p> <p>Upgrade: <a href="#">Task 3</a></p>	<p>No change</p>	<hr/>

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Application server administrator	<p>You will use this account to</p> <ul style="list-style-type: none"> <li>• Create and configure the application server domain/profile for IBM FileNet P8</li> <li>• Start the application server before launching Content Engine Setup</li> <li>• Specify the username and password to Content Engine Setup to log on to the application server</li> <li>• Deploy the EAR file during initial Content Engine installation and when deploying Content Engine to other servers</li> </ul> <p>Content Engine Setup will use this account to configure JDBC data sources and connection pools for the GCD.</p>	<p>Install: <a href="#">Task 5a</a> (Web Sphere)</p> <p>Install: <a href="#">Task 5b</a> (Web Logic)</p> <p>Install: <a href="#">Task 5c</a> (JBoss)</p> <p>Install: <a href="#">Task 9</a></p> <p>Install: <a href="#">Task 14a</a> (Web Sphere 5.1.x)</p> <p>Install: <a href="#">Task 14b</a> (Web Sphere 6.0.x)</p> <p>Install: <a href="#">Task 14c</a> (Web Sphere 6.1.x)</p> <p>Install: <a href="#">Task 14d</a> (Web Logic 8.1.x)</p> <p>Install: <a href="#">Task 14e</a> (Web Logic 9.2.x)</p>	New for 4.0.0	

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Content Engine Setup account (Windows)	The account you will use to log on to a machine to launch Content Engine Setup. This account must have at least Local Administrator permissions on the machine.	Install: <a href="#">Task 9</a>	No longer needs to be member of Domain Users.	_____
Content Engine Setup account (UNIX)	The account you will use to log on to a machine to launch Content Engine Setup. This account must have read/write/execute permissions on the IBM FileNet, application server instance, and temp directories.  <b>NOTE</b> Only the root account can run Content Engine Setup on AIX® machines.	Install: <a href="#">Task 9</a>	New for 4.0.0	_____

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Content Engine system user	<p>An application server administrative account that is stored in the CEMPBoot.properties file. This account is used first to create the GCD, and thereafter is the account that Content Engine runs as. It is the account Content Engine will use to establish a connection with the application server, access the application server's JNDI tree, and look up the data sources for accessing the GCD.</p> <p><b>WARNING</b> If you are deploying Content Engine on WebSphere 6.1.x with federated repositories and with multiple realms in your P8 domain, be sure that no two realms contain the same short name for this user; otherwise, this user will not be able to create the GCD.</p>	Install: <a href="#">Task 9</a>	New for 4.0.0	
Content Engine operating system user	The user under which Content Engine Server executes (typically, the user that starts Content Engine Server).	Install: <a href="#">Task 40</a>  Upgrade: <a href="#">Task 3</a>		

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
GCD administrators	<p>The accounts that serve in the role of GCD administrator. Use one of these accounts to:</p> <ul style="list-style-type: none"> <li>• Create the GCD</li> <li>• Launch the Configure New Domain Permissions wizard the first time you start Enterprise Manager (see <a href="#">“To configure permissions for a FileNet P8 domain” on page 191</a>).</li> </ul>	<p>Install: <a href="#">Task 9</a></p> <p>Install: <a href="#">Task 16</a></p>	No change	_____
Object store administrators	<p>Each account will administer an object store and will have Full Control access to it. You will need to specify these accounts when you run the Create an Object Store wizard in Enterprise Manager (see <a href="#">“To create an object store” on page 259</a>).</p>	<p>Install: <a href="#">Task 16</a></p>	No change	_____

2. Create a new (or designate an existing) directory server account, which Content Engine Setup will prompt you for, to use when connecting to and searching within the directory server. This account must have at least the following permissions:

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Directory service user (Content Engine/Windows)	Used by Content Engine to connect to Active Directory. Requires at least the following permissions:  Member of the Pre-Windows 2000 Compatible Access Group in each desired domain in the Active Directory forest.	Install: <a href="#">Task 1a</a>  Install: <a href="#">Task 16</a>	New for 4.0.0	_____
Directory service user (Content Engine/ADAM)	Used by Content Engine to connect to a single Microsoft ADAM partition. Requires at least the following permissions:  Ability to see the other users in the partition. (For a procedure, see the entry for the ADAM directory service user in <a href="#">FileNet P8 Administration &gt; Enterprise-wide Administration &gt; FileNet P8 Security &gt; Users and groups.</a> )	Install: <a href="#">Task 1b</a>  Install: <a href="#">Task 16</a>	New for Content Engine 4.0.1	_____
Directory service user (Content Engine/Sun)	Used by Content Engine to connect to the Sun Java System directory server. Requires at least the following permissions:  Read, Search, Compare	Install: <a href="#">Task 1c</a>  Install: <a href="#">Task 16</a>	No change	_____

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Directory service user (Content Engine/Novell)	Used by Content Engine to connect to Novell eDirectory. Requires at least the following permissions:  Read, Compare	Install: <a href="#">Task 1d</a>  Install: <a href="#">Task 16</a>	No change	_____
Directory service user (Content Engine/IBM)	Used by Content Engine to connect to IBM Tivoli Directory Server. Requires at least the following permissions:  Read, Search, Compare	Install: <a href="#">Task 1e</a>  Install: <a href="#">Task 16</a>	New for 4.0.0	_____

**To create Process Engine accounts (Windows)**

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

There are two options:

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

Process Engine Setup will prompt asking if you want to configure aliases:

- A no answer indicates that Process Engine Setup should create the default accounts and groups.
- A yes answer brings up a second screen allowing you to define aliases for each of the default accounts and groups.

If the accounts and groups do not already exist, Process Engine Setup will create them. If the accounts and groups already exist, Process Engine Setup will use them during the install. If you are pre-defining the users and groups, be sure to assign the fnsu user alias to the aliased fnadmin, fnusr, and fnop groups before running Process Engine Setup, otherwise Process Engine Setup may encounter errors.

If you choose to create an alias for one Process Engine account, you must create aliases for all accounts, even if you set the alias to the default value of a Process Engine account.

If the database is DB2, the f\_sw and f\_maint accounts, or their aliases, must be predefined as OS users. See [“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92](#) for additional information.

1. Optionally assign aliases to the following users and groups:

User/ Group	Description	Required by	Change status from 3.5.x	Name/Alias
fnadmin (group account)	Operating system group whose members have all privileges on Process Engine files and databases	Install: <a href="#">Task 19a</a>  Install: <a href="#">Task 22</a>	No change	_____
fnusr (group account)	Operating system group whose members have non- administrator privileges on Process Engine files and databases	Install: <a href="#">Task 19a</a>	No change	_____
fnop (group account)	Operating system group whose members have operator non-administrator privileges on Image Services utilities	Install: <a href="#">Task 19a</a>	No change	_____
fnsu (user account)	Operating system user account who executes Process Engine software	Install: <a href="#">Task 4b</a>  Install: <a href="#">Task 19a</a>  Install: <a href="#">Task 22</a>  Upgrade: <a href="#">Task 8b</a>  Upgrade: <a href="#">Task 11</a>	No change	_____



User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
f_sw (user account)	Database runtime user (OS user of DB2)	Install: <a href="#">Task 3c</a>  Install: <a href="#">Task 4c</a>  Appendix: "Process Engine SQL Scripts"	No change	
f_maint (user account)	Database maintenance user (OS user of DB2)	Install: <a href="#">Task 3c</a>  Install: <a href="#">Task 4c</a>  Appendix: "Process Engine SQL Scripts"	No change	

**NOTE** The f\_sw and f\_maint database users will be created if the SQL scripts for SQL Server and Oracle databases are manually executed before Process Engine Setup. In this case, Process Engine Setup will only define the default IBM FileNet passwords for these users, or their aliases. See "[Process Engine SQL Scripts](#)" on page 647 for details on the SQL scripts.

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

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

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

Default Name	Description
<peinstaller>	User who will run Process Engine Setup
fnsf or alias	Primary IBM FileNet software user  <b>NOTE</b> Set the password to BPMtemp1pzwd. This password is case-sensitive. You can change the password after installing Process Engine.

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

#### **To create Process Engine accounts (UNIX)**

---

On UNIX platforms, Process Engine accounts, or their aliases, must always be created before installing the Process Engine software. If the database is DB2, the f\_sw and f\_maint accounts, or their aliases, must be created as OS users. See [“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92](#) for additional information.

**NOTE** The f\_sw and f\_maint database users will be created if the SQL scripts for SQL Server and Oracle databases are manually executed before Process Engine Setup. In this case, Process Engine Setup will only define the default IBM FileNet passwords for these users or their aliases. See [“Process Engine SQL Scripts” on page 647](#) for details on the SQL scripts.

See the appropriate Process Engine tasks in this guide for specific information on how to create these accounts on your operating system.

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
fnadmin (group account)	Operating system group account whose members have all privileges on Process Engine files and databases	Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)	No change	_____
fnusr (group account)	Operating system group account whose members have non-administrative privileges on Process Engine files and databases	Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)	No change	_____
fnop (group account)	Operating system group account whose members have operator non-administrator privileges on Image Services utilities	Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)	No change	_____

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
fns (user account)	Operating system user account who executes Process Engine software	Install: <a href="#">Task 4b</a>  Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)  Install: <a href="#">Task 22</a>  Upgrade: <a href="#">Task 8a</a>  Upgrade: <a href="#">Task 11</a>	No change	_____
f_sw (user account)	Database runtime user (OS user of DB2)	Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)  Upgrade: <a href="#">Task 11</a>	No change	_____
f_maint (user account)	Database maintenance user (OS user of DB2)	Install: <a href="#">Task 19b</a> (Solaris)  Install: <a href="#">Task 19c</a> (AIX)  Install: <a href="#">Task 19d</a> (HP-UX)	No change	_____

**To create Oracle accounts for Process Engine (UNIX)**

If the Oracle database is local to Process Engine, the following operating system user and group should already exist as a result of installing the Oracle software. Process Engine Setup will prompt for the values. If the database is remote, the user and group may need to be defined on the server where Process Engine will be installed. When the Oracle user and group have been created, modify the account information as described.

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

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Oracle Database Administrators Group; default = <i>&lt;dba&gt;</i>	Members act as database administrators  Action: Modify, if database is local. Create, if database is remote.  Members: fnsr, Oracle user; default = <i>&lt;oracle&gt;</i>	Install: <a href="#">Task 19b</a>  Install: <a href="#">Task 19c</a>  Install: <a href="#">Task 19d</a>	No change	

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Oracle User; default = <i>&lt;oracle&gt;</i>	Action: Modify, if database is local. Create, if database is remote.  User Type: Database  Primary Group: Oracle Database Administrators Group; default = <i>&lt;dba&gt;</i>  Secondary Group: fnsr	Install: <a href="#">Task 19b</a>  Install: <a href="#">Task 19c</a>  Install: <a href="#">Task 19d</a>	No change	

**To create other Process Engine accounts**

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

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Process Engine service user	Process Engine uses the service username when connecting the Content Engine server. This user must belong to the Process Engine Administrator group	Install: <a href="#">Task 22</a>	No longer used to communicate directly to the directory server.	_____
Process Engine administrators group	Members of this group automatically have administrative privileges for Process Engine.	Install: <a href="#">Task 22</a>	No change	_____
Process Engine configuration group	(Optional) A valid group name. Members of this group automatically have configuration privileges for the PE workflow database.  If this group is used to configure security on Process Task Manager, members of this group or of the Process Engine Administrator Group can make configuration changes to the workflow database. If the Process Engine Configuration group is not used during this configuration, anyone can make these changes.	Install: <a href="#">Task 22</a>	No change	_____

**To create Application Engine accounts**

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

User/Group	Description	Required by	Change status from 3.5.x	Name/Alias
Application Engine setup account (Windows)	The account you will use to log on to a Windows machine and launch Application Engine Setup. This account must have at least local administrator permissions on the machine.	Install: <a href="#">Task 24</a>	No change	_____
Application Engine setup account (UNIX)	The account you will use to log on to a UNIX machine and launch Application Engine Setup. This account must have write access to the <b>/bin</b> directory and read, write, and execute access to the directory where you will install Application Engine.	Install: <a href="#">Task 24</a>	No change	_____
Application server account	This account will have permissions to deploy an application. The account may be the same as the Application Engine setup account.	Install: <a href="#">Task 25a</a>  Install: <a href="#">Task 25b</a>		_____
Application Engine Administrators	These accounts will serve in the role of Application Engine administrator. You will specify these accounts as members of the AE administrator role in " <a href="#">Set Application Engine Bootstrap Preferences</a> " on <a href="#">page 403</a> . These accounts must have passwords.	Install: <a href="#">Task 30</a>  Install: <a href="#">Task 35</a>  Install: <a href="#">Task 47</a>  Upgrade: <a href="#">Task 19</a>	No change	_____

**NOTE** All IBM FileNet Workplace accounts, as well as accounts for other client applications and expansion products that use Content Engine or Application Engine, must have passwords.



## Task 3a: Verify that Microsoft SQL Server Is Installed for IBM FileNet P8

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

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

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

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

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

If you are installing IBM FileNet P8 as a *new* application (that is, *not* upgrading from version 3.5.x, do all the procedures in this task. If you are upgrading Content Engine 3.5.x to version 4.0.x, do *only* the following procedures in this task:

- [“To create a Microsoft SQL database for the GCD” on page 83](#)
- [“To enable XA Transactions” on page 84](#)
- [“To configure the JDBC Distributed Transaction Components” on page 84](#)
- [“To configure user-defined roles” on page 84](#)

### NOTES

- If your Microsoft SQL Server database will be remote, then after completing the procedures in this topic you might also need to complete the procedures in [“Verify that Microsoft SQL Server Client Is Installed for IBM FileNet P8” on page 98](#) on the machine where you will install Process Engine.
- The following components can optionally share a MS SQL Server instance:
  - Content Engine (object store databases and GCD database)
  - Process Engine
  - Rendition Engine
- A database should not be shared by multiple Content Engine object stores.
- Process Engine and the optional Process Analyzer cannot share an instance. For information on SQL Server requirements for Process Analyzer Engine, see the *Process Analyzer Installation and Upgrade Guide*.

- Record the values for the following settings as you work through the database installation. This information must be entered during subsequent installations of IBM FileNet P8 components. Be aware that Process Engine Setup allows only alphanumeric characters and underscores.
  - Server name
  - Instance name (for example, P8\_inst)
  - Dedicated database name (for example, VWdb)
  - Dedicated file group (for example, vwdata\_fg)
  - TCP/IP port number assigned

### **To install and configure Microsoft SQL Server**

---

**CAUTION** Do this procedure *only if* you are installing version 4.0.x of IBM FileNet P8 as a new application. If you are upgrading from version 3.5.x, skip to [“To create a Microsoft SQL database for the GCD” on page 83](#).

- Create a database instance for use by IBM FileNet P8 software, or verify that such an instance already exists.
- If creating a new instance, indicate an appropriate name based on whether Content Engine (object store), GCD, Process Engine, or Rendition Engine will use the instance. Be aware of the following rules for instance names:
  - The name cannot exceed 16 characters.
  - The first character cannot be numeric or '\$'.
  - The name cannot contain special characters, except underscores or periods.
  - The name cannot contain spaces.
  - The name cannot be Default or MSSQLServer.
- Verify the authentication mode you specify is for Mixed Mode.
- Select the database collation settings. Specify one of the following:
  - **Dictionary order, case-insensitive, for use with 1252 Character Set** (or any case-insensitive MS SQL Server collation). Case-insensitive collation is the Microsoft default and the setting most used in IBM FileNet P8 environments (because it offers search results without regard to character case).
  - **Dictionary order, case-sensitive, for use with 1252 Character Set** (or any case-sensitive MS SQL Server collation). Select case-sensitive MS SQL Server collation only if you are sure your site actually requires (and will continue to require) searches that must differentiate upper-case from lower-case characters (in property choice lists, folder names, etc.).

**CAUTION** Select your MS SQL Server collation setting carefully. Switching collation settings after installation can be difficult and time-consuming, especially if you want to switch from case-sensitive to case-insensitive collation after significant user activity. Also, be aware that if you have a case-sensitive database, and you want to perform a case-insensitive search

(programmatically or otherwise), you will likely encounter serious performance degradation on MS SQL Server because the database cannot use column (that is, property) indexes in these cases.

- (SQL Server 2000 only) Assign a TCP/IP port number for the instance. If your instance is the only instance on the server, you can accept port 1433 (the default value). If you are creating a named instance you must assign a static port number.
- (SQL Server 2005 only) The TCP/IP port number cannot be assigned during installation of SQL Server 2005. You must assign the port number after installation is complete by using the SQL Server Configuration Manager application to modify the network configuration.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. Verify that the required service pack has been installed before proceeding. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- If you want to disable the Named Pipes protocol for the database instance to be used by Process Engine, wait until after Process Engine installation and configuration is complete. Disabling this protocol too early might cause Process Engine initialization to fail.

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

#### **To create a Microsoft SQL database for the GCD**

---

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

#### **To create a Microsoft SQL database for Content Engine object store(s)**

---

**CAUTION** Do this procedure *only if* you are installing version 4.0.x of IBM FileNet P8 as a new application. If you are upgrading from version 3.5.x, skip to [“To enable XA Transactions” on page 84](#).

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

#### **To create the Process Engine database**

---

Create a SQL Server database for Process Engine. The default name assigned by the Process Engine Setup program is VWdb. Create the database with an initial size of 200MB, minimum. Assign a new file name for the database (for example, vw\_data). Specify a filegroup (for example, vwdata\_fg.) IBM recommends that the Primary filegroup is not used.

**NOTE** You will need the database and filegroup names for the Process Engine installation.

#### **To modify the tempdb database for Process Engine**

---

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

### To enable XA Transactions

---

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

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

### To configure the JDBC Distributed Transaction Components

---

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

1. Download the Microsoft SQL Server 2005 JDBC Driver that is referenced in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
2. Copy the **sqljdbc\_xa.dll** from the JDBC installation directory to the **Program Files\Microsoft SQL Server\80\Tools\Binn** directory if there is only a default instance or the **Program Files\Microsoft SQL Server\MSSQL\$<instance name>Binn** if you are using a named instance. If you are on a 32-bit processor, use the **sqljdbc\_xa.dll** file in the x86 folder. If you are on a 64-bit processor, use the **sqljdbc\_xa.dll** file in the x64 folder.
3. Log on as a database administrator and execute the database script **xa\_install.sql** on every SQL Server instance that will participate in distributed transactions. This script installs **sqljdbc\_xa.dll** as an extended stored procedure and creates the SqlJDBCXAUser role in the Master database.

**CAUTION** Use SQL Server database credentials, not Windows credentials, to log on. Windows Integrated Logon to SQL Server is not supported with IBM FileNet P8.

### To configure user-defined roles

---

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

To grant permissions to a specific user to participate in distributed transactions with the JDBC driver, add the user to the SqlJDBCXAUser role. The user must be assigned to the master database.

Refer to [“Specify IBM FileNet P8 Accounts” on page 60](#) for more information on the Content Engine SQL Server user requirements.

**NOTE** If the SQL Server login you are using for Content Engine is the default sa user, these steps are not necessary.

If you are upgrading from version 3.5.x of IBM FileNet P8, continue at [“To install and configure an application server for Content Engine Server” on page 490](#) in [“Upgrade Content Engine Software” on page 487](#).

## Task 3b: Verify that Oracle Server Is Installed for IBM FileNet P8

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

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

In the shared configuration, the IBM FileNet P8 components use the same database, but different tablespaces. You can also share the database with other (non-IBM FileNet) applications. In the dedicated configuration, Content Engine, Process Engine, and Rendition Engine use separate databases.

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

If you are installing IBM FileNet P8 as a *new* application (that is, *not* upgrading from version 3.5.x, do all the procedures in this task. If you are upgrading Content Engine 3.5.x to version 4.0.0, do *only* the procedure “[To create tablespaces for the GCD](#)” on page 87.

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

### NOTES

- If your Oracle database will be remote, then after completing the procedures in this topic you must also complete the procedures in “[Verify that Oracle Client Is Installed for IBM FileNet P8](#)” on [page 100](#) on each machine where you will install Content Engine or Process Engine. For Rendition Engine instructions, see the IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).
- Make sure the machine that will host the database satisfies all pre-installation requirements specified in the Oracle9i or Oracle 10g installation documentation.
- For Content Engine and Process Engine, IBM FileNet P8 supports the Oracle Advanced Security functionality of secure data transfer across network protocol boundaries.
- If you will be installing Process Engine on a UNIX machine hosting the Oracle database, be sure that the value of the Oracle environment variable ORACLE\_HOME (the path name for the Oracle Server software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installer will not find the Oracle software, causing the installation to fail.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. To download this guide from the IBM support page, see “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on

[page 21](#). The Oracle patches are available at [OracleMetalink](#). The Oracle patch-installation procedure may be less complicated if done before you create any databases.

- Transaction Processing is the required configuration type for the database that supports Content Engine. Choose this configuration type if your database will be dedicated to Content Engine or shared with Process Engine.
- IBM FileNet suggests that you record the values for the following settings as you work through the database installation. This information must be entered during subsequent installations. Be aware that Process Engine Setup allows only alphanumeric characters and underscores.
  - Oracle Home
  - Global Database Name
  - Oracle temporary tablespace name
  - Oracle data tablespace name
  - Oracle index tablespace name (optional)
  - Oracle SID

### To install an Oracle database engine

---

**CAUTION** Do this procedure *only if* you are installing version 4.0.0 of IBM FileNet P8 as a new application. If you are upgrading from version 3.5.x, skip to “[To create tablespaces for the GCD](#)” on [page 87](#).

The following procedure shows the minimal choices (specific to the needs of Content Engine and Process Engine) for installing a database engine. Consult the Oracle9i or Oracle 10g installation documentation for complete preinstallation requirements and instructions.

- For Oracle 9i, choose the following from the list of available product components.
  - Oracle9i Server
  - Oracle Net Services
    - Oracle Net Listener
  - Oracle9i Development Kit
    - Oracle Call Interface (OCI)
  - (Windows) Oracle Windows Interfaces
    - Oracle Services for Microsoft Transaction Server
  - Oracle9i Documentation (recommended)
- For Oracle 10g, choose the following from the list of available product components.
  - Oracle10g Server
  - Oracle Net Services
    - Oracle Net Listener

- Oracle Call Interface (OCI)
- (Windows) Oracle Windows Interfaces
  - Oracle Services for Microsoft Transaction Server
- Oracle10g Documentation (recommended)
- If you are going to install Process Engine on this machine, verify/add/edit/uncomment the following lines in the file **sqlnet.ora** (create the file if it doesn't exist) while the Oracle services/processes are stopped:

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

**NOTE** If Oracle is configured to use LDAP, TNSNAMES must appear in the names\_directory\_path ahead of LDAPNAMES.

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

1. Install the latest Oracle patch sets, as specified in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
2. Start the listener and the Oracle database service/processes (Windows/UNIX) if they haven't started automatically.

### To create an Oracle database

---

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

IBM FileNet requires the following settings:

- Database configuration type
 

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

Dedicated Server Mode
- Database character set
 

Choose a database character set as specified in [“IBM FileNet P8 Database Character Sets” on page 653](#).

### To create tablespaces for the GCD

---

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



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

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

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

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

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

Tablespace Name	Tablespace Type	Minimum Size (MB)	Description
<gcd>	Permanent	100	Permanent tablespace for the GCD
<tempgcd>	Temporary	200	Temporary tablespace for the GCD

If you are upgrading from version 3.5.x of FileNet P8, continue at [“To install and configure an application server for Content Engine Server” on page 490](#) in [“Upgrade Content Engine Software” on page 487](#).

**To create Oracle tablespaces for Content Engine object stores**

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

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

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

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



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

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

Tablespace Name	Tablespace Type	Minimum Size (MB)	Description
<objectstore1>	Permanent	200	Permanent tablespace for object store
<tempobjectstore1>	Temporary	400	Temporary tablespace for object store

### To create tablespaces for Process Engine

---

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

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

Tablespace Name	Tablespace Type	Minimum Size (MB)	Description
vwdata_ts	Permanent	200	Default name of the dedicated IBM FileNet default tablespace
vwtemp_ts	Temporary	400	Default name of the dedicated IBM FileNet temporary tablespace
vwindex_ts (optional)	Permanent	200	Default name of the optional index tablespace

**NOTE** If you don't create vwindex\_ts, vwdata\_ts will be used for indexes.

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

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

- For the oracle user, set the following:
  - ORACLE\_SID
  - ORACLE\_HOME
  - Set PATH to:
    - \$ORACLE\_HOME/bin** (HP-UX and 32-bit Solaris only)
  - Set LD\_LIBRARY\_PATH to:
    - \$ORACLE\_HOME/lib** (HP-UX and 32-bit Solaris only)
  - Set LD\_LIBRARY\_PATH to:
    - \$ORACLE\_HOME/lib32** (64-bit Solaris only)
  - Set LIBPATH to:
    - \$ORACLE\_HOME/lib32:\$ORACLE\_HOME/lib** (AIX only)
  - Set SHLIB\_PATH to:
    - \$ORACLE\_HOME/lib32** (HP-UX only)
- For the root user, set ORACLE\_HOME.

### To configure automatic transaction recovery

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

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

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

### To edit the listener.ora file

IBM recommends that you use the Oracle utilities to configure the listener, which creates the **listener.ora** file. To manually update the **listener.ora** file (for example, to change the database

name), refer to the following example, where *sunblock* is the dbname and *BLK* is the ORACLE\_SID.

```
LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP)(HOST = vwnetra1.filenet.com)(PORT = 1521))
      )
    )
  )
)

SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = sunblock)
      (ORACLE_HOME = C:\oracle\ora92)
      (SID_NAME = BLK)
    )
  )
)
```

## Task 3c: Verify that DB2 Server Is Installed for IBM FileNet P8

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

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

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

If you are installing IBM FileNet P8 as a *new* application (that is, *not* upgrading from version 3.5.x), do all the procedures in this task. If you are upgrading Content Engine 3.5.x to version 4.0.0, do *only* the procedure to create a GCD database as documented in [“To create the DB2 tablespaces” on page 96](#).

This task uses 3rd-party software. Be sure to read all pertinent documentation from the 3rd-party software vendor. Contact your system/network/database administrator with 3rd-party vendor or environment questions before continuing. When this procedure has been completed, proceed to [“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103](#).

### NOTES

- Only a remote DB2 database is supported for Process Engine.
- On UNIX Process Engine servers you must install DB2 version 8 client software, regardless of whether the remote database server is running DB2 version 8 or 9.
- Content Engine and Process Engine should not share a tablespace but can share a database. For maintenance and support reasons, IBM recommends separate databases.
- The Content Engine GCD and every object store must have dedicated databases.
- Content Engine and Process Engine have different requirements for DB2 users and groups.
- See [“IBM FileNet P8 Database Character Sets” on page 653](#) for information on character set requirements.

IBM recommends using Database Managed Space (DMS) for user and user temp tablespaces for both Content Engine and Process Engine.

IBM recommends that you record the values for the following settings as you work through the database installation. This information must be entered during subsequent installations of Process Engine and Content Engine. Be aware that Process Engine Setup allows only alphanumeric characters and underscores.

- DB2 Server name

- DB2 server database instance name(s) (e.g. P8\_inst)
- Content Engine and Process Engine dedicated database names (e.g. VWdb)
- Dedicated tablespace names (e.g. vwdata)
- DB2 instance port numbers
- Process Engine runtime database user (f\_sw or alias) password (Process Engine only)
- Process Engine maintenance database user (f\_maint or alias) password (Process Engine only)
- User ID and password for Content Engine DB2 user

### **To create DB2 users and groups**

---

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

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

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

- Fenced User
- Fenced User primary group

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

Each instance should have its own home file system.

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

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

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

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

When the database authentication type is set to SERVER or SERVER\_ENCRYPT, two additional operating system users must be created for Process Engine on the database server where the

DB2 database resides. See [“To create additional OS users and groups \(Process Engine, SERVER or SERVER\\_ENCRYPT authentication only\)”](#) on page 96.

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

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

---

### To install DB2 ESE

---

As **root** user, create a temporary file system with 2.0 GB of free space to contain the tar.Z and the uncompressed installation file.

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

When the installation is finished, view the status report or go to **/tmp** to view all DB2 install logs to ensure there are no errors.

Make note of the TCP/IP port number assigned to the instance or instances, as the port number will be needed during the DB2 client configuration steps. The port number assigned can be found in the **/etc/services** file, associated with the DB2 instance(s) just created.

After a successful installation, the DB2 instance should be up and running. Continue with the next section.

---

### To create DB2 instance(s)

---

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

---

### To set TCP/IP as the default protocol

---

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

```
db2set DB2COMM=tcPIP
```

---

### To determine page size and user fields

---

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

The important thing to remember is that **the total row length of all the fields (including system and user) cannot be larger than the page size**. The DB2 page size you select when you create the database must be large enough to hold at least one complete record. DB2 cannot retrieve a partial record or spread a single record onto two pages.

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

Log on as the database instance owner as defined earlier. At least three databases must be created, one for Process Engine, one for the Content Engine GCD, and one for a single Content Engine object store. Every object store requires a separate DB2 database.

The database name needs to be unique and from 1 to 8 characters long. For example, **VWdb** This is the default database name in the Process Engine installer.

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

```
APPLHEAPSZ 2560
```

**To create the DB2 tablespaces**

---

IBM FileNet Tablespaces	Actual Assigned Name	Minimum Size (MB)	Actual Created Size	Minimum Page Size (KB)
user temporary ts(for PE)		40		must match pagesize of PE data tablespace
system temporary (for PE)		40		must match pagesize of PE data tablespace
vwdata_ts (for PE)		200		8 (recommended)
GCD_ts (for the GCD database)		256		32 (required)
cedata_ts (for a single CE object store)		512		32 (required)
user temporary ts (for CE)		40		32 (required)
system temporary ts (for CE)		40		32 (required)

**To create additional OS users and groups (Process Engine, SERVER or SERVER\_ENCRYPT authentication only)**

---

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

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



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

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

After creating the new users and setting their group memberships, log off as **root** user, log on as each of the new users, and change the password to avoid connection problems the first time they're used.

### **To grant database permissions for f\_sw and f\_maint (Process Engine, CLIENT authentication only)**

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

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

## Task 4a: Verify that Microsoft SQL Server Client Is Installed for IBM FileNet P8

Only Process Engine has a possible requirement for SQL Server Client software installation. A number of SQL scripts must be executed for the SQL Server database to create stored procedures. Those scripts can either be executed manually on the SQL Server database before running Process Engine Setup, or they can be executed from Process Engine Setup. If the scripts are executed from Process Engine Setup, SQL Server Client software must be installed on the Process Engine server. If the scripts are executed manually on the remote database server before you install Process Engine, you need not install SQL Server Client. If SQL Server Client software is installed on the Process Engine Server, that software can be removed after successful installation and configuration of Process Engine software.

See [“Process Engine SQL Scripts” on page 647](#) for details on execution of the SQL Scripts.

### NOTES

- Remote SQL Server connections for Process Engine at run time are handled through an ODBC data source.
- You have the option of installing Microsoft SQL Server and Process Engine servers in different Active Directory forests. In this case, however, you must use SQL Authentication rather than Windows Authentication.

### To install MS SQL Server Client software for remote database access (optional for PE)

---

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

#### Install the MS SQL Server Client software (SQL Server 2000)

1. Log on with an account that has local administrator privileges on the computer where the MS SQL Server client software will be installed.
2. Install the SQL Server **Client Tools Only**.
3. When the installation is complete, start the **Client Network Utility** and clear the **Automatic ANSI to OEM conversion** on the *DB-Library Options* tab.
4. Test the database connection.

#### Install the MS SQL Server Client software (SQL Server 2005)

1. Log on with an account that has local administrator privileges on the computer where the MS SQL Server client software will be installed.
2. Install **Workstation components, Books Online and development tools** and from the Advanced options, select **Client Components**.

#### Install SQL Server patches and service packs

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs. To download this guide from the IBM support page,

see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs”](#) on page 21. Verify that the required service pack has been installed before proceeding.

## Task 4b: Verify that Oracle Client Is Installed for IBM FileNet P8

The procedure in this task describe how to install the Oracle Client software to prepare for the installation of Process Engine. For information regarding installation of Oracle Client and Rendition Engine, see IBM FileNet P8 guide [FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

### NOTES

- Install Oracle Client on any machine that will host Process Engine, Rendition Engine, or any other IBM FileNet P8 component except Content Engine (such as Enterprise Manager) that needs to access an Oracle database.
- Make sure that the machine where you will install Oracle Client satisfies all pre-installation requirements specified in the Oracle9i or Oracle 10g installation documentation.
- If you will be installing Process Engine on a UNIX machine hosting Oracle Client software, be sure that the value of the Oracle environment variable ORACLE\_HOME (the pathname for the Oracle Client software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installer will not find the Oracle software, causing the installation to fail.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and Service Packs. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#). The Oracle patches are available at [OracleMetalink](#). The Oracle patch-installation procedure may be less complicated if done before you create any databases.

### To install the Oracle client

---

The following procedure shows the minimal choices (specific to the needs of Process Engine) for installing a database client. Consult the Oracle9i or Oracle 10g installation documentation for complete preinstallation requirements and instructions. From the Available Products screen, click **Oracle9i Client**.

1. For Oracle 9i, choose the following from the list of available product components.
  - Oracle9i Client
  - Oracle Network Utilities
  - Oracle Database Utilities
  - SQL\*PLUS
  - (Windows) Oracle Windows Interfaces
    - Oracle Services for Microsoft Transaction Server
2. For Oracle 10g, choose the following from the list of available product components.
  - Oracle10g Client
  - Oracle Network Utilities

- Oracle Database Utilities
  - SQL\*PLUS
  - (Windows) Oracle Windows Interfaces
    - Oracle Services for Microsoft Transaction Server
3. Using Oracle Net Configuration Assistant, test the connection to the Oracle database server with an appropriate Oracle user and password.
  4. If you are going to install Process Engine software on this machine, and your remote Oracle database uses the Unicode character set AL32UTF8, then for each user who will access Process Engine software on the machine, set the value of the Oracle environment variable parameter NLS\_LANG to reflect the PE-supported locale and (non-Unicode) character set on the machine.
    - (Windows) Set/modify the value of the NLS\_LANG key via System Properties in the Control Panel.
    - (UNIX) Add NLS\_LANG to the shell environment login files for each user who will be logging onto the machine to run IBM FileNet P8 software.

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

5. If you are going to install Process Engine software on this machine to connect to a remote Oracle database, set the value of an environment variable for the *oracle* user to a default connect identifier, such as the Oracle net service name or the database service name. The name of the environment variable depends on which operating system is on this machine:
  - (Windows) LOCAL
  - (UNIX) TWO\_TASK
6. If you are going to install Process Engine software on a UNIX machine to connect to a remote Oracle database, set the following environment variables in the startup file of the default *oracle* user.
  - ORACLE\_SID
  - ORACLE\_HOME
  - Set PATH to:
    - **\$ORACLE\_HOME/bin** (HP-UX and 32-bit Solaris only)
  - Set LD\_LIBRARY\_PATH to:
    - **\$ORACLE\_HOME/lib** (HP-UX and 32-bit Solaris only)
  - Set LD\_LIBRARY\_PATH to:
    - **\$ORACLE\_HOME/lib32** (64-bit Solaris only)
  - Set LD\_LIBRARY\_PATH\_64 to:
    - **\$ORACLE\_HOME/lib** (64-bit Solaris only)

- Set LIBPATH to:

**\$ORACLE\_HOME/lib32:\$ORACLE\_HOME/lib** (AIX only)

- Set SHLIB\_PATH to:

**\$ORACLE\_HOME/lib32** (HP-UX only)

7. If you are going to install Process Engine on this machine, verify/add/edit/uncomment the following lines in the file **sqlnet.ora** (create the file if it doesn't exist) while the Oracle services/processes are stopped:

```
NAMES_DIRECTORY_PATH=(TNSNAMES)
SQLNET.AUTHENTICATION_SERVICES=(NTS)
```

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

8. Install all required Oracle patches, as specified in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#). These patches are available at [OracleMetaLink](#).

## Task 4c: Verify that DB2 Client Is Installed for IBM FileNet P8

The procedures in this task describe how to install the DB2 Client software to prepare for the installation of Process Engine. Content Engine does not require installation of DB2 client software in the 4.0.x version of IBM FileNet P8 software.

In this task you will:

- Install and configure the DB2 client software
- Create several operating system users
- Create a client instance to connect to the remote DB2 server instance (Process Engine, UNIX only)

### NOTES

Before starting this task you will need the following information:

- The database name on the remote database server for the Process Engine database
- DB2 instance port numbers from the remote database server.
- Process Engine runtime database user (f\_sw or alias) password
- Process Engine maintenance database user (f\_maint or alias) password

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

### To install DB2 client software

---

To install the DB2 Client, log on as **root** user (UNIX) or as **Administrator** (Windows).

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

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

When the installation is finished, view the status report or go to **/tmp** (UNIX) or **\\My Documents\DB2log** (Windows) to view all DB2 install logs to ensure there are no errors.

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

---

Operating system users and groups must be created for the client instance on UNIX platforms only. Two users and two groups are required to create and use DB2 client instance for Process Engine on a UNIX platform. You may specify your own user and group names as long as they adhere to system naming rules and DB2 naming rules.

- Instance Owner - one is required for Process Engine.

- Instance Owner primary group - one is required for Process Engine.  
 The **instance owner** home directory is where the DB2 instance will be created.
- Fenced User - Process Engine only
- Fenced User primary group - Process Engine only  
 The **fenced user** runs user-defined functions (UDFs) and stored procedures outside the address space used by the DB2 database.

Each instance should have its own home file system.

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

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

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

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

When the database authentication type is set to SERVER or SERVER\_ENCRYPT, the following two additional operating system users must be created on the database server where the DB2 database resides.

**To create client instances for Process Engine( UNIX only)**

---

Depending upon your configuration, create at least one 32-bit instance for Process Engine.

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

---

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

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

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

After creating the new users and setting their group memberships, log off as **root** user, log on as each of the new users, and change the password to avoid connection problems the first time



they're used. For Client authentication these are the users and passwords that will be entered into Process Engine Setup to be used by the Process Engine software. In a Process Engine farm environment, when using DB2 Client authentication, it is required (at this time) to make sure all Process Engine runtime user names are the same on each Process Engine server in the farm.

### **To catalog the DB2 server node**

---

Reboot the server and log on as the instance owner on the Process Engine.

Use the db2ca tool, or catalog the DB2 server node as follows:

```
db2 catalog tcpip node <server alias> remote <server name> server <server side  
instance tcpip port #>
```

For example:

```
db2 catalog tcpip node aix20nod remote hqvwais20 server 60004
```

### **To catalog the Process Engine databases**

---

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

For example:

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

### **To verify the connection to the DB2 database**

---

Since the DB2 database is located on a remote server, verify that TCP/IP communications have been configured successfully on both server and client computers.

You can use the DB2 Configuration Assistant on a Windows client (enter **db2ca** at a command prompt) to connect the DB Client to the DB2 database on the remote database server. You can also use db2 at a command prompt.

### **To verify the ability to log on to the DB2 database**

---

After successfully connecting to the remote DB2 database, you can check the connection using either the Command Line Processor (CLP) or db2. Launch the Command Line Processor and enter for example:

```
db2 connect to <database_name> user <PE runtime user> using <password>
```

where:

<database\_name> is your Process Engine DB2 database name

<PE runtime user> is the Process Engine runtime user (f\_sw or alias)

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

## Task 5a: Configure an Application Server for Content Engine (WebSphere)

Content Engine requires a profile if it is to be deployed on WebSphere Application Server. A default profile (called *default*) is part of the initial WebSphere Application Server installation. You can use this profile, or create another one.

### NOTES

- If you intend to install AddOns (extensions to IBM FileNet P8 core components), and your Content Engine database will be Oracle, your Oracle JDBC Driver file requirements may be more restrictive. For the required version and patch number, see the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

Do the following procedure to specify the WebSphere environment variable for the JDBC driver corresponding to the type of database (DB2, Oracle, or SQL Server) where the Global Configuration Data (GCD) database will reside.

### To specify the WebSphere environment variables

---

1. Refer to the *IBM FileNet P8 Hardware and Software Requirements* for information on how to obtain the JDBC driver file for the database type that you need for the GCD or for an object store you will be creating later. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
2. Depending on your database types, copy the appropriate JDBC driver file to a directory on the WebSphere machine.

For example, you can copy the JDBC driver file to one of the following locations:

- (UNIX) **/opt/jars**
- (Windows) **C:\jars**

**CAUTION** Do not copy the JDBC driver file to either of the following directories.

- (UNIX) **...WebSphere/AppServer/lib/ext**
  - (Windows) **...WebSphere\AppServer\lib\ext**
3. If your WebSphere environment includes managed nodes, do the following:
    - a. Access the latest Content Engine installation software and extract the file **FNDShelper.jar**.
    - b. Copy **FNDTShelper.jar** to the **<WebSphere\_Install\_Path>/lib/ext** directory on each managed node.
  4. If you already have a profile for Content Engine to use, continue at [Step 5](#); or run the command script at one of the following (default) locations to create a new profile:
    - (AIX) **/usr/IBM/WebSphere/AppServer/wasprofile.sh**
    - (Other UNIX) **/opt/IBM/WebSphere/AppServer/wasprofile.sh**

- (Windows) **C:\Program Files\IBM\WebSphere\AppServer\wasprofile.bat**

**NOTE** Remember the profile name, as you will need to specify it when installing Content Engine.

5. Start the WebSphere administrative console and log on to your profile.
6. Navigate to one of the following:
  - (WebSphere 5.1.x) Environment > Manage WebSphere Variables
  - (WebSphere 6.0.x) Environment > WebSphere Variables
  - (WebSphere 6.1.x) Environment > WebSphere Variables
7. (WebSphere 5.1.x or 6.0.x) Click **Cell**.
8. (WebSphere 6.1.x) Click **Node**. From the drop-down list, choose your profile for Content Engine.
9. Depending on your database and the operating system of the WebSphere machine, create a JDBC environment variable and set its path to the location you specified in [Step 2](#):

Database	JDBC Environment Variable
SQL Server	MSSQLSERVER_JDBC_DRIVER_PATH
Oracle	ORACLE_JDBC_DRIVER_PATH
DB2	DB2UNIVERSAL_JDBC_DRIVER_PATH

10. Navigate to one of the following:
  - (WebSphere 5.1.x) Environment > Manage WebSphere Variables
  - (WebSphere 6.0.x) Environment > WebSphere Variables
  - (WebSphere 6.1.x) Environment > WebSphere Variables
 and then click **Node** (WebSphere 5.1.x or 6.0.x) or **Cell** (WebSphere 6.1.x).
11. Repeat [Step 9](#).
12. Navigate to Application server > server1 > Java and Process Management > Process Definition > Java Virtual Machine. In the Generic JVM arguments box, set the value to the following:
 

```
-Xms512m -Xmx1024m
```
13. Save your changes to the master configuration
14. If any of your object stores will be of a database type that differs from those whose JDBC environment variables you have already specified in this procedure, return to [Step 1](#); otherwise, continue at [Step 15](#).
15. Stop and start WebSphere Application Server.

### To adjust the WebSphere transaction timeout value (optional)

---

Content Engine relies on WebSphere's transaction-timeout value, whose default may be too short for some standard or administrative processes (such as adding an expansion product or upgrading to the latest version of Content Engine). You can increase the value via the WebSphere administrative console, as follows:

1. Log on to the WebSphere administrative console.
2. Navigate to Servers > Application Servers > *server\_instance* > Container Services > Transaction Service, where *server\_instance* is the name of the WebSphere server instance you created for Content Engine.
3. Change the value of the Total transaction lifetime timeout parameter to 600 (seconds) or more.

**CAUTION** A large enough timeout value is critical if you are upgrading Content Engine from version 3.5.x; otherwise, the upgrade may fail.

4. Save your change.

### To use WebSphere 6.1 federated user repository

---

A multi-realm installation of Content Engine with WebSphere is possible only with WebSphere 6.1. Earlier releases of WebSphere cannot support multiple user repositories. WebSphere 6.1's new federated user repository feature provides the ability to map entries from multiple individual user repositories into a single virtual repository.

**NOTE** In a change from earlier versions, Content Engine Setup no longer enables or checks the status of WebSphere's own internal security. Enabling global security (WebSphere 5.1 and 6.0) or administrative security (WebSphere 6.1) is now the responsibility of the WebSphere administrator and can be done anytime before or after running Content Engine Setup. But you must enable global security before putting an IBM FileNet P8 system into production.

1. Start the WebSphere 6.1 Administrative Console and log on to your profile.
2. Configure federated user repositories. Refer to your WebSphere documentation for detailed instructions.

**NOTE** You must enter information about all these user repositories while running Content Engine Enterprise Manager's Directory Configuration Wizard, one time for each user repository. See ["To configure directory service authentication" on page 187](#).

3. Select both **Enable administrative security** and **Enable application security**.

**CAUTION** Do **not** select **Enable Java 2 security** as this will add a security layer that Content Engine cannot address.

4. Confirm that the WebSphere administrative account that you plan to enter into the Content Engine Setup program is unique and not duplicated in any user repository that will be included in the federated user repository.

**CAUTION** Entering a non-unique WebSphere administrative account into Content Engine Setup could potentially lock you out of your WebSphere 6.1 console as soon as you enable Administrative Security.

5. Stop and start WebSphere Application Server. This will turn on administrative security.

Important points to remember if you have configured federated user repository:

- When you run Content Engine Setup and you are asked which components to install, do **not** select the Application Server Authentication Provider component, as this will replace and overwrite your existing federated user repository configuration with whatever WebSphere authentication settings you subsequently enter into the Content Engine Setup program.
- As explained in [“To configure directory service authentication” on page 187](#), you will run Enterprise Manager’s Directory Configuration Wizard once for each of the federated user repositories.

## Task 5b: Configure an Application Server for Content Engine (WebLogic)

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

Before installing and deploying Content Engine on a WebLogic machine, you need to create a WebLogic domain and install JDBC drivers. (The drivers must be installed on the WebLogic machine whether your database is collocated or not.) If you are running WebLogic 8.1.5 on Windows, you must install a BEA WebLogic patch. You need to install two Microsoft components if WebLogic is on a Windows machine. The steps are as follows:

1. Depending on your version of WebLogic, do one of the following procedures to create a WebLogic domain for Content Engine:
  - [“To configure WebLogic 8.1.x” on page 110](#)
  - [“To configure WebLogic 9.2.x” on page 112](#)
2. Depending on your database, do one of the following procedures to install the JDBC drivers:
  - [“To install WebLogic JDBC drivers \(DB2\)” on page 115](#)
  - [“To install WebLogic JDBC drivers \(MS SQL Server\)” on page 115](#)
  - [“To install WebLogic JDBC drivers \(Oracle\)” on page 116](#)
3. (Windows only) Install the following Microsoft components on the WebLogic machine:
  - Microsoft .NET 2.0 Framework
  - Microsoft Web Services Enhancements 3.0
4. If your authentication provider will utilize single sign-on (SSO) create the LDAP provider now, using the WebLogic Server Administration Console.

### **To configure WebLogic 8.1.x**

---

1. Start WebLogic Configuration Wizard.
2. In the Create or Extend a Configuration screen, click **Create a new WebLogic configuration**, and then click **Next**.
3. In the Select a Configuration Template screen, select Basic WebLogic Server Domain in the Templates pane, and then click **Next**.
4. In the Choose Express or Custom Configuration screen, click **Express** (to accept all default settings in the template) or **Custom** (to enable you to modify default settings), and then click **Next**.
5. In the Configure Administrative User and Password screen, type an (internal WebLogic) administrative username (at least eight characters) and password, and then click **Next**.
6. In the Configure Server Start Mode and Java SDK screen do the following and then click **Next**:
  - a. Click **Production Mode**.

- b. Click **BEA Supplied SDKs** and select the version of the JRockit SDK specified in the section Supported platforms for server components > Application Engine/Content Engine > Application/Web server layer of the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

**NOTE** When running in Production mode, you must supply a username and a password to stop and start WebLogic.

7. In the Create WebLogic Configuration screen, type a configuration (domain) name (*FNCEDomain*, for example, in this procedure) and location, and then click **Create**.
8. In the Creating Domain screen, after the message “Configuration Created Successfully!” appears, click **Done**.
9. If you haven’t already done so, go to [http://java.com/en/download/windows\\_ie.jsp](http://java.com/en/download/windows_ie.jsp) to download and install the latest Java Runtime Environment for the operating system where WebLogic runs.
10. Start WebLogic Server Administration Console.
11. (Optional) When running Content Engine Setup (see [“Install and Deploy Content Engine” on page 160](#)) you can have it create a WebLogic authentication provider (by installing the Application Server Authentication Provider component).

If you want Content Engine Setup to create the provider, skip to [Step 12](#); otherwise, you must create the provider *before* running Content Engine Setup. To create the provider:

- a. Navigate within the tree view of the WebLogic Server Administration Console to Security > Realms > myrealm > Providers > Authentication.
  - b. In the list view, click one of the supported authentication providers and specify the required configuration information.
12. Do the following to allow or prohibit logons to *FNCEDomain* by LDAP-authenticated users who are not in *FNCEDomain*’s active security realm:
    - a. Navigate in the tree view to *FNCEDomain* > Security > Realms > myrealm > Providers > Authentication > DefaultAuthenticator and click the General tab.
    - b. In the Control Flag drop-down list, choose SUFFICIENT to allow logons by users not in the active security realm; choose REQUIRED to prohibit such logons.
    - c. Click **Apply**.
  13. If WebLogic is using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers, as follows, so that the most-frequently-used provider is first in the list, and the least-frequently-used is last.

**NOTE** Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.

- a. In WebLogic Server Administration Console, select Security > Realms in the tree view and then click the realm whose providers you are going to reorder.
- b. Click the Providers:Authentication tab and then click **Reorder**.
- c. Use the arrow buttons to reorder the providers as needed.

d. Save your changes.

14. In the Details tab of the authentication provider to be used by Content Engine, set the values (specified for performance reasons) that control searches within the authentication provider, as shown in the following table:

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

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

15. Specify the following heap sizes for the JVM:

- Initial Java heap size (-Xms): 512 MB
- Maximum Java heap size (-Xmx): 1024 MB

Do [Step 16](#) and [Step 17](#) only for WebLogic 8.1.5 on Windows; otherwise, skip to [Step 18](#).

16. Go to <http://support.bea.com> and download the patch file **CR247206\_810sp5.jar** to a directory on the Windows WebLogic machine, for example **C:\WebLogic815Patch**.

17. In the file **startWebLogic.cmd** for the domain you created, insert the following line immediately after the first line that starts with `set CLASSPATH=...`

```
set LDAP_JAR=C:\WebLogic815Patch\CR247206_810sp5.jar;%CLASSPATH%
```

18. Continue at one of the following procedures, depending on your database:

- (DB2) ["To install WebLogic JDBC drivers \(DB2\)" on page 115](#)
- (MS SQL Server) ["To install WebLogic JDBC drivers \(MS SQL Server\)" on page 115](#)
- (Oracle) ["To install WebLogic JDBC drivers \(Oracle\)" on page 116](#)

### To configure WebLogic 9.2.x

1. Start WebLogic Configuration Wizard.
2. In the Welcome screen, click **Create a new WebLogic domain**, and then click **Next**.
3. In the Select Domain Source screen,
  - a. Click **Generate a domain configured automatically to support the following BEA products**.
  - b. Select the WebLogic Server (Required) check box, if it isn't already selected.
  - c. Click **Next**.
4. In the Configure Administrator User and Password screen, specify an (internal WebLogic) administrative username (at least eight characters), a password, a description, and then click **Next**.
5. In the Configure Server Start Mode and JDK screen do the following, and then click **Next**:



- a. Click **Production Mode** and **BEA Supplied SDKs**.
- b. Refer to the section Supported platforms for server components > Application Engine/Content Engine > Application/Web server layer of the *IBM FileNet P8 Hardware and Software Requirements* for the supported version of JRockit SDK or the JDK, depending on the operating system where WebLogic is hosted. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- c. Select one of the following, depending on the operating system where WebLogic is hosted:
  - (Windows) The version of the JRockit SDK specified in the Third Party Support Information section of the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
  - (AIX) IBM SDK 1.5.0 at /usr/java5
6. In the Customize Environment and Services Settings screen, click **Yes**, and then click **Next**.
7. In the Configure the Administration Server screen, accept the default values for the each field, including SSL not enabled, and then click **Next**.
8. In the Configure Managed Servers screen, click **Next**.
9. In the Configure Machines screen, click **Next**.
10. In the Review WebLogic Domain screen, review the information and click **Next**.
11. In the Create WebLogic Domain screen, type a domain name (*FNCEDomain*, for example, in this procedure) and location, and then click **Create**.
12. In the Creating Domain screen, after the message “Domain Created Successfully!” appears, click **Done**.
13. Start WebLogic Server Administration Console.
14. (Optional) When running Content Engine Setup (see [“Install and Deploy Content Engine” on page 160](#)) you can have it create a WebLogic authentication provider (by installing the Application Server Authentication Provider component).

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

If you want Content Engine Setup to create the provider, skip to [Step 15](#); otherwise, you must create the provider *before* running Content Engine Setup. To create the provider:

- a. Navigate within the tree view of the WebLogic Server Administration Console to Services > Security Realms.
- b. In the list view, click myrealm, click the Providers tab, and then click **New**.
- c. Specify the required configuration information for the authentication provider.

- d. In the Provider Specific tab, set the values (specified for performance reasons) that control searches within the authentication provider, as shown in the following table:

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

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

15. Do the following to allow or prohibit logons to *FNCEDomain* by LDAP-authenticated users in the DefaultAuthenticator who are not in *FNCEDomain*'s active security realm:

- Start WebLogic Server Administration Console and navigate in the tree view to *FNCEDomain* > Security > Realms > myrealm > Providers > DefaultAuthenticator.
- In the Control Flag drop-down list, choose SUFFICIENT to allow logons by users not in the active security realm (such as users in your authentication domain); choose REQUIRED to prohibit such logons.

If you choose REQUIRED, all the users you wish to authenticate for Content Engine must exist not only in the LDAP directory, but also exist as WebLogic users who are in the Default Authenticator provider.

- Click **Apply**.

16. If WebLogic is using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers so that the most-frequently-used provider is first in the list, and the least-frequently-used is last:

- In WebLogic Server Administration Console, navigate in the tree view to *FNCEDomain* > Security > Realms > myrealm > Providers > Authentication Providers.
- Click **Re-order the Configured Authentication Providers**.
- Use the arrow buttons to reorder the providers as needed.
- Click **OK**.

**NOTE** Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.

17. Specify the following heap sizes for the JVM:

- Initial Java heap size (-Xms): 512 MB
- Maximum Java heap size (-Xmx): 1024 MB

18. Continue at one of the following procedures, depending on your database:

- (DB2) ["To install WebLogic JDBC drivers \(DB2\)" on page 115](#)
- (MS SQL Server) ["To install WebLogic JDBC drivers \(MS SQL Server\)" on page 115](#)
- (Oracle) ["To install WebLogic JDBC drivers \(Oracle\)" on page 116](#)

### To install WebLogic JDBC drivers (DB2)

---

1. Obtain the latest version of the Redistributable DB2 JDBC Driver Type 4 driver v8 from the [IBM web site](http://www.ibm.com) (<http://www.ibm.com>) by searching for “JDBC Type 4”.
2. Add the following JAR files to the WebLogic classpath:
  - **db2jcc.jar**
  - **db2jcc\_license\_cu.jar**

For example,

```
set CLASSPATH=%CLASSPATH%;c:\db2\jdbc\db2jcc.jar;c:\db2\jdbc\db2jcc_license_cu.jar
```

3. Stop and then start WebLogic Server.

### To install WebLogic JDBC drivers (MS SQL Server)

---

1. Download and unzip Microsoft SQL Server Driver 2005 JDBC Driver from Microsoft Support to a directory `<jdbc_path>` on your application server machine, such as:
  - (UNIX) **/opt/jars**
  - (Windows) **C:\jars**
2. (WebLogic on Windows) Edit the file **startWebLogic.cmd** (by default, in the directory **C:\bea\user\_projects\domains\FNCEDomain**) for the WebLogic domain you created, as follows:
  - (WebLogic 8.1.5) Insert the following lines immediately after the line you inserted in [Step 17](#) of the procedure “[To configure WebLogic 8.1.x](#)” on page 110:

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

- (WebLogic 8.1.6) Insert the following lines immediately after the line you inserted in [Step 17](#) of the procedure “[To configure WebLogic 8.1.x](#)” on page 110:

```
set JDBC_PATH=<jdbc_path>\sqljdbc_1.1\enu\sqljdbc.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%
```

- (WebLogic 9.2.x) Insert the following lines immediately after the first occurrence of the line `set CLASSPATH=%CLASSPATH%;...`

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

3. (WebLogic on AIX) Add the following line to the file **setDomainEnv.sh** file:

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

4. (WebLogic on UNIX) Edit the file **startWebLogic.sh** by inserting the following two lines immediately after the first occurrence of the line `CLASSPATH=...`

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

5. Stop and then start WebLogic Server.

### To install WebLogic JDBC drivers (Oracle)

---

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

**NOTE** If you intend to install AddOns (extensions to IBM FileNet P8 core components), and your Content Engine database will be Oracle, your Oracle JDBC Driver file requirements may be more restrictive. For the required version and patch number, see the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on page 21.

3. From the Oracle web site, apply the patches Oracle Patch Ojdbc14.
4. Edit the file `startWebLogic.cmd` in your WebLogic domain `<WL_DomainName>` by adding the following line immediately after the first line that starts with `set CLASSPATH`.
  - (Windows)

```
set JDBC_PATH=<jdbc_path>\ojdbc##.jar
set CLASSPATH=%JDBC_PATH%;%CLASSPATH%;
```
  - (UNIX)

```
JDBC_PATH=<jdbc_path>/ojdbc##.jar
CLASSPATH=$JDBC_PATH:$CLASSPATH;
```
5. Stop and then start WebLogic Server.

### To adjust the WebLogic 8.1.x transaction timeout value (optional)

---

On WebLogic 8.1.x, Content Engine relies on the transaction-timeout value, whose default may be too short for some standard or administrative processes (such as adding an expansion product or upgrading to the latest version of Content Engine). You can increase the value via WebLogic Server Administration Console, as follows:

1. Log on to WebLogic Server Administration Console.
2. On the left-side pane, open the Services > JTA node.
3. Change the value of Timeout Seconds to at least 600 seconds.
4. Save your change.

### To adjust the WebLogic 9.2 transaction timeout value (optional)

---

On WebLogic 9.2, Content Engine relies on the transaction-timeout value, whose default may be too short for some standard or administrative processes (such as adding an expansion product or upgrading to the latest version of Content Engine). You can increase the value via WebLogic Server Administration Console, as follows:

1. Log on to WebLogic Server Administration Console.
2. Click **Lock and Edit**.
3. On the left-side pane, open the **Services > JTA** node.
4. Change the value of **Timeout Seconds** to at least 600 seconds.
5. Save your change.

## Task 5c: Configure an Application Server for Content Engine (JBoss)

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

### To configure JBoss for Content Engine

---

1. Navigate to the JBoss directory **JBOSS\_DIST/server**, which contains configuration file sets.
2. Create a new configuration file set by copying the **default** configuration file set to a new directory (called **server1** in this procedure) within the **server** directory.
3. Edit the **run.conf** configuration file, located at **JBOSS\_DIST/bin**, as follows:
  - a. Add a line to specify the path to the JDK on the machine where JBoss is installed, as shown in the following example:

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

```
-Xms128m -Xmx512m
```

to

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

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

to

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

where **<jboss\_install\_dir** is the directory where JBoss is installed.
  - b. Save your edit.
5. (DB2 only) Download the DB2 JDBC Driver referenced in the *IBM FileNet P8 Hardware and Software Requirements* and place it in CLASSPATH by copying the driver to directory **JBOSS\_DIST/server/server1/lib**. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).
6. (MS SQL Server only) Download the Microsoft SQL Server 2005 JDBC Driver, **sqljdbc.jar**, referenced in the *IBM FileNet P8 Hardware and Software Requirements* and place it in CLASSPATH by copying it to directory **JBOSS\_DIST/server/server1/lib**. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).

7. (Oracle only) Download the Oracle JDBC Driver, **ojdbc14.jar**, that is referenced in the *IBM FileNet P8 Hardware and Software Requirements* and place it in CLASSPATH by copying it to directory **JBOSS\_DIST/server/server1/lib**. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
8. If it isn't already running, start JBoss as follows, and leave the command window open:
  - (UNIX) `./run.sh -c server1`
  - (Windows) `run.bat -c server1`

## Manually Configure JBoss Data Sources for the GCD (Optional)

The Global Configuration Data (GCD) requires data sources for database connectivity. Create the data sources manually, as shown in the following procedures; or have Content Engine Setup create them, in which case skip these procedures and go to [“Installation Tasks” on page 128](#).

The procedures for creating data sources assume the following environment for your JBoss installation (substitute your own environment in its place where applicable):

- **JBOSS\_DIST** is the home directory of your JBoss installation
- **server1** is the JBoss server instance on which you will install and deploy Content Engine.

### To configure JBoss 4.0.x database connectivity (DB2)

---

1. Copy the files **FNGCD-Unix-ds.xml** and **FNGCD-Unix-xa-ds.xml** in directory **JBOSS\_DIST/samples** to **JBOSS\_server/server1/deploy** and rename the file **FNGCD-ds.xml** and **FNGCD-xa-ds.xml**, respectively.
2. In the non-XA data source file, edit the section titled `<sample datasource for DB2>` in the file **FNGCD-ds.xml**:
  - a. Set the content of the `<jndi-name>` element to `FNGCDDS`.
  - b. Set the content of the `<connection-url>` element to

```
jdbc:db2://<JBoss_host_name>:<DB_port>/<dbname>
```

where `<dbname>` is the name of the database that will contain the GCD.
  - c. Change the content of the `<user-name>` and `<password>` elements to that of the user and password of the tablespace to be used for the GCD.
  - d. Save your edits.
3. In the XA data source file, edit the section titled `<sample datasource for DB2>` in the file **FNGCD-Unix-xa-ds.xml**:
  - a. Set the content of the `<jndi-name>` element to `FNGCDDSXA`.
  - b. Set the content of the `<xa-datasource-property name="ServerName">` element to the IP address of the machine where DB2 is installed.
  - c. Set the content of the `<xa-datasource-property name="DatabaseName">` element to the name of the database that will contain the GCD.

- d. Set the content of the `<xa-datasource-property name="PortNumber">` element to the port used by DB2.
- e. Change the content of the `<user-name>` and `<password>` elements to that of the user and password of the tablespace to be used for the GCD.
- f. Save your edits.

### To configure JBoss 4.0.x database connectivity (MS SQL Server)

---

1. Copy the files **FNGCD-Unix-ds.xml** and **FNGCD-Unix-xa-ds.xml** in directory **JBOSS\_DIST/samples** to **JBOSS\_server/server1/deploy** and rename the file **FNGCD-ds.xml** and **FNGCD-xa-ds.xml**, respectively.
2. In the non-XA data source file, edit the section titled `<sample datasource for MS SQLServer>` in the file **FNGCD-ds.xml**:
  - a. Set the content of the `<jndi-name>` element to `FNGCDDS`.
  - b. Set the content of `<connection-url>` element to
 

```
jdbc:sqlserver://<dbserver>:<db_port>;DatabaseName=<dbname>
```

 where `<dbname>` is the name of the database that will contain the GCD.
  - c. Change the content of the `<user-name>` and `<password>` elements to that of the user and password of the database to be used for the GCD.
  - d. Save your edits.
3. In the XA data source file, edit the section titled `<sample datasource for MS SQLServer>` in the file **FNGCD-xa-ds.xml**:
  - a. Set the content of the `<jndi-name>` element to `FNGCDDSXA`.
  - b. Set the content of the `<xa-datasource-property name="ServerName">` element to the IP address of the machine where MS SQL Server is installed.
  - c. Set the content of the `<xa-datasource-property name="DatabaseName">` element to the name of the database that will contain the GCD.
  - d. If MS SQL Server is not using the default port number (1433), add this XML element:
 

```
<xa-datasource-property name="PortNumber">dbport</xa-datasource-property>
```

 where `dbport` is the port number used by the MS SQL Server database instance.
  - e. Change the content of the `<user-name>` and `<password>` elements to that of the user and password of the database to be used for the GCD.
  - f. Save your edits.

### To configure JBoss 4.0.x database connectivity (Oracle)

---

1. Copy the files **FNGCD-Unix-ds.xml** and **FNGCD-Unix-xa-ds.xml** in directory **JBOSS\_DIST/samples** to **JBOSS\_server/server1/deploy** and rename the file **FNGCD-ds.xml** and **FNGCD-xa-ds.xml**, respectively.



2. In the non-XA data source file, edit the section titled <sample datasource for Oracle> in the file **FNGCD-Unix-ds.xml**:
  - a. Set the content of the <jndi-name> element to FNGCDDS.
  - b. Set the content of <connection-url> element to

```
jdbc:oracle:thin:@<dbserver>:<db_port>:<Oracle_service_name>
```

where *dbserver* is the name of the database machine and *port* is the port used by the database that will contain the GCD.
  - c. Change the content of the <user-name> and <password> elements to that of the user and password of the tablespace to be used for the GCD.
  - d. Save your edits.
3. In the XA data source file, edit the section titled <sample datasource for Oracle> in the file **FNGCD-xa-ds.xml**:
  - a. Set the content of the <jndi-name> element to FNGCDDSXA.
  - b. Set the content of the <xa-datasource-property name="URL"> element to

```
jdbc:oracle:thin:@<dbserver>:<db_port>:<Oracle_service_name>
```

where *dbserver* is the name of the database machine and *port* is the port used by the database that will contain the GCD.
  - c. Change the content of the <xa-datasource-property name="User"> and <xa-datasource-property name="Password"> elements to that of the user and password of the tablespace to be used for the GCD.
  - d. Save your edits.

## Encrypt LDAP Password (Optional)

Do the following procedure to encrypt the password used by JBoss to query an LDAP directory server.

### NOTES

- Substitute your own value (at least eight characters) for the Salt attribute in place of the value `twasalt12` used in this procedure.
- Substitute your own obscuring and encrypted LDAP passwords in place of the passwords `Magic` and `Secret` used in this procedure.
- Assume the path to the JBoss **conf** directory is `.../server/myserver/conf`, where *myserver* is the name of the JBoss server. Use your own value in place of *myserver*.

### To encrypt the LDAP password

---

1. Edit the file **jboss-service.xml** in the **conf** directory by adding the following text at the end of the Security section of the file, just before the Transactions section:

```
<mbean code="org.jboss.security.plugins.JaasSecurityDomain"
  name="jboss.security:service=JaasSecurityDomain,domain=ServerMasterPassword">
```

```
<constructor>
  <arg type="java.lang.String" value="ServerMasterPassword"/>
</constructor>
<!-- The opaque master password file used to decrypt the encrypted database
password key -->
<attribute name="KeyStorePass">
  {CLASS}org.jboss.security.plugins.FilePassword:${jboss.server.home.dir}/
  conf/server.password
</attribute>
<attribute name="Salt">twsalt12</attribute>
<attribute name="IterationCount">13</attribute>
</mbean>
```

2. Choose an obscuring password in place of `Magic`, the password used in this procedure.
3. Run the following commands from the command prompt to create the **server.password** file:

```
cd %JBOSS_HOME%\server\myserver\conf
```

If an old **server.password** file exists, delete it:

```
del server.password
```

Create the **server.password** file by running the following command (without carriage returns):

```
java -cp ..\lib\jbossx.jar org.jboss.security.plugins.FilePassword twsalt12 13
"Magic" server.password
```

4. Run the following two commands, substituting your obscuring password and real LDAP password in place of `Magic` and `Secret`, respectively:

```
cd %JBOSS_HOME%\server\myserver\conf
java -cp ..\lib\jbossx.jar org.jboss.security.plugins.PBEUtils twsalt12 13 "Magic"
"Secret"
```

Copy to the clipboard the encrypted password (`Fm/hKKlyXZj`, for example) displayed on the command console.

5. Edit **login-config.xml** in the **conf** directory, as follows:
  - a. Paste the encrypted password (`Fm/hKKlyXZj`, for example) for `LdapExtLoginModule`, as shown below:

```
<module-option name="bindCredential">Fm/hKKlyXZj</module-option>
```

- b. Add the following XML element below the line you added in step 4a:

```
<module-option name="jaasSecurityDomain">
  jboss.security:service=JaasSecurityDomain,domain=ServerMasterPassword
</module-option>
```

## Encrypt GCD Data Source Passwords (Optional)

To encrypt the GCD passwords (XA and non-XA), do the procedure [“To encrypt data source passwords”](#) on page 250 in [“Configure Content Engine Application Server Database Connectivity \(JBoss 4.0.x\)”](#) on page 247.

## Task 6a: Configure an Application Server for Application Engine (WebSphere)

You must install WebSphere Application Server on the machine where you are going to install and deploy Application Engine.

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

### **To Configure WebSphere for Application Engine**

---

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

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

## Task 6b: Configure an Application Server for Application Engine (WebLogic)

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

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

### To Configure WebLogic for Application Engine

---

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

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

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

3. Make a backup copy of the application server startup script.

Backup **startWebLogic.cmd** for Windows or **startWebLogic.sh** for UNIX.

**NOTE** If you are not using a WebLogic domain, backup **startWLS.cmd** for Windows or **startWLS.sh** for UNIX.

4. Edit the application server startup script MEM\_ARGS settings.

Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

**NOTE** If the MEM\_ARGS variable doesn't exist, add it to the startup script.

- For all systems except those using JRockit JAVA.

Append the following to the MEM\_ARGS variable:

```
-XX:MaxPermSize=<size>m
```

where <size> is the value, in MB, of the MaxPermSize.

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).

- For systems using JRockit JAVA.

Append the following to the MEM\_ARGS variable:

```
-Xgc:gencon
```

5. (WebLogic 8.1.5 on Windows with container-managed authentication) Install the patch file **CR247206\_810sp5.jar** from <http://support.bea.com>.
  - a. Go to <http://support.bea.com> and download the patch file **CR247206\_810sp5.jar** to a directory on the Windows WebLogic machine, for example **C:\WebLogic815Patch**.
  - b. In the file **startWebLogic.cmd** for the domain you created, insert the following line immediately after the first line that starts with `set CLASSPATH=...`

```
set LDAP_JAR=C:\WebLogic815Patch\CR247206_810sp5.jar;%CLASSPATH
```

## Task 6c: Configure an Application Server for Application Engine (JBoss)

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

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

### To Configure JBoss for Application Engine

---

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

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

2. Make a backup copy of the application server startup script.

Backup **run.bat** (Windows) or **run.sh** (UNIX).

3. Edit the application server startup script Java settings.

- a. Add a line to specify the path to the JDK on the machine where JBoss is installed, as shown in the following example (Windows):

```
set JAVA_HOME=C:\Program Files\Java\jdk1.5.0_06
```

**NOTE** If your JDK is different from version 1.5.0, substitute your version for the one listed above.

- b. Update the JAVA\_OPTS memory settings.

Adjusting this setting prevents the application server from running out of memory, a condition in which users would not be able to log in to Workplace.

In the `JAVA_OPTS` line, change the `-Xms` and `-Xmx` values (**bold**) for your configuration.

Example (Windows):

```
set JAVA_OPTS=%JAVA_OPTS% -Xms128m -Xmx512m
```

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).

- c. Save your edits.

4. Start JBoss as follows, and leave the command window open:

- (UNIX) `./run.sh`
- (Windows) `run.bat`

5. (Optional) Disable JBoss logging.

In development mode, JBoss creates a large number of HTTP Access, "INFO", "DEBUG" and "TRACE" log messages. This can cause unexpected behavior in the deployed IBM FileNet software. Using the following procedure, you can limit this type of excessive JBoss logging.

**NOTE** When logging is disabled, error messages will still be displayed in the JBoss console.

- a. Edit the **log4j.xml** file (`<JBoss_home>/server/default/conf/log4j.xml`).
  - i. Change all threshold values and priority values from "INFO", "DEBUG", or "TRACE" to "ERROR".
  - ii. Delete or comment out the "Preserve messages in a local file" to turn off the server log.
- b. To turn off HTTP access logging, open **jboss-service.xml** with a text editor and delete or comment out the "Access logger" section.

Location of jboss-service.xml:

- JBoss 4.05 - `<JBoss_Home>/server/default/deploy/jbossweb-tomcat55.sar/META-INF`

c. Open **web.xml** and change the logVerbosityLevel to "FATAL".

Location of web.xml:

- JBoss 4.05 - `<JBoss_Home>/server/default/deploy/jbossweb-tomcat55.sar/conf`

d. Restart the JBoss server.

# Installation Tasks

## To install the core IBM FileNet P8 Platform components

---

1. Install IBM FileNet P8 Platform documentation on your application server. Do one of the following:
  - [Task 7a on page 130](#) (WebSphere)
  - [Task 7b on page 135](#) (WebLogic)
  - [Task 7c on page 141](#) (JBoss)
2. Install Content Search Engine. Do [Task 8 on page 147](#).
3. Set up Content Engine. Do the following:
  - a. Install and deploy Content Engine. Do [Task 9 on page 160](#).
  - b. Install Content Engine software updates. Do [Task 10 on page 177](#).
  - c. Complete Post-Install Content Engine Configuration. Do [Task 11 on page 178](#).
  - d. Install Enterprise Manager. Do [Task 12 on page 183](#).
  - e. Establish the IBM FileNet P8 Domain and Global Configuration Data (GCD). Do [Task 13 on page 187](#).
  - f. Configure database connectivity on the Content Engine application server. Do one of the following:
    - [Task 14a on page 193](#) (WebSphere 5.1.x)
    - [Task 14b on page 206](#) (WebSphere 6.0.x)
    - [Task 14c on page 220](#) (WebSphere 6.1.x)
    - [Task 14d on page 236](#) (WebLogic 8.1.x)
    - [Task 14e on page 241](#) (WebLogic 9.2.x)
    - [Task 14f on page 247](#) (JBoss 4.0.x)
  - g. Prepare storage areas for object stores. Do [Task 15 on page 252](#).
  - h. Create object stores. Do [Task 16 on page 258](#).
  - i. Verify the Content Engine installation. Do [Task 17 on page 263](#).
4. Install Content Search Engine Software Updates. Do [Task 18 on page 265](#).
5. Set up Process Engine. Do the following:
  - a. Install Process Engine. Do one of the following:
    - [Task 19a on page 266](#) (Windows 2003)
    - [Task 19b on page 285](#) (Sun Solaris)



- [Task 19c on page 304](#) (AIX)
  - [Task 19d on page 320](#) (HP-UX)
- b. Install Process Engine Software Updates. Do [Task 20 on page 338](#).
  - c. Install the Latest Content Engine Client Files on Process Engine Servers. Do [Task 21 on page 339](#).
  - d. Configure Process Task Manager. Do [Task 22 on page 340](#).
  - e. Complete Process Engine Configuration. Do [Task 23 on page 342](#).
6. Set up Application Engine. Do the following:
- a. Install Application Engine. Do [Task 24 on page 344](#).
  - b. Configure Application Engine. Do one of the following:
    - [Task 25a on page 356](#) (WebSphere)
    - [Task 25b on page 373](#) (WebLogic)
    - [Task 25c on page 381](#) (JBoss)
  - c. Install Application Engine Software Updates. Do [Task 26 on page 383](#).
  - d. Install Content Engine and Process Engine client file updates. Do [Task 27 on page 384](#).
  - e. Deploy Application Engine. Do one of the following:
    - [Task 28a on page 385](#) (WebSphere)
    - [Task 28b on page 391](#) (WebLogic)
    - [Task 28c on page 394](#) (JBoss)

## Task 7a: Install IBM FileNet P8 Platform Documentation (WebSphere)

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

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

### NOTES

- You may collocate the IBM FileNet P8 documentation on an IBM FileNet P8 Application Engine or IBM FileNet P8 Content Engine machine.
- To ensure proper documentation search functionality, make sure that JavaScript™ support is enabled on each user's browser client.
- Because of the number of possible network and web configurations, contact your network or web administrator for specific system, software, and security requirements.
- The requirements described in this topic may also exist for associated custom application help, or for any customizations you may have applied to the IBM FileNet P8 documentation. For further details on customizing or localizing IBM FileNet P8 documentation, see the file [ecm\\_help/transkit/P8DocTransKit.htm](#) on the IBM FileNet P8 Documentation package or wherever you install the IBM FileNet P8 documentation.
- If you are upgrading your IBM FileNet P8 Platform installation and have downloaded a refresh of the documentation from the [IBM Information Management support page on www.ibm.com](#), see [“Upgrade IBM FileNet P8 Documentation” on page 478](#) for instructions on upgrading the documentation.
- Under WebSphere, IBM FileNet P8 documentation must be installed and deployed as a WAR file (ecm\_help.war). You cannot deploy as an EAR file because in that case the IBM FileNet P8 help remains packaged and the internal documentation search engine will not function as expected.

### Overview of Procedures

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

- [“Install IBM FileNet P8 Platform Documentation” on page 131](#) -- Begin here if you are installing the IBM FileNet P8 Platform documentation for the first time.
- [“Install Expansion Product Documentation” on page 131](#) -- Begin here if you have already installed the IBM FileNet P8 Platform documentation and need to add on the documentation for IBM FileNet expansion products (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets).

- [“Update Documentation Search Index” on page 133](#) -- Perform this procedure only after you have added on documentation for your expansion products. Otherwise, if you install only IBM FileNet P8 Platform documentation, which has a baseline search index, you may skip this procedure.
- [“Complete and Verify the Documentation Installation” on page 134](#) -- Perform this final procedure in all cases, but only after you have installed the IBM FileNet P8 Platform documentation, added on documentation for your expansion products, and (where required) regenerated the search index.

## ***Install IBM FileNet P8 Platform Documentation***

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

**NOTE** Depending on your operating system and WebSphere versions, your WebSphere screens may be slightly different than those documented in the examples below.

### **To install the IBM FileNet P8 Platform documentation on WebSphere**

---

1. Access the IBM FileNet P8 Platform Documentation package.
2. Copy the IBM FileNet P8 **ecm\_help.war** file from the IBM FileNet P8 documentation package to a location on the local hard drive (for example, **.../p8docs/ecm\_help**).
3. If it is not running, start the WebSphere server, and start the WebSphere administrative console.
4. From the WebSphere administrative console, select Applications > Install a New Application.
5. Enter the Local Path location or Browse to the **ecm\_help.war** file.
6. Enter a Context Root, then click **Next** and follow all WebSphere installation screens.

**NOTE** The **Context Root** (for example, **ecm\_help**) and installation directory can be user-defined, so substitute your actual values when prompted if you do not want to take the WebSphere default values.

7. Save your changes to the WebSphere Master Configuration.
8. Continue as follows:
  - If you need to add documentation for any IBM FileNet P8 expansion product, go on to the procedure in the following topic, [“Install Expansion Product Documentation” on page 131](#).
  - If you have no further documentation to install, go on to the procedure in the topic [“Complete and Verify the Documentation Installation” on page 134](#).

## ***Install Expansion Product Documentation***

Use the procedure in this topic to install documentation for expansion products onto an existing IBM FileNet P8 Platform documentation server. If you have no such documentation to add on, skip to the procedure [“Complete and Verify the Documentation Installation” on page 134](#).

## To install expansion product documentation

---

1. Determine the expansion product documentation media source or location:
  - For most expansion products, use the *Documentation* package included as part of the particular software.
  - For Content Federation Services for Image Services, use the **Documentation** directory in the software package.
2. If the IBM FileNet P8 Platform documentation application (ecm\_help) is running, stop it. Verify that no processes are accessing the documentation web application.

**NOTE** On UNIX or Windows, stop the ecm\_help application from the WebSphere administrative console.

3. Copy the expansion product documentation to IBM FileNet P8 Platform documentation server, as follows:
  - (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 Platform documentation server.  

```
cp -r <mount_location> <target_destination>
```

**WARNING** Care should be taken when copying folders in UNIX. Dropping-and-dragging of folders replaces any existing folder(s) of the same name. Also note that your switch (`-R`) requirement may be different from the example shown. Contact your system administrator if you have questions about proper syntax.
  - (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
  - (Content Federation Services for Image Services) Copy the `cfs_guide.pdf` file from the **Documentation** directory on the software package into the **ecm\_help/cfs\_help** directory deployed on the existing IBM FileNet P8 Platform documentation site.

**NOTE** Repeat this step for each of your expansion products or custom applications. You can copy more than one expansion product documentation set to the documentation application server before continuing (for example, IBM FileNet P8 eForms and IBM FileNet P8 Portlets) so you end up with one **ecm\_help** directory structure containing multiple sets of expansion product documentation files added to it.

4. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM P8 Platform and various functional expansion products. Check the documentation page on the [IBM Information Management support page on www.ibm.com](#) for the latest versions of these guides. See “Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21 for details.
5. Go on to the procedure in the following topic, “Update Documentation Search Index” on page 133.

## Update Documentation Search Index

Perform this procedure only if you have refreshed the core IBM FileNet P8 Platform documentation, or installed expansion product documentation onto your IBM FileNet P8 Platform documentation server. Otherwise, skip to the procedure in the following topic, [“Complete and Verify the Documentation Installation” on page 134](#).

**NOTE** Any time you update the documentation search index, a backup of the files in the existing **Index/core** directory will be automatically copied to the **Index/IndexOld** subdirectory. You can reapply these backed-up files to the **core** subdirectory (after first removing the new files created there) if you need to return to your previous indexed state.

### To update the documentation search index

---

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

**NOTE** Make sure you have copied the help for all your various expansion products to a designated application server location containing the IBM FileNet P8 Platform documentation. Otherwise, you will have to repeat this procedure if you add additional documentation later.

2. Open a command prompt or terminal on the application server.
3. From the command line, navigate to the **search** subdirectory under the application root directory, for example, **ecm\_help**.
4. Using a text editor, open the search-indexing script file that is appropriate for your operating system:

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

(UNIX) **indexFiles.sh**

(Windows) **indexFiles.bat**

5. If necessary, set the JAVA\_HOME variable in the script file with the path to your JRE installation. The default examples are:

(UNIX) `JAVA_HOME="/usr/java/j2sdk1.4.1_02"`

(Windows) `SET JAVA_HOME=c:\j2sdk1.4.2`

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

6. Save your changes and close the text editor.

**NOTE** If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (`chmod 755`) to the **indexFiles.sh** file.

7. From the ecm\_help/search folder, execute the updated search-indexing script file that is appropriate to your application server operating system.

(UNIX) **indexFiles.sh**

(Windows) **indexFiles.bat**

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

8. Go on to the next procedure.

## Complete and Verify the Documentation Installation

Perform this procedure after you have installed (and, if necessary, reindexed) the IBM FileNet P8 documentation on the application server.

### To complete and verify the documentation installation on WebSphere

---

1. Start the IBM FileNet P8 documentation application from the WebSphere administrative console, Applications > Enterprise Applications.
2. Verify that the WebSphere application server and the new IBM FileNet P8 documentation web site are running:
  - a. From your web browser, access the following URL. The documentation's Help Directory should open.

**`http://<docserver>:<port#>/<contextRoot>/`**

where:

*docserver* is the name of the Java web server.

*port#* is the port number (for example, 9080).

*contextRoot* is the value of the Map to URL field that you specified when you deployed the IBM FileNet P8 Platform documentation application. If you specified **`/ecm_help`**, then the *contextRoot* is **`ecm_help`**.

Example: **`http://yourdocserver:9080/ecm_help/`**

**NOTE** You can use multi-part root directories (for example, **`/docs/ecm_help`**) if your application server supports them.

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

**NOTE** When it is time to configure the online help location for the various IBM FileNet P8 components, either while running Setup programs or later via site preferences settings, use the URL in [Step a on page 134](#).

## Task 7b: Install IBM FileNet P8 Platform Documentation (WebLogic)

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

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

### NOTES

- You may collocate the IBM FileNet P8 Platform documentation on an IBM FileNet P8 Application Engine or IBM FileNet P8 Content Engine machine.
- To ensure proper documentation search functionality, make sure that JavaScript support is enabled on each user's browser client.
- Because of the number of possible network and web configurations, contact your network or web administrator for specific system, software, and security requirements.
- The requirements described in this topic may also exist for associated custom application help, or for any customizations you may have applied to the IBM FileNet P8 Platform documentation. For further details on customizing or localizing IBM FileNet P8 Platform documentation, see the file [ecm\\_help/transkit/P8DocTransKit.htm](#) in the IBM FileNet P8 Platform Documentation package or wherever you install the IBM FileNet P8 Platform documentation.
- If you are upgrading your IBM FileNet P8 Platform installation and have downloaded a refresh of the documentation from the [IBM Information Management support page on www.ibm.com](#), see ["Upgrade IBM FileNet P8 Documentation" on page 478](#) for instructions on upgrading the documentation.
- Under WebLogic, you must install and deploy IBM FileNet P8 documentation as a flat-file directory (example, /ecm\_help). You cannot deploy the documentation as a WAR file or EAR file because, in both cases, the IBM FileNet P8 help remains packaged and the internal documentation search engine will not function as expected.
- For WebLogic application server, you cannot directly deploy from the package. Therefore, you must first copy the IBM FileNet P8 Platform documentation files from the package to the PC where WebLogic is running, as described in this procedure.
- Depending on your operating system and WebLogic versions, your WebLogic screens may be slightly different than those documented in the examples below.

### Overview of Procedures

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

- ["Install the IBM FileNet P8 Platform Documentation" on page 136](#) -- Begin here if you are installing the IBM FileNet P8 Platform documentation for the first time.

- [“Install Expansion Product Documentation” on page 137](#) -- Begin here if have already installed the IBM FileNet P8 Platform documentation and need to add on the documentation for IBM FileNet expansion products (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets).
- [“Update the Documentation Search Index” on page 138](#) -- Perform this procedure only after you have added on documentation for your expansion products. Otherwise, if you install only IBM FileNet P8 Platform documentation, which has a baseline search index, you may skip this procedure.
- [“Complete and Verify the Documentation Installation” on page 140](#) -- Perform this final procedure in all cases, but only after you have installed the IBM FileNet P8 Platform documentation, added on documentation for your expansion products, and (where required) regenerated the search index.

## ***Install the IBM FileNet P8 Platform Documentation***

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

### **To install the IBM FileNet P8 Platform documentation on WebLogic 8.x**

---

1. Access the IBM FileNet P8 Platform Documentation package.
2. Copy the **/ecm\_help** folder structure from the package to a location on the local hard drive (for example, **...P8docs/ecm\_help**).
3. If it is not running, start the WebLogic Server.
4. Start the WebLogic Server Administration Console.
5. From the WebLogic Server Administration Console, select **<mydomain>** > Deployments > Web Application Modules.
6. Click **Deploy a new Web Application Module...**
7. From the right pane of the WebLogic Server Administration Console, browse to and select the radio button for the **ecm\_help** folder.
8. Click **Target Module**.
9. Click **Deploy**.
10. Continue as follows:
  - If you need to add on documentation for expansion products, go on to the procedure in the following topic, [“Install Expansion Product Documentation” on page 137](#).
  - If you have no further documentation to install, then go on to the procedure in the topic [“Complete and Verify the Documentation Installation” on page 140](#).



### **To install the IBM FileNet P8 Platform documentation on WebLogic 9.x**

---

1. If this is a Windows server, verify that Microsoft Windows Internet Information Services (IIS Admin Service, Simple Mail Transport Protocol (SMTP), and World Wide Web Publishing Service are stopped and set to *manual*.
2. Access the IBM FileNet P8 Platform Documentation package.
3. Copy the **/ecm\_help** folder structure from the package to a location on the local hard drive (for example, **...P8docs/ecm\_help**).
4. If it is not running, start the WebLogic Server.
5. Start the WebLogic Server Administration Console.
6. From the WebLogic Server Administration Console, select **<mydomain>** > Deployments > Web Application Modules.
7. Click **Deployments**.
8. Click **Lock and Edit**.
9. Click **Install**.
10. Navigate to the documentation location.
11. Select the radio button for the **ecm\_help** folder and click **Next**.
12. Select **Install this deployment as an application** and click **Next**.
13. Name the site and click **Finish**.
14. Click **Activate Changes**.
15. Select the site (check) and select **Start > Servicing all requests** and click **Yes**.
16. Continue as follows:
  - If you need to add on documentation for expansion products, go on to the procedure in the following topic, [“Install Expansion Product Documentation” on page 137](#).
  - If you have no further documentation to install, then go on to the procedure in the topic [“Complete and Verify the Documentation Installation” on page 140](#).

### ***Install Expansion Product Documentation***

Use the procedure in this topic to install documentation for expansion products into an existing IBM FileNet P8 Platform documentation application. If you have no such documentation to add on, skip to the procedure [“Complete and Verify the Documentation Installation” on page 140](#).

#### **To install expansion product documentation**

---

1. Determine the documentation media source or location:
  - For most expansion products, use the *Documentation* package included as part of the particular software.

- For Content Federation Services for Image Services, use the **Documentation** directory on the software package.

**NOTE** Make sure you have copied the help for all your various expansion products to a designated application server location containing the IBM FileNet P8 Platform documentation. Otherwise, you will have to repeat this procedure if you add additional documentation later.

2. Identify the directory where you deployed the IBM FileNet P8 Platform documentation application. This is the directory location where you copied the expanded **ecm\_help**, or extracted **ecm\_help.war** to the **ecm\_help** directory.
3. Stop the WebLogic server where you deployed the IBM FileNet P8 Platform documentation application.
4. Copy the expansion product documentation to the IBM FileNet P8 Platform documentation server, as follows:
  - (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 Platform documentation server.  

```
cp -r <mount_location> <target_destination>
```

**WARNING** Care should be taken when copying folders in UNIX. Dragging-and-dropping of folders replaces any existing folder(s) of the same name. Also note that your switch (`-r`) requirements may be different from the example shown. Contact your system administrator if you have questions about proper syntax.
  - (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
  - (Content Federation Services for Image Services) Copy the **cfs\_guide.pdf** file from the **Documentation** directory in the software package into the **ecm\_help/cfs\_help** directory on the IBM FileNet P8 Platform documentation server.

**NOTE** Repeat this step for each of your expansion products. You can copy more than one expansion product documentation set to the documentation application server before continuing (for example, IBM FileNet P8 eForms and IBM FileNet P8 Portlets) so you end up with one **ecm\_help** directory structure containing multiple sets of expansion product documentation files added to it.

5. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM P8 Platform and various functional expansion products. Check the documentation page on the [IBM Information Management support page on www.ibm.com](#) for the latest versions of these guides. See “Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21 for details.
6. Go on to the procedure in the following topic, “[Update the Documentation Search Index](#)” on page 138.

## Update the Documentation Search Index

Perform this procedure only if you have refreshed the core IBM FileNet P8 Platform documentation, or installed expansion product documentation onto your IBM FileNet P8 Platform

documentation server. Otherwise, skip to the procedure in the following topic, [“Complete and Verify the Documentation Installation” on page 140](#).

**NOTE** Any time you update the documentation search index, a backup of the files in the existing **Index/core** directory will be automatically copied to the **Index/IndexOld** subdirectory. You can reapply these backed-up files to the **core** subdirectory (after first removing the new files created there) if you need to return to your previous indexed state.

### To update the documentation search index

---

1. Make sure the WebLogic server where you deployed the IBM FileNet P8 Platform documentation application is stopped. For the search indexer to run, no other processes can be accessing the IBM FileNet P8 Platform documentation.
2. Open a command prompt or terminal on the application server.
3. From the command line, navigate to the **search** subdirectory under the application root directory, for example, **ecm\_help**.
4. Using a text editor, open the search-indexing script file that is appropriate to your operating system:

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

(UNIX) **indexFiles.sh**

(Windows) **indexFiles.bat**

5. If necessary, set the JAVA\_HOME variable in the script file with the path to your JRE installation. The default examples are:

(UNIX) `JAVA_HOME="/usr/java/j2sdk1.4.1_02"`

(Windows) `SET JAVA_HOME=c:\j2sdk1.4.2`

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

6. Save your changes and close the text editor.

**NOTE** If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (`chmod 755`) to the **indexFiles.sh** file.

7. Run the following updated search-indexing script file that is appropriate to your operating system.

(UNIX) **indexFiles.sh**

(Windows) **indexFiles.bat**

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

8. Go on to the next procedure.

## Complete and Verify the Documentation Installation

Perform this procedure after you have installed (and, if necessary, reindexed) the IBM FileNet P8 documentation on the application server.

### To complete and verify the documentation installation on WebLogic

---

1. If the WebLogic server where you deployed IBM FileNet P8 documentation application is not running, start it.
2. Verify that the application server and the new IBM FileNet P8 documentation web site are running, as follows:
  - a. From your web browser, access the following URL. The documentation's Help Directory should open.

**`http://<docserver>:<port#>/<contextRoot>/`**

where:

*docserver* is the name of the Java web server.

*port#* is the port number (for example, 7001).

*contextRoot* is the value of the Map to URL field that you specified when you deployed the IBM FileNet P8 Platform documentation application. If you specified **`/ecm_help`**, then the *contextRoot* is **`ecm_help`**.

Example: **`http://yourdocserver:7001/ecm_help/`**

**NOTE** You can use multi-part root directories (for example, **`/docs/ecm_help`**) if your application server supports them.

- b. Click the **Search** link on the Help Directory toolbar. The documentation Search page should open.
- c. Select one of the Search query result links. The associated help page should open.

**NOTE** When it is time to configure the online help location for the various IBM FileNet P8 components, either while running Setup programs or later via site preferences settings, use the URL in [Step a on page 140](#).

## Task 7c: Install IBM FileNet P8 Platform Documentation (JBoss)

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

You must install the IBM FileNet P8 Platform documentation on a web application server if you intend to access online help from within any IBM FileNet P8 applications (for example, Workplace, Enterprise Manager, Process Task Manager, and Records Manager) or use the full-text search feature.

### NOTES

- You may install the IBM FileNet P8 Platform documentation and Application Engine on the same machine.
- A valid Java JSDK must be installed for the JBoss application server.
- To ensure proper documentation search functionality, enable JavaScript support on the user's browser client.
- Because of the number of possible network and web configurations, contact your network or web administrator for specific system, software, and security requirements.
- The requirements described in this topic may also exist for associated custom application help, or for any customizations you may have applied to the IBM FileNet P8 Platform documentation. This documentation requirement may also exist for associated custom applications. For further details on customizing or localizing IBM FileNet P8 Platform documentation, see the file [ecm\\_help/transkit/P8DocTransKit.htm](#) on the IBM FileNet P8 Platform Documentation package or wherever you install the IBM FileNet P8 Platform documentation.
- If you are upgrading your IBM FileNet P8 Platform installation and have downloaded a refresh of the documentation from the [IBM Information Management support page on www.ibm.com](#) see ["Upgrade IBM FileNet P8 Documentation" on page 478](#) for instructions on upgrading the documentation.

### Overview of procedures

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

- ["Install the IBM FileNet P8 Platform documentation" on page 142](#) - Begin here if you are installing the IBM FileNet P8 Platform documentation for the first time.
- ["Install functional expansion documentation" on page 142](#) - Begin here if have already installed the IBM FileNet P8 Platform documentation and need to add on the documentation for IBM FileNet functional expansions (e.g., Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, Records Manager, or IBM FileNet P8 Portlets).

- [“Update the documentation search index” on page 143](#) - Perform this procedure only after you have added on documentation for your functional expansions. Otherwise, if you install only IBM FileNet P8 Platform documentation, which has a baseline search index, you may skip this procedure.
- [“Complete and verify the documentation installation” on page 145](#) - Perform this final procedure in all cases, but only after you have installed the IBM FileNet P8 Platform documentation, added on documentation for your functional expansions, and (where required) regenerated the search index.

## ***Install the IBM FileNet P8 Platform documentation***

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

### **To install the IBM FileNet P8 Platform documentation**

---

1. Access the IBM FileNet P8 Platform Documentation package.
2. Create a directory for the IBM FileNet P8 Platform documentation. For example:

(UNIX) `/home/usr/apps/jboss_ver/server/default/deploy/p8_docs.war`

(Windows) `C:\jboss-4.0.5\server\default\deploy\p8_docs.war`

#### **NOTES**

- The installation directory must have a `.war` extension as shown above.
  - The installation directory can be user-defined, so substitute your actual location for these commands.
3. Unpack the `ecm_help.war` file or copy the `ecm_help` directory contents to the directory you just created.
  4. Continue as follows:
    - If you need to install documentation for functional expansions, go on to the procedure in the following topic, [“Install functional expansion documentation” on page 142](#).
    - If you have no further documentation to install, go on to the procedure in the topic [“Complete and verify the documentation installation” on page 145](#).

## ***Install functional expansion documentation***

Use the procedure in this topic to install documentation for functional expansions onto an existing IBM FileNet P8 Platform documentation server. If you have no such documentation to add on, skip to the procedure [“Complete and verify the documentation installation” on page 145](#).

### **To install functional expansion documentation**

---

1. Determine the functional expansion documentation media source location:

- For most functional expansions, use the *Documentation* package included as part of the particular software.
  - For Content Federation Services for Image Services, use the **Documentation** directory on the software package.
2. If the application server is running, use the following command to stop it on the machine where the documentation is installed, and verify that no processes are accessing the documentation web application:

```
(UNIX) <path> /shutdown.sh -S
```

```
(Windows) <path> \shutdown.bat -S
```

3. Copy the functional expansion documentation to IBM FileNet P8 Platform documentation server, as follows:

- (UNIX) When copying functional expansions, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 Platform documentation server.

```
cp -irv --reply=no <mount_location> <target_destination>
```

**WARNING** Care should be taken when copying folders in UNIX. Dropping-and-dragging of folders 'replaces' any existing folder(s). Also note that your switch (`-irv`) requirements may be different from the example shown. Contact your system administrator if you have questions about proper syntax.

- (Windows) Use a copy command from a command prompt or drop-and-drag the files to the destination.
- (Content Federation Services for Image Services) Copy the **cfs\_guide.pdf** from the **Documentation** directory in the software package into the **ecm\_help/cfs\_help** directory deployed on the existing IBM FileNet P8 Platform documentation server.

**NOTE** Repeat this step for each of your functional expansions or custom applications. You can copy more than one functional expansion documentation set to the documentation application server before continuing (for example, eForms) so you end up with one **ecm\_help** directory containing multiple sets of functional expansion files added in.

4. Download the latest web-posted updates of installation and upgrade guide PDFs for the IBM P8 Platform and various functional expansion products. Check the documentation page on the [IBM Information Management support page on www.ibm.com](#) for the latest versions of these guides. See “Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21 for details.
5. Continue with “Update the documentation search index” on page 143.

## Update the documentation search index

Perform this procedure only if you have refreshed the core IBM FileNet P8 Platform documentation, or installed functional expansion documentation onto your IBM FileNet P8 Platform documentation server. Otherwise, skip to the procedure in the following topic, “[Complete and verify the documentation installation](#)” on page 145.

**NOTE** Any time you update the documentation search index, a backup of the files in the existing **Index/core** directory will be automatically copied to the **Index/IndexOld** subdirectory. You can reapply these backed-up files to the **core** subdirectory (after first removing the new files created there) if you need to return to your previous indexed state.

### To update the documentation search index

---

1. If the JBoss server is running, use the following command to stop it on the machine where the documentation is installed, and verify that no processes are accessing the documentation web application:

(UNIX) `<path> /shutdown.sh -S`

(Windows) `<path> \shutdown.bat -S`

**NOTE** Make sure you have copied the help for all your various functional expansions to a designated application server location containing the IBM FileNet P8 Platform documentation. Otherwise, you will have to repeat this procedure if you add new help later.

2. Open a command prompt (or terminal) on the application server.
3. From the command prompt (or terminal), navigate to the search subdirectory under your `ecm_help` root directory.
4. Using a text editor, open the search-indexing script file that is appropriate to your application server operating system:

**NOTE** You may need to set the permissions on this file as it is set to read-only on the documentation package.

(UNIX) `indexFiles.sh`

(Windows) `indexFiles.bat`

5. Modify the `JAVA_HOME` variable in the script file with the path to your JRE installation. The default examples are:

(UNIX) `JAVA_HOME="/usr/java/j2sdk1.4.1_02"`

(Windows) `SET JAVA_HOME=c:\j2sdk1.4.2`

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

6. Save your changes and close the text editor.

**NOTE** If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (`chmod 755`) to the **indexFiles.sh** file.



7. Run the following updated search-indexing script file that is appropriate to your application server operating system.

(UNIX) `indexFiles.sh`

(Windows) `indexFiles.bat`

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

8. Go on to the next procedure.

## Complete and verify the documentation installation

Perform this procedure after you have installed (and, if necessary, reindexed) the IBM FileNet P8 Platform documentation on the application server.

### To complete and verify the JBoss documentation installation

---

1. Start the IBM FileNet P8 Platform documentation web application by running the following JBoss command:

(UNIX) `run.sh`

(Windows) `run.bat`

2. Verify that the application server and the new IBM FileNet P8 Platform documentation web site are running, as follows:
  - a. From your web browser, access the following URL. The documentation's Help Directory should open.

**`http://<docserver>:<port#>/<contextRoot>/`**

where:

**`docserver`** is the name of the Java web server.

**`port#`** is the port number (def: 80 or 8080).

**`contextRoot`** is the value of the Map to URL field that you specified when you deployed the IBM FileNet P8 Platform documentation application. If you specified `/ecm_help`, then the **`contextRoot`** is `ecm_help`.

Example: **`http://yourdocserver:80/ecm_help/`**

**NOTE** You can use multi-part root folders (e.g., `/docs/ecm_help`) if your application server supports them.

- b. Click the **Search** link on the Help Directory toolbar. The documentation Search page should open.
- c. Select one of the Search query result links. The associated help page should open.

**NOTE** When it is time to configure the online help location for the various IBM FileNet P8 components, either while running Setup programs or later via site preferences settings, use the URL in [Step a on page 145](#).



## Task 8: Install and Configure Content Search Engine

This task describes how to install and configure the initial Administration Server for the IBM FileNet P8 Content Search Engine, an optional component based on the Autonomy K2 product. In effect, you will set up an Autonomy K2 Master Administration Server required for any single and multi-server K2 configuration.

The Autonomy K2 software that underlies the IBM FileNet P8 Content Search Engine has many inherent features not discussed in the IBM FileNet P8 Help that you might want to configure. For details see the Autonomy documentation set that is installed on the Autonomy K2 Master Administration Server. The K2 PDF/HTML documentation set is located at: [http://<hostname>:9990/verity\\_docs/](http://<hostname>:9990/verity_docs/) after installation. This K2 documentation set is not searchable from the IBM FileNet P8 Help but does have its own internal index and search functionality.

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

### NOTES

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

## To install Autonomy K2 Master Administration Server on Windows

---

1. Create the required accounts and groups with related permissions as specified in [“To create Content Search Engine accounts on page 60”](#).

2. Access the host machine and log onto the directory service as K2 Operating System User .

**NOTE** Ensure K2 Operating System User is an operating system administrator on this machine.

3. Insert the IBM FileNet Autonomy K2 installation CD and extract the contents of **K2-win.zip** to the following location:

**C:\Program Files\filenet\contentengine\verity\**

4. Set the Java\_Home environment variable as follows:

- a. Open the System control panel.
- b. Click the **Advanced** tab.
- c. Click **Environment Variables**.
- d. Click **New** under System Variables.
- e. Set the variable information as follows:

Variable name: **Java\_Home**

Variable value: `<Java1.5.0xx_JDK_install_path>`

5. Open **C:\Program Files\filenet\contentengine\verity\config.vcnf** in a text editor and make the following modifications in the file:

- Replace the instance of `<myMode>` with Master.
- Replace all instances of `<myLocalHostName>` with the machine name you are installing on.
- Replace all instances of `<myMasterHostName>` with the machine name you are installing on.
- Replace all instances of `<installDir>` with **C:\Program Files\filenet\contentengine\verity**.
- Replace the instance of `<JavaHome>` with the path to the installed Java 1.5.0\_xx JDK.

6. Open a command line and change directory to **C:\Program Files\filenet\contentengine\verity**.

7. Enter the following at the command line:

```
k2\_nti40\bin\vconfig -cfg "C:\Program
Files\filenet\contentengine\verity\config.vcnf" -dir "C:\Program
Files\filenet\contentengine\verity" -verbose -log log.txt
```

The Autonomy K2 Administration Server service will be installed and running at the completion of the vconfig command.

8. Close the command window.

**NOTE** It is important that you close this command window instance.

9. Update the documentation XML file and set the access path:
  - a. Access **C:\Program Files\filenet\contentengine\verity\appserver\conf\Catalina\localhost**.
  - b. Right-click the **k2\_docs.xml** file and select **Edit**.
  - c. Replace `<installDir>` with **C:\Program Files\filenet\contentengine\verity**.

For example:

```
<Context path="/verity_docs" docBase="c:\program
files\filenet\contentengine\verity\data/docs"
    debug="0" reloadable="true" crossContext="false">
</Context>
```

10. Update the K2 Dashboard shortcut:
  - a. Access **C:\Program Files\filenet\contentengine\verity**.
  - b. Right-click the **K2 Dashboard** shortcut and select **Properties**.
  - c. Replace `<myhostname>` with the name of the machine on which you are installing.
11. Update the K2 Doc shortcut:
  - a. Access **C:\Program Files\filenet\contentengine\verity**.
  - b. Right-click the **K2 doc** shortcut and select **Properties**.
  - c. Replace `<myhostname>` with the name of the machine on which you are installing.
12. Update **web.xml**:
  - a. Access **C:\Program Files\filenet\contentengine\verity\data\docs\WEB-INF**.
  - b. Right-click **web.xml** and select **Edit**.
  - c. Replace `<myHostName>` with the machine name on which you are installing.
13. Open a new command line window and change directory to **C:\Program Files\filenet\contentengine\verity\appserver\bin**.

14. Enter the following command:

```
service.bat install k2
```

15. Enter the following command to launch the Tomcat application server:

```
startup
```

16. Set K2 Administration Server service to run as K2 Operating System User, as follows:
  - a. Access Component Services.
  - b. Stop the Verity K2 6.1.1 Administration Server service.
  - c. Change the logon settings and set the service to Log On as K2 Operating System User.
  - d. Start the Verity K2 6.1.1 Administration Server service.

- See ["To configure Autonomy K2 for Content-Based Retrieval on page 153"](#) to configure this Autonomy K2 installation.

**To install Autonomy K2 Master Administration Server on UNIX**

**NOTE** For HP-UX installations, manually configure the kernel with following parameters before you begin the Autonomy K2 Master Administration Server installation:

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

- Create the required accounts and groups and related permissions as specified in ["To create Content Search Engine accounts on page 60"](#).
- Log onto the UNIX machine as the K2 Operating System User.
- Insert the IBM FileNet Autonomy K2 installation CD and extract the contents of K2-*<platform>*.tar.gz to **/opt/verity** using the following commands:
  - `gzip -d <platform>.tar.gz`
  - `tar -xvf <platform>.tar`

**NOTE** If you install to a directory other than **/opt/verity**, you must create a soft link using the following command:

```
ln -s <InstallPath> /opt/verity
```

- Change directory to **/opt/verity**.
- Edit **/opt/verity/config.vcnf** as follows:
  - Replace the instance of *<myMode>* with Master.

- Replace all instances of *<myLocalHostName>* with the name of the machine on which you are installing.
- Replace all instances of *<myMasterHostName>* with the name of the machine on which you are installing.
- Replace the instance of *<JavaHome>* with the path to the installed Java 1.5.0\_xx JDK.

6. Set the Java\_home variable:

```
JAVA_HOME=<java_install_path>/jdk1.5.0_xx
export JAVA_HOME
```

**NOTE** Place the above entry in a user's **.profile** file to make the setting available each time the user logs in.

7. Append the following environment variables according to your platform:

(HP-UX)

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

(AIX)

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

(Solaris)

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

(Linux)

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

8. Change directory to **/opt/verity/**

9. Enter the following command:

```
k2/<platform>/bin/vconfig -cfg "/opt/verity/config.vcnf" -dir "/opt/verity" -
verbose -log log.txt
```

Substitute one of the following for *<platform>*:

- *\_ssol26* (Solaris 8.0, 9.0 or 10.0)
- *\_hpux* (HP-UX 11i with -AA compiler flag)
- *\_rs6k43* (AIX 5.2 and 5.3)
- *\_ilnx21* (Red Hat Advanced Server 3.0 and 4.0, SUSE 8 and 9)

The Autonomy K2 Administration Server service and Tomcat will be installed and running at the completion of the vconfig command.

**NOTE** To manually start or stop the Autonomy K2 Administration Server service, use the follow commands, according to your environment:

(HP-UX)

Start Service:

```
nohup /<device >/<verity_install_dir>/k2/_hpux/bin/k2adminstart &
```

Stop Service:

```
/<device >/<verity_install_dir>/k2/_hpux/bin/k2adminstop
```

(AIX)

Start Service:

```
nohup /<device >/<verity_install_dir>/k2/_rs6k43/bin/k2adminstart &
```

Stop Service:

```
/<device >/<verity_install_dir>/k2/_rs6k43/bin/k2adminstop
```

(Solaris)

Start Service:

```
nohup /<device>/<verity_install_dir>/k2/_ssol26/bin/k2adminstart &
```

Stop Service:

```
/<device>/<verity_install_dir>/k2/_ssol26/bin/k2adminstop
```

(Linux)

Start Service:

```
nohup /<device >/<verity_install_dir>/k2/_ilnx21/bin/k2adminstart &
```

Stop Service:

```
/<device >/<verity_install_dir>/k2/_ilnx21/bin/k2adminstop
```



10. See ["To configure Autonomy K2 for Content-Based Retrieval on page 153"](#) to configure this Autonomy K2 installation.

### **To install additional Autonomy K2 Administration Servers**

---

You will likely need to install additional Administration Servers to handle the indexing load. The Master Administration Server is the main hub for configuring all the servers you install. The additional servers are managed through the K2 Dashboard that is installed with the Master Administration Server.

To install additional Autonomy K2 Administration Servers, follow the procedure outlined in [Task 46 "Install Additional Content Search Engine Servers" on page 447](#), then follow the procedure below to configure the new Administration Server.

### **To configure Autonomy K2 for Content-Based Retrieval**

---

This procedure outlines how to create and configure the minimum set of server services required on the K2 Master Administration Server, and on additional Administration Servers you may install, for IBM FileNet P8 Content Engine. All servers are configured through the Master Administration Server Dashboard.

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. If multiple Brokers are assigned, then if one goes down the others will be used. But, each broker must have all K2 Servers (search servers) attached that are needed to access collections (index areas). The Content Engine Server will not call multiple brokers and merge the results.
- If you add additional Index Servers and K2 Servers (search servers), they will not be activated until you enable them through Enterprise Manager. See ["To enable additional K2 Index Servers and Search Servers on page 401"](#) for details.
- Each K2 Administration Server must contain a Ticket Server for Content Engine.
- For good stability and performance, it is recommended that Broker Servers be attached to local Ticket Servers for security on each machine.

### **NOTES**

- When naming particular servers you create with this procedure, it's a good idea to indicate the type of server you've created. Otherwise, when you configure Content Engine, discerning which server is which could be confusing. For example, `<servername>_broker` to indicate that this is a Broker Server.
  - 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.
  - The following procedures suggest the recommended ports, however, you may choose which ports to use. Multiple servers will require multiple ports.
1. Configure the Autonomy K2 Dashboard to use SSL security. The Autonomy K2 Dashboard web application, by default, uses a non-SSL web site and sends username and password information

in plain text. For information on how to modify your Tomcat web applications to use SSL, access <http://tomcat.apache.org/tomcat-5.5-doc/ssl-howto.html>

2. Access the K2 Dashboard by launching your browser and entering: *http://<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. If multiple Administration Servers are installed, select the server on which you want to create the service and click **continue**.
  - d. Enter the following information on the *Configure basic settings for the new K2 Index Server* page:
    - Service Alias: *<server\_name>\_index\_server*
    - Port: 9960
  - e. Click **Next** to continue with the installation.
  - f. Enter the following information on the *Configure threads for the K2 Index Server* page:
    - Synchronous Threads: 25
    - Asynchronous Threads: 3
    - Access Type: Authorized Administrator
  - g. 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. If multiple Administration Servers are installed, select on which of these you want to create the service and click **continue**.
  - d. Enter the following information on the *Configure basic settings for the new K2 Broker* page:
    - Service Alias: *<servername>\_broker*

- Port: 9900
  - e. 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. If multiple Administration Servers are installed, select which Administration Server you want to create the service on and click **continue**.
  - d. Enter the following information on the *Configure basic settings for the new K2 Server* page:
    - Service Alias: `<server_name>_search_server`
    - Port: 9920
  - e. Click **Next**.
  - f. Click **Next** on the *Set security options for this service* page.
  - g. Enter the following information on the Attach to K2 Brokers page:
    - Select the K2 Brokers that will provide access to this service: Select the K2 Broker you created in step 3 from the drop-down menu, `<servername>_broker`.
  - h. Click **Finish**.
7. Import the IBM FileNet Styleset.
- NOTE** Using a stop words file will increase your performance by about 30%. You can put a file named **style.stp** into the stylefiles it will not full-text index short words like a, the, and. This involves copying vdk30.stp over to stylefile directory and renaming it to **style.stp**. See the Verity documentation for more details.
- 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.
  - d. Enter the following information on the Import page:
    - Style Set Alias: **FileNet\_FileSystem\_PushAPI**
    - Gateway Type: `--Auto-detect--`
    - Source Administration Server (if multiple servers are installed). Choose which server.
    - Source Path:
      - (Windows) **C:\Program Files\filenet\contentengine\verity\data\stylesets\FileNet\_FileSystem\_PushAPI**
      - (UNIX) **/opt/verity/data/stylesets/FileNet\_FileSystem\_PushAPI**
  - e. Click **Import**.
8. Create a K2 Ticket Server.

- a. Click **K2 Ticket Servers** under System View.
  - b. Click **Add a K2 Ticket Server** under Actions on the K2 Ticket Server Summary page.
  - c. If multiple K2 Administration Servers are installed, select on which you want to create the service and click **continue**.
  - d. Enter the following information on the *Configure basic settings for the new K2 Ticket Server* page:
    - Service Alias: `<server_name>_ticket_server`
    - Port: 9910
  - e. Click **Next**.
  - f. Enter the following information on the *Configure the login module to use with this K2 Ticket Server* page:
    - Select which Login Module type to use with this K2 Ticket Server:
      - Windows
      - UNIX

**NOTE** LDAP Ticket Servers are not currently supported.
    - Default Domain (Windows only): Enter the domain on which this K2 Server is authenticated.
  - g. Click **Next**.
  - h. 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.
  - i. Click **Finish**.
  - j. Click **Edit Properties**.
  - k. Click Windows Login Module.
  - l. Check **Use Local Credentials**.
  - m. Check Enable Built-in Groups.
  - n. Click **Modify**.
9. Set Autonomy K2 Administration Security.
    - a. Click the K2 Ticket Server you created.
    - b. Click **Manage Administration Security** under *Actions*.
    - c. Enter the following information on the Manage Administration Security page:
      - Select a K2 Ticket Server to configure for administration security: From the drop-down menu, select the K2 Ticket Server you just created.

- User Name: Enter the K2 Operating System User. For UNIX installs, this is the user you logged in as to run the install in [Step 2 on page 150](#). For more detail on required accounts, see ["To create Content Search Engine accounts on page 60"](#).
- Password: Enter the authentication password.
- Default Domain: Enter the domain on which this user and K2 Server are authenticated.

d. Click **Modify**.

K2 will authenticate the user using 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 in again as the Dashboard Administrator to complete the configuration.

10. Launch the K2 Dashboard and log in.

11. Restart K2 services:

- a. Under *Notifications* on the K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You'll need to perform either a **Quick Restart** or a **Full Restart**.

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

12. Enable additional K2 Admin Users (optional).

- a. From the K2 Dashboard home page, click **Administration Servers**.
- b. Click **Manage K2 Administrative Users**.
- c. Click **Add User** on the Manage K2 Administrative Users page.
- d. Enter the name of an authenticated user on the directory service that you want to make a K2 Administrator and click **Add**.

13. Enable security on the K2 services you have created.

- a. From the K2 Dashboard home page, click **K2 Ticket Servers**.
- b. Click your ticket server `<ticket_servername>`.
- c. Click **Manage K2 Broker/K2 Server Security** in the *Services Secured by this K2 Ticket Server* section at the bottom of the page.
- d. Click the **K2 Servers** button on the Manage K2 Broker/K2 Server Security page.
- e. One-at-a-time, click the service listed in the window on the right to enable security.
- f. Click the **K2 Brokers** button on the Manage K2 Broker/K2 Server Security page.
- g. One-at-a-time, click each service listed in the window on the right to enable security.
- h. Click **Modify** to save your changes.

14. Restart K2 services, as follows:

- a. Click **Home** in the top left corner of the page.

- b. Under Notifications on the Verity K2 Dashboard Home page, a number of servers are listed as requiring a restart. Click **Start/Stop this Service** to access the settings page and follow the instructions listed there. You'll need to perform either a **Quick Restart** or a **Full Restart**.

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

15. Install additional locales from CD. Complete this step only if you require locales other than English.

Windows:

- a. Create a directory <Program Files>\Common Files\InstallShield\Universal\<Win platform>\x86\<machine\_name>\Gen1\\_vpddb

<Program Files> is the system Program Files folder, such as C:\Program Files

<Win platform> is the Windows version. Valid options:

- Windows XP
- Windows 2000
- Windows 2003

<machine\_name> machine name on which you are installing

Example: C:\Program Files\Common Files\InstallShield\Universal\Windows 2003\x86\myMachine\Gen1\\_vpddb

- b. Locate the vpd.script file in the K2 CLI setup root folder. Modify all the variable instances in the string as follows:

- Replace <K2InstallDir> with the K2 root folder. For example, C:\Program Files\FileNet\ContentEngine\verity
- Replace <myHostName> with the machine name.
- Replace <WinVersion> with the Windows version.

Ensure the entries in the vpd.script file match the folder path in [Step a](#).

- c. Copy the vpd.script to the location in [Step a](#) of this procedure.
- d. Run the installer. Navigate to the decompressed locale installer location and execute the respective installer launcher. The locale installer will locate the K2 installation and start, based on the settings you completed above.

**NOTE** The locale installer requires the following license key: 2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS

UNIX:

To install locales from the Locale CD after Content Search Engine (Autonomy K2) has been installed, prepare your system with the following:

export variable VERITY\_CFG

VERITY\_CFG=/opt/verity/k2/common/verity.cfg

```
export VERITY_CFG
```

To run the installer, navigate to the decompressed locale installer location and execute the respective installer launcher. The locale installer will locate the K2 installation and start, based on the settings you completed above.

**NOTE** The locale installer requires the following license key: 2UV4MPT-2KPEQBJ-1D6A6KT-2KPE6KT-2KPE6KS

16. If you are adding Content Search Engine to an already functioning and updated IBM FileNet P8 Platform system, skip to [Task 18 Install Content Search Engine Software Updates](#) to apply any available service packs, interim fixes, or fix packs and then complete [Task 29 Configure Content Engine for Content-Based Retrieval](#). Otherwise, move on to [Task 9 "Install and Deploy Content Engine" on page 160](#).

## Task 9: Install and Deploy Content Engine

The procedures in this task show how to install and deploy Content Engine on an application server. Do these procedures whether you are installing Content Engine as a new application or upgrading from an old version.

**CAUTION** Before doing any procedures in this task, be sure you have completed all the applicable tasks in [“Prerequisite Tasks” on page 53](#).

After upgrading from an old version, you will also need to run utilities to migrate items (such as databases and object stores) from the old version of Content Engine to the new version (see [“Upgrade Content Engine Software” on page 487](#)).

You can install Content Engine interactively, via Content Engine Setup, or silently, via a command-line interface.

Before doing a silent install, you must first record a *response file* (also known as an *options file*) by running Content Engine Setup interactively, typically in a development or test setting. You then use the response file to run a silent install of Content Engine in a production environment.

In your production environment, install Content Engine first on the application server master (administrative) node. You can then use the administrative console of the application server to deploy Content Engine to other managed servers (see [“Deploy Content Engine to Other Application Servers” on page 437](#)).

You can install some or all of the following Content Engine components on a given machine, depending on the specifics of your configuration:

- **Content Engine Server.** Your first installation of Content Engine must include this component.
- **Application Server Authentication Provider.** If your LDAP provider uses single sign-on, create the provider (through the administrative console of your application server) prior to installing Content Engine.
- **Administration Tools - Enterprise Manager and COM Compatibility Client API.** These are .NET clients that can be installed only on a Windows machine. (In this document, the terms “Enterprise Manager” and “IBM FileNet Enterprise Manager” refer to the same component and are used interchangeably.)

Your initial installation of Content Engine does not need to include Enterprise Manager. However, until you install Enterprise Manager, you will not be able to administer Content Engine.

**CAUTION** You cannot install version 4.0.x of Enterprise Manager on a machine where version 3.5.x of the component is already installed.

**NOTE** On each HP-UX 11 or HP-UX 11i machine where you will install Content Engine (a JVM-based component), ensure that the values of the kernel parameters *max\_thread\_proc* (maximum number of threads per process) and *nkthread* (maximum number of kernel threads in the system) are larger than their default values, which are too small for IBM FileNet P8. Refer to [“Operating System Considerations” on page 25](#) for information on specifying appropriate values.

At the HP web site, refer to the document titled *HP-UX Programmer’s Guide for Java 2* for tools to determine values of these two kernel parameters that suffice for IBM FileNet P8.



The machine where you will install Content Engine Server is referred to in the following procedure as the *Content Engine machine*. The procedure assumes you have a WebLogic domain called *FNCEDomain*, a WebSphere profile called *default*, or a JBoss server instance. Substitute your own domain, profile, or server name in place of a default value.

**WARNING** If you need to rename a machine, do so before installing Content Engine on it.

**To run Content Engine Setup**

1. Log on to the machine where you will run the Content Engine Setup program. Logging on as a member of the local Administrators group (Windows) or the root user (UNIX) gives you sufficient privileges to run the program. On most operating systems, you can run the program with lesser privileges, as shown in the following table:

Operating System	User Requirements for Running Content Engine Setup
AIX	Must be root user
HP-UX 11 HPUX 11i Linux Solaris	<ul style="list-style-type: none"> <li>• Must be a member of the adm group</li> <li>• Must have read/write/execute permission to the following:                             <ul style="list-style-type: none"> <li>– Device/location where Content Engine is to be installed</li> <li>– Device/location where the application server has been installed</li> </ul> </li> <li>• Must have write permission to the directories where you will create file storage areas, index areas, and content caches</li> <li>• Must have write permission on the <b>/tmp</b> directory</li> <li>• Must be the user (or a member of the group) that launched the application server instance where you will deploy Content Engine</li> </ul>
Windows	<ul style="list-style-type: none"> <li>• Must be a member of the Local Administrators group</li> <li>• If application server is WebLogic, must also have write permission to directory <b>&lt;WebLogic_Install_Directory&gt;/user_projects/domains/FNCEDomain</b></li> </ul>

2. (WebLogic and WebSphere) Stop and start the application server on the Content Engine.

**NOTE** On UNIX, you can start WebSphere as the root user, or as a normal user with read/write/execute permissions on **<WebSphere\_Install\_Directory>/profiles/<profile\_name>**.

3. (WebSphere on Windows only) If Windows Active Directory is going to be your directory service, set the primary DNS server IP address on your Content Engine machine to the IP address of the machine where DNS is installed.

If you previously uninstalled and undeployed version 4.0.x of Content Engine on this machine, delete the following directory if you are going to run Content Engine Setup interactively:

**C:\Program Files\Common Files\InstallShield\Universal\FileNet\ContentEngine**

- (UNIX-based application server only) Use the UNIX utility program `umask` to set the default file-creation permissions mask for the JVM instance that will host Content Engine Server so that the owner (the user running JVM) and the members of the owner's group have read/write/execute access permissions, and all others have no access:

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

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

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

- Access the Content Engine software package, and navigate to the **ContentEngine** directory.

Do one of the following to start Content Engine Setup, provided you have *not* yet applied a service pack or interim fix to the Content engine software on this machine:

- If you are doing an interactive (non-silent) install, run one of the commands in the following table:

Platform	Command line
AIX	P8CE-4.0.0-AIX.bin
HP-UX 11	P8CE-4.0.0-HPUX.bin
HP-UX 11i	On WebLogic 9.2 (replacing <b>/opt/java 1.5</b> with the path name to your version 1.5 JVM): P8CE-4.0.0-HPUXi.bin -is:javahome /opt/java1.5  On non-WebLogic 9.2: P8CE-4.0.0-HPUXi.bin
Linux	P8CE-4.0.0-Linux.bin
Solaris (SPARC)	P8CE-4.0.0-Sol.bin
Windows	P8CE-4.0.0-Win.exe

- If you are going to record a response file for a subsequent silent install, run one of the commands in the following table, replacing **/opt/response.txt** or **C:\response.txt** with the path and name of your response file, provided you have *not* yet applied a service pack or interim fix to the Content engine software on this machine:

Platform	Command line
AIX	P8CE-4.0.0-AIX.bin -options-record '/opt/response.txt'
HP-UX 11	P8CE-4.0.0-HPUX.bin -options-record '/opt/response.txt'
HP-UX 11i	On WebLogic 9.2 (replacing <b>/opt/java 1.5</b> with the path name to your version 1.5 JVM):  P8CE-4.0.0-HPUXi.bin -is:javahome /opt/java1.5 -options-record '/opt/response.txt'  On non-WebLogic 9.2:  P8CE-4.0.0-HPUXi.bin -options-record '/opt/response.txt'
Linux	P8CE-4.0.0-Linux.bin -options-record '/opt/response.txt'
Solaris (SPARC)	P8CE-4.0.0-Sol.bin -options-record '/opt/response.txt'
Solaris (x86)	P8CE-4.0.0-Solx86.bin -options-record '/opt/response.txt'
Windows	P8CE-4.0.0-Win.exe -options-record 'C:\response.txt'

However, if you are running Content Engine Setup on a machine *after* having applied a service pack or interim fix to the Content Engine software on that machine, you must disable version checking by supplying two additional parameters to the command that starts Content Engine Setup:

```
-W updateCheck.active=false -W installCheck.active=false
```

For example, to do an interactive install on a Windows machine, run this command:

```
P8CE-4.0.0-Win.exe -W updateCheck.active=false -W installCheck.active=false
```

**NOTE** The first Content Engine Setup screen may not appear for 30 seconds or more after you launch the command.

**CAUTION** If it detects an earlier version of Content Engine software on this machine, Content Engine Setup displays a warning message: If version 3.5.x of Enterprise Manager is running on the machine, and you want to retain it to prepare object stores for upgrade to version 4.0.x, do *not* select .NET Clients for installation in the Select Components wizard screen (see [“Select Components” on page 164](#)). After you have upgraded all object stores to version 4.0.x, you can run Content Engine Setup again on this machine to install .NET clients.

6. Complete the Content Engine Setup screens as follows:

In this screen...	Perform this action...
Welcome	<p>Click <b>Next</b> to proceed with the Content Engine installation.</p> <p><b>NOTE</b> Click <b>Back</b> at any time while running Content Engine Setup to make changes in any previous screens. Click <b>Cancel</b> to exit Content Engine Setup.</p>
License Agreement	<p>Review and accept the license agreement, and then click <b>Next</b>.</p>
Specify Installation Location and Local Host Name	<p>Accept the default (or browse to another) Content Engine installation directory, specify the local host name or IP address, and then click <b>Next</b>.</p> <p><b>NOTE</b> (UNIX) The account running Content Engine Setup must have read, write, and execute access to both the installation directory and the <b>/tmp</b> directory.</p>
Select Components	<p>From the list on the screen, select components to be installed and then click <b>Next</b>:</p> <ul style="list-style-type: none"> <li>Content Engine Server</li> </ul> <p>You must select this component in your initial installation of Content Engine software.</p>

In this screen...	Perform this action...
Select Components	<ul style="list-style-type: none"> <li>Application Server Authentication Provider                     <p>Select this component if you want the installer to create and configure J2EE application server authentication in a later screen of Content Engine Setup.</p> <p><b>WEBSHERE</b> WebSphere 6.0.x and earlier support only a single user registry for authentication. Do not select this component if you have already configured WebSphere's <b>Global Security &gt; LDAP User Registry</b>. If you do, Content Engine Setup will overwrite that configuration with information provided during Content Engine installation. This could have unintended consequences on your environment, including any non-IBM FileNet P8 applications already deployed on this application server. <b>CAUTION</b> For WebSphere 6.1: do not select this component if you have already configured WebSphere authentication; if you are using WebSphere 6.1 federated user repositories, you should already have done the procedure <a href="#">"To use WebSphere 6.1 federated user repository"</a> on page 108.</p> <p><b>WEBLOGIC</b> WebLogic supports multiple security realms and multiple authentication providers per realm. If you select this component, Content Engine Setup will add a new authentication provider to any that already exist. In WebLogic 8, Content Engine uses the default security realm with the default name of "myrealm". In WebLogic 9, Content Engine uses the default security realm; if there is more than one, the <b>Summary of Security Realms</b> in the WebLogic Administration Console displays which is the default.</p> <p><b>JBOSS</b> JBoss supports multiple security realms. If you select this component, Content Engine Setup will edit the JBoss login-config.xml file and add several new &lt;authentication&gt; sections to join any that already exist.</p> <p><b>NOTE</b> If you wish to configure the application server's authentication yourself, rather than having Content Engine Setup do so, exit Content Engine Setup now and perform these tasks as described in the administration guide of your application server, and then run Content Engine Setup again, making sure not to select this component to install.</p> </li> </ul>

In this screen...	Perform this action...
Select Components	<ul style="list-style-type: none"> <li>• .NET Clients                             <p>The .NET client software is installable only on Windows. If Content Engine Setup is running on UNIX, you can install the .NET software on a Windows machine (see <a href="#">“Install an Additional Instance of Enterprise Manager” on page 427</a>) after Content Engine Setup finishes.</p> <p><b>NOTE</b> If you intend to install Records Manager, you must select the COM Compatibility Clients API component.</p> </li> <li>• Java Clients                             <p>If you created object stores in a previous version of Content Engine, select Object Store 3.5 to 4.0 Upgrade Tool to install the tool needed to upgrade them. This component is installable only on Windows. You will run this component in <a href="#">“Upgrade Content Engine Software” on page 487</a>.</p> </li> </ul>
.NETFramework and WSE Requirements	<p>This screen appears only if all the following conditions apply:</p> <ul style="list-style-type: none"> <li>• Content Engine Setup is running on Windows</li> <li>• Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0 are not both installed on the application server where you are running Content Engine Setup</li> <li>• You selected .NET Clients or one of its subcomponents in <a href="#">“Select Components” on page 164</a></li> </ul> <p>If you have later or compatible versions of these components, click <b>Next</b>; otherwise, click <b>Cancel</b> to exit Content Engine Setup and install the two components, as you cannot proceed beyond this screen.</p>
Specify Documentation URL <p><b>NOTE</b> This screen appears only if Content Engine Setup is running on Windows.</p>	<p>This URL is for accessing IBM FileNet P8 documentation from Enterprise Manager and other IBM FileNet P8 components.</p> <p>If you have already installed IBM FileNet P8 documentation, as instructed earlier, and want to test access to it now, type its URL, click <b>Test</b>, and then click <b>Next</b>.</p> <p>If you prefer to specify the URL after Content Engine Setup finishes, click <b>Next</b> and continue at <a href="#">“Choose Application Server” on page 166</a>.</p>
Choose Application Server	<p>Click the option button for the application server where you are installing Content Engine.</p> <p>If applicable, select an application server version from the drop-down list, and then click <b>Next</b>.</p>

In this screen...	Perform this action...
Review Installation Summary	<p>Verify your component selections and click <b>Install</b> to start installing Content Engine.</p> <p>Depending on your application server, continue at one of the following screens:</p> <ul style="list-style-type: none"> <li>• <a href="#">“Specify WebLogic Directories” on page 167</a></li> <li>• <a href="#">“Specify WebSphere Directory” on page 168</a></li> <li>• <a href="#">“Jboss Configuration” on page 169</a></li> </ul>
Specify WebLogic Directories	<p>Accept the default pathname, or browse to another location, for the following three directories:</p> <ul style="list-style-type: none"> <li>• WebLogic Root Directory</li> <li>• WebLogic Configuration Tool Directory</li> <li>• WebLogic Domains Directory</li> </ul> <p>Click <b>Next</b>.</p>

In this screen...	Perform this action...
Specify WebLogic Information	<p>Specify the following information and then click <b>Next</b>:</p> <ul style="list-style-type: none"> <li>• Domain Name: The name of the WebLogic domain (default is <i>FNCEDomain</i>) where this Content Engine instance will be installed.                             <p><b>NOTE</b> The WebLogic domain name is case sensitive. An incorrectly specified name will cause installation of Content Engine to fail.</p> </li> <li>• Administration Server Name: The WebLogic Server instance (administration server) where Content Engine will be deployed.                             <p><b>NOTE</b> If you are installing Content Engine in a managed server environment, specify the host name of the Administrator node.</p> </li> <li>• WebLogic Administration Server Port: The port (7001, by default) assigned to WebLogic.</li> <li>• Administrator Login Name: The user name (weblogic, by default) for logging on to WebLogic Server Administration Console. This is the same user you specified in one of the following:                             <ul style="list-style-type: none"> <li>– <a href="#">Step 5 of “To configure WebLogic 8.1.x” on page 110</a> or</li> <li>– <a href="#">Step 4 of “To configure WebLogic 9.2.x” on page 112</a></li> </ul> </li> <li>• Administrator Password: The password (at least eight characters) for the administrator login name.</li> </ul> <p>If WebLogic is not already running, Content Engine Setup tries to start it (this will take a few minutes) and displays a message saying so. Click <b>OK</b>.</p> <p>If WebLogic won't start, Content Engine Setup prompts you to run the tool startWebLogic:</p> <ul style="list-style-type: none"> <li>• (Windows) startWebLogic.cmd</li> <li>• (UNIX) startWebLogic.sh</li> </ul> <p>Continue at <a href="#">“GCD JNDI Configuration” on page 169</a>.</p>
Specify WebSphere Directory	<p>Specify the location where WebSphere Application Server is installed and then click <b>Next</b>.</p> <p><b>NOTE</b> When installing into a managed Server environment using Network Deployment or Deployment Manager, specify the installation location of WebSphere Network Deployment or Deployment Manager, respectively.</p>



In this screen...	Perform this action...
Specify WebSphere Information	<p>Specify the following information and then click <b>Next</b>.</p> <ul style="list-style-type: none"> <li>• Profile: The name of the WebSphere profile you created for this instance of Content Engine.</li> <li>• Cell: The name of WebSphere (multi-machine) environment.</li> <li>• Node: The name of the machine on which WebSphere is managed.</li> </ul> <p><b>NOTE</b> If you are installing Content Engine in a managed server configuration using Network Deployment or Deployment Manager, specify the host name of the Deployment Manager node.</p> <ul style="list-style-type: none"> <li>• Server: The name of the WebSphere instance.</li> <li>• SOAP Port: SOAP port assignment.</li> <li>• HTTP Port: HTTP port assignment.</li> </ul>
Specify WebSphere Administrator Account	<p>If WebSphere Global Security (Administrative Security in version 6.1) is turned on, type the name and password of the WebSphere administrator. If it is not turned on, type any name. Click <b>Next</b> and continue at <a href="#">“GCD JNDI Configuration” on page 169</a>.</p>
Jboss Configuration	<p>Specify the directory where JBoss is installed and then click <b>Next</b>.</p> <p>Specify the instance in the JBoss Server box,                      Specify 8080 the HTTP port 8080.</p>
GCD JNDI Configuration	<p>Click either the option button for Content Engine Setup to create new XA and non-XA JDBC Providers (WebSphere only) and/or XA and non-XA data sources/connection pools, or the option button for Content Engine Setup to use existing ones, and then click <b>Next</b>.</p> <p>If you choose to use existing ones, be sure to specify the exact, case-sensitive JNDI names for the data sources. If you wish to create new ones, specify the JNDI names you want Content Engine Setup to create.</p>
Configure JDBC	<p>From the drop-down lists, choose the database type and (for WebSphere and WebLogic only) the corresponding JDBC driver, and then click <b>Next</b>.</p>

In this screen...	Perform this action...
<p>Configure JDBC Connection Pools</p>	<p><b>CAUTION</b> Your database/tablespace for the GCD must already exist. If you haven't created it yet, exit from Content Engine Setup now and continue at one of the following, depending on your database vendor:</p> <ul style="list-style-type: none"> <li>• <a href="#">“Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81</a></li> <li>• <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85</a></li> <li>• <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a></li> </ul> <p>Specify the following information and then click <b>Next</b>.</p> <ul style="list-style-type: none"> <li>• <b>JDBC Connection Pool Name:</b> The name of the connection pool used for local (non-distributed) connections to the GCD database. Content Engine Setup will create an additional connection pool, based on this name, for global (distributed) connections.</li> <li>• <b>Database Name:</b> The name of the database to be used by the GCD. (For Oracle, the database name must be the SERVICE_NAME value in the <b>listener.ora</b> file).</li> <li>• <b>Database Host Name:</b> The name or IP address of the machine where the database is located.</li> <li>• <b>Database Port:</b> The port used by the database/tablespace.</li> </ul>
<p>Specify Database User</p>	<p>Specify the user name (ce_login, by default) and password that Content Engine will use to access the GCD database, and then click <b>Next</b>.</p> <p>If you did not choose Application Server Authentication Provider in <a href="#">“Select Components” on page 164</a>, skip to <a href="#">“Setup Content Engine Bootstrap Properties” on page 174</a>.</p>
<p>Specify Authentication Provider</p> <p><b>NOTE</b> This screen appears only if you selected to install the Application Server Authentication Provider component.</p>	<p>From the drop-down list, choose an authentication provider. If your application server is WebLogic, also specify a provider name. Select or clear the Kerberos checkbox as needed, and then click <b>Next</b>.</p> <p><b>NOTE</b> (WebLogic) Do not specify a provider name already being used on the WebLogic instance; otherwise, Content Engine Setup will overwrite that authentication provider.</p> <p><b>CAUTION</b> If you are running Content Engine Setup as part of an upgrade to version 4.0.x, your authentication provider must be the same one you used in version 3.5.x.</p>

In this screen...	Perform this action...
<p>Configure Authentication Provider</p> <p><b>NOTE</b> This screen appears only if you selected to install the Application Server Authentication Provider component.</p>	<p>Specify the following authentication provider information:</p> <ul style="list-style-type: none"> <li>• Host: The machine hosting the authentication provider.</li> <li>• Port: The port on which the authentication provider listens. This port must be non-SSL while Content Engine Setup is running. After Content Engine Setup finishes, you can set up SSL communication (see <a href="#">“Set Up Content Engine and Client Transport SSL Security” on page 412</a>).</li> <li>• Directory Service User: The distinguished name of the user that Content Engine will use to connect to the directory server. (This is the user you will specify in <a href="#">“To configure directory service authentication” on page 187</a>.)</li> </ul> <p>If your authentication provider is Windows Active Directory, an example distinguished name is:</p> <pre>cn=administrator,cn=users,dc=redwood,dc=local</pre> <ul style="list-style-type: none"> <li>• Password: The password for this user account.</li> </ul> <p>Click <b>Next</b>.</p>

In this screen...	Perform this action...
<p>Specify Authentication Provider User Account Parameters</p> <p><b>NOTE</b> This screen appears only if you selected to install the Application Server Authentication Provider component.</p>	<p><b>NOTE</b> For details on these parameters, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Enterprise-wide Administration &gt; FileNet P8 Security &gt; Directory Service Providers</a>, and navigate to the section about your directory service provider.</p> <p>Specify the following information, which will be used to by the authentication provider to authenticate users, and click <b>Next</b>.</p> <ul style="list-style-type: none"> <li>• User Base DN                     <p>If your authentication provider is Windows Active Directory, an example distinguished name is:</p> <pre>cn=users,dc=redwood,dc=local</pre> </li> <li>• User Name Attribute</li> <li>• User From Name Filter</li> </ul> <p><b>Weblogic and JBoss</b> By entering a User Base DN here you will create an authentication realm that will join any that already exist. You can add additional realms when Content Engine Setup is finished.</p> <p><b>WebSphere</b> Websphere 6.0.x and earlier supports only a single authentication realm, so entering account parameters here establishes that realm. However, later on when you run the Directory Configuration Wizard, you can run it multiple times to establish multiple authorization Base DNs as long as they are all included within the single authentication Base DN provided here. WebSphere 6.1 supports federated user repositories, explained in <a href="#">“To use WebSphere 6.1 federated user repository” on page 108</a>.</p> <p><b>All application servers</b> To configure multiple realms following installation, refer to IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Enterprise-wide Administration &gt; FileNet P8 Security &gt; How to &gt; Configure for multiple realms</a>.</p> <p>Click <b>Next</b>.</p>

In this screen...	Perform this action...
<p>Specify Authentication Provider Group Account Parameters</p> <p><b>NOTE</b> This screen appears only if you selected to install the Application Server Authentication Provider component.</p>	<p><b>NOTE</b> For details on these parameters, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Enterprise-wide Administration &gt; FileNet P8 Security &gt; Directory Service Providers</a>, and navigate to the section about your directory service provider.</p> <p>Specify the following information, which will be used by the authentication provider to authenticate groups, and then click <b>Next</b>:</p> <ul style="list-style-type: none"> <li>• Group Base DN                     <p>If your authentication provider is Windows Active Directory, an example distinguished name is</p> <pre>cn=users,dc=redwood,dc=local</pre> </li> <li>• Group From Name Filter</li> <li>• Static Group Name Attribute</li> </ul> <p>Depending on your application server, continue at one of the following:</p> <ul style="list-style-type: none"> <li>• (WebSphere) <a href="#">“Specify WebSphere Administrative Login” on page 173</a></li> <li>• (WebLogic or JBoss) <a href="#">“Setup Content Engine Bootstrap Properties” on page 174</a>.</li> </ul>
<p>Specify WebSphere Administrative Login</p>	<p>Type the user name of an account in your authentication provider to serve as a WebSphere administrator, and then click <b>Next</b>.</p>

In this screen...	Perform this action...
Setup Content Engine Bootstrap Properties	<p>Specify the administrator login name and password, needed to create a FileNet P8 domain and then click <b>Next</b>:</p> <p><b>CAUTION</b> Content Engine will fail to start if you specify the wrong login name or password, or if these values are altered on the directory server after Content Engine is installed.</p> <p>To recover from such a condition, refer to the P8 Help topic <a href="#">FileNet P8 Administration &gt; Enterprise-wide Administration &gt; FileNet P8 Security &gt; Security tools and procedures &gt; CE Bootstrap properties</a>.</p> <ul style="list-style-type: none"> <li>Administrator Login Name: The account (referred to in documentation as the “Content Engine system user”) under which Content Engine will run. It must be a directory service account. For WebSphere and WebLogic, it will also serve as an application server console administrator. You will use this account to create the IBM FileNet P8 domain later on.</li> <li>Administrator Password: The password of the Administrator Login account.</li> </ul>
Specify GCD Master Key	<p>Specify and confirm a word or phrase for encrypting sensitive GCD entries, and then click <b>Next</b>.</p> <p><b>NOTE</b> Store the master key in a secure location, as it is not retrievable. You will need it to install updated EAR files in interim fixes, or to access the GCD from applications built with non-IBM FileNet APIs.</p>
Compatibility COM API (CCL) Configuration  <b>NOTE</b> This screen appears only if you chose to install Client Connectivity (“ <a href="#">Select Components</a> ” on <a href="#">page 164</a> ).	<p>Accept the default values or enter your own values for the following parameters to allow .NET clients to connect to Content Engine, and then click <b>Next</b>:</p> <ul style="list-style-type: none"> <li>Connection: The protocol (HTTP or HTTPS) for communicating with Content Engine</li> <li>Server: The name of the machine where you are installing Content Engine</li> <li>Application Server: The type of application server on which Content Engine is to be deployed</li> <li>Path</li> <li>URL</li> </ul>

In this screen...	Perform this action...
Verify Installation	<p>Review the summary information about the Content Engine installation and click <b>Finish</b> to exit Content Engine Setup.</p> <p>If you have run Content Engine Setup as part of an upgrade, apply the latest Content Engine service pack before using Enterprise Manager to configure the IBM FileNet P8 domain.</p>

7. If you ran Content Engine Setup to create a response file for a silent install of Content Engine, continue at [“To install and deploy Content Engine silently” on page 176](#).
8. (WebLogic only) Stop and start the application server when Content Engine Setup exits, and then continue at [“Install Content Engine Software Updates” on page 177](#).
9. (WebSphere 5.1 and 6.0 only) Enable global security.
10. (WebSphere 6.1 only) Enable administrative security and application security, and disable Java 2 security.
11. (WebSphere only) Stop and start the application server. If the application server fails to start, or if you want to verify the LDAP user registry settings:
  - a. Navigate to the directory containing the wsadmin (Windows) or wsadmin.sh (UNIX) command on your WebSphere installation path:
    - (WebSphere 5.1.x) `<Install_path>/bin`
    - (WebSphere 6.0.x) `<Install_path>/profiles/<profile>/bin`
    - (WebSphere 6.1.x) `<Install_path>/profiles/bin`
  - b. Run the following wsadmin command, as follows:
    - (UNIX) `wsadmin.sh -conntype NONE`
    - (Windows) `wsadmin -conntype NONE`
  - c. At the `wsadmin>` prompt, type:
 

```
wsadmin> securityoff
wsadmin> exit
```
  - d. Restart the application server, verify global security LDAP user registry modifications and repeat [Step 9](#) or [Step 10](#).
12. (JBoss only) Verify the following:
  - The **FileNet/ContentEngine** directory exists and contains a nonempty file named **servers.xml**
  - The **FileNet/ContentEngine/lib** directory contains the file **Engine-jb.ear**.
13. Continue at [“Install Content Engine Software Updates” on page 177](#).

**To install and deploy Content Engine silently**

**NOTE** Before doing a silent install of Content Engine, you must have recorded a response file in “To run Content Engine Setup” on page 161.

Before you run Content Engine Setup, you can manually edit the response file to suit your installation environment. See “Encrypt Passwords for Silent Installations and Upgrades” on page 642 for a description of this file. Once you have completed your edits, do the following steps:

1. On the machine where you want to install Content Engine, log on as a user satisfying the requirements specified in Step 1 on page 161.
2. (WebLogic and WebSphere only) Start the application server on the Content Engine machine if it is not already running.

**NOTE** On UNIX, you can start WebSphere Application Server as the root user, or as a normal user with read/write/execute permissions on *<WebSphere\_Install\_Directory>/profiles/<profile\_name>*.

3. (WebSphere on Windows only) If Windows Active Directory is going to be your directory service, set the primary DNS server IP address on your Content Engine machine to the IP address of the machine where DNS is installed.
4. Access the Content Engine software package, and navigate to the **ContentEngine** directory.

Run one of the commands shown in the following table to silently install Content Engine (replacing */opt/response.txt* or *C:\response.txt* with the path to your response file):

Platform	Command line
AIX	P8CE-4.0.0-AIX.bin -options '/opt/response.txt' -silent
HP-UX 11	P8CE-4.0.0-HPUX.bin -options '/opt/response.txt' -silent
HP-UX 11i	On WebLogic 9.2 (replacing <i>/opt/java 1.5</i> with the path name to your version 1.5 JVM): P8CE-4.0.0-HPUXi.bin -is:javahome /opt/java1.5 -options '/opt/response.txt' -silent On non-WebLogic 9.2: P8CE-4.0.0-HPUXi.bin -options '/opt/response.txt' -silent
Linux	P8CE-4.0.0-Linux.bin -options '/opt/response.txt' -silent
Solaris (SPARC)	P8CE-4.0.0-Sol.bin -options '/opt/response.txt' -silent
Solaris (x86)	P8CE-4.0.0-Solx86.bin -options '/opt/response.txt' -silent
Windows	P8CE-4.0.0-Win.exe -options 'C:\response.txt' -silent

5. Continue at Step 8 on page 175.



## Task 10: Install Content Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Content Engine.

### To install the Content Engine software updates

---

1. To download the latest software update, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. P8 CE 4.0.1 Service Pack
  - b. Any subsequent interim fixes (typically optional)

# Task 11: Complete Post-Install Content Engine Configuration

## Install Centera Shared Libraries

If you are going to access Centera fixed content devices with your IBM FileNet P8 system, you must install the Centera shared libraries on each machine where you deploy Content Engine Server. The following two procedures install the libraries on UNIX and Windows machines:

- [“To install Centera shared libraries \(UNIX\)” on page 178](#)
- [“To install Centera shared libraries \(Windows\)” on page 179](#)

Do the procedure corresponding to the operating system that hosts Content Engine Server.

### To install Centera shared libraries (UNIX)

---

**NOTE** The path names and environment variables in this procedure are appropriate for the AIX operating system and may differ for other UNIX systems.

1. If you haven't already done so, log on as a user with write permission on the directory in which you will install the Centera shared libraries, for example **/usr/local**.
2. Access the Content Engine installation software Insert and mount the Content Engine installation CD in the CD drive.
3. Navigate to the **Centera/install** directory and copy this directory to a temporary location (where you have execute permission), from which you will install the Centera libraries.

Note that the **Centera** directory contains **install** and **lib** subdirectories. The **install** directory contains the installation and setup scripts (**install.sh** and **setCenteraLibPath.sh**); the **lib** directory contains the Centera shared libraries.

4. Navigate to the **Centera/install** directory in the temporary location and do the following substeps to install the Centera shared libraries:
  - a. Launch the installation script **install.sh**:

```
./install.sh
```

- b. The script will prompt you to specify an installation directory, as in the following example:

```
Introduce install directory [/usr/local/Centera_SDK]:
```

Accept the default installation directory, **/usr/local/Centera\_SDK**, or specify another one, and then press <Enter>. You do not need to create the installation directory in advance; if the directory doesn't exist, the script will create it.

If your UNIX system is Solaris or HP-UX 11i, **install.sh** creates subdirectories **lib/32** and **lib/64** under the installation directory you specify, and installs both the 32-bit and 64-bit libraries. On other UNIX systems, **install.sh** creates only subdirectory **lib/32**.

- Navigate to **Centera/install** directory at the temporary location, open the file **setCenteraLibPath.sh** in a text editor, and verify that it contains the following lines:

```
CENTERA_LIB_PATH=/usr/local/Centera_SDK/lib/32
LIBPATH=$LIBPATH:$CENTERA_LIB_PATH
export LIBPATH
```

**NOTE** On HP-UX 11i, **SHLIB\_PATH** appears in place of **LIBPATH**.

- Edit **setCenteraLibPath.sh** and save your changes, as follows:
  - If you did not accept the default installation directory in [Step 4](#), replace `/usr/local/Centera_SDK` with the value you did specify.
  - If your UNIX system is Solaris or HP-UXZ 11i, the default value for the first line in [Step 5](#) is as follows:

```
CENTERA_LIB_PATH=/opt/Centera_SDK/lib/32
```

If your Solaris or HP-UX 11i system is 64-bit, change the line above to this:

```
CENTERA_LIB_PATH=/opt/Centera_SDK/lib/64
```

- Open your application server startup script in a text editor:
  - (WebSphere) **setupCmdLine.sh**
  - (WebLogic 8.1.x) **startWebLogic.sh**
  - (WebLogic 9.2) **setDomainEnv.sh**
  - (JBoss) **run.sh**
- Add a line to the application server startup script to run the **setCenteraLibPath.sh** script and save your edits. Continuing the example from [Step 4](#):
 

```
./usr/local/Centera_SDK/setCenteraLibPath.sh
```
- Stop and start the application server to run the startup script.
- If you did this procedure as part of your first deployment of Content Engine Server, continue at ["To verify the GCD data sources and connection pools" on page 180](#). If you did this procedure as part of deploying Content Engine to an additional server, exit this task.

### To install Centera shared libraries (Windows)

---

- If you haven't already done so, log on as a user with write permission on the drive where you will install the Centera shared libraries, for example **C:**.
- Access the Content Engine installation software.
- Navigate to the **Centera** directory and copy this directory to a temporary location, from which you will install the Centera libraries.

Note that the **Centera** directory contains **install** and **lib** subdirectories. The **install** directory contains the installation and setup scripts; the **lib** directory contains the Centera shared libraries.

4. Navigate to the **Centera\install** directory at the temporary location and run the installation script **install.bat** to install the Centera shared libraries, as follows:

```
install.bat INSTALL_DIR
```

where *INSTALL\_DIR* is the installation directory where you want to install the Centera shared libraries, for example, C:\Centera\_SDK.

You do not need to create the installation directory in advance; if the directory doesn't exist, the installation script will create it.

5. Navigate to the **Centera\install** directory at the temporary location, open the file **setCenteraLibPath.bat** in a text editor, and verify that it contains the following lines:

```
set CENTERA_LIB_PATH=C:\Centera_SDK\lib  
set PATH=%PATH%;%CENTERA_LIB_PATH%
```

If you did not specify C:\Centera\_SDK as the installation directory in [Step 4](#), change C:\Centera\_SDK in the first line above to the location you did specify.

6. Open your application server startup script in a text editor:
  - (WebSphere) **setupCmdLine.cmd**
  - (WebLogic 8.1.x) **startWebLogic.cmd**
  - (WebLogic 9.2) **setDomainEnv.cmd**
  - (JBoss) **run.bat**
7. Add the following line to your application server startup script to run the **setCenteraLibPath.bat** script and save your edits. Continuing the example:

```
set CENTERA_LIB_PATH=C:\Centera_SDK\lib  
set PATH=%PATH%;%CENTERA_LIB_PATH%
```

8. Stop and start the application server to run the startup script.

## Verify the Data Sources and Connection Pools

### To verify the GCD data sources and connection pools

---

1. If WebLogic is your application server, do the following:
  - a. If you ran Content Engine Setup on a machine where a browser is installed, WebLogic Server Administration Console automatically starts after you click **Finish** in [“Verify Installation” on page 175](#); otherwise, on a machine where a browser is installed, start WebLogic Server Administration Console by going to the following web address:

```
http://<WebLogic_Machine_Name>:7001/console
```

where *WebLogic\_Machine\_Name* is the name of the WebLogic machine on which you ran Content Engine Setup.

- b. Log on to WebLogic Server Administration Console as the user you specified in [“Specify WebLogic Information” on page 168](#), and verify that the data sources and connection pools you specified in Content Engine Setup have been created.

- c. Log out of WebLogic Server Administration Console.
- d. Stop and start WebLogic Application Server.

**NOTE** If there are too few connection pools for a GCD data source, an IBM FileNet P8 client application will encounter errors. Depending on the number of object stores you will create, and the number of users who will concurrently access the GCD, you may want to modify, via WebLogic Administration Console, the number of connection pools (XA and non-XA) for each GCD data source. As a general rule, this number should be  $\langle \text{number of object stores} \rangle * 16 + \langle \text{number of concurrent users} \rangle$ .

2. If WebSphere is your application server, do the following:
  - a. Stop and start WebSphere.
  - b. Log on to WebSphere administrative console as the user you specified in [“Specify WebSphere Administrative Login”](#) on page 173.
  - c. Verify that the data sources and connection pools you created are present.
  - d. Navigate to Applications > Enterprise Applications and verify that FileNetEngine is in the Enterprise Applications list and has a right-pointing green arrow in the Application Status column.
3. If JBoss is your application server, do the following:
  - a. Verify that the data sources and connection pools you created are present.
  - b. Verify that Content Engine is deployed.
4. Continue at [“To verify that Content Engine has deployed”](#) on page 181.

## Verify That Content Engine Has Deployed

### To verify that Content Engine has deployed

---

Do the following steps to verify that Content Engine is deployed on your application server:

1. Open a browser on the Content Engine machine and navigate to the following location:

`http://localhost:<port>/FileNet/Engine`

where the default value of `<port>` is one of the following:

- (WebSphere) 9080
- (WebLogic) 7001
- (JBoss) 8080

2. Verify that this location displays Content Engine information similar to the following:

Startup Message	P8 CEMP Startup: 4.0 dap430.1089 Copyright (c) 2006, 2007 IBM Corporation. All rights reserved.
-----------------	---

3. Start version 4.0.x of Content Engine Server if it is not already running.

**NOTE** An upgrade of Content Engine to version 4.0.x will fail if you do not start Content Engine Server.

4. If you have not previously selected Administration Tools - Enterprise Manager for installation in [“Select Components” on page 166](#), continue at [“Install Enterprise Manager” on page 183](#). Otherwise, continue at [“To create a FileNet P8 domain” on page 187](#).

## ***Enable Log4J Logging API (Optional)***

### **To enable the Log4j logging API (optional)**

---

To allow Log4j logging by IBM FileNet P8 client applications, the **log4j.properties.client** file is in the **FileNet/ContentEngine/config/samples** directory. Before you can use Log4J logging in an IBM FileNet P8 client application, you need to move and rename this file as follows:

1. Navigate to the directory **FileNet/ContentEngine/config/samples**, which contains the log4j files.
2. Copy the file to the **FileNet/ContentEngine/config** directory.
3. Rename the file you moved, **log4j.properties.client**, to **log4j.properties**.
4. Depending on whether you ran Content Engine Setup as an upgrade to an earlier version of Content Engine, or as a new installation, continue as follows:
  - If you are upgrading, continue at procedure [“To prepare NetApp SnapLock Volumes for the upgrade” on page 493](#) in [“Upgrade Content Engine Software” on page 487](#).
  - If you are running a new installation, continue at [“Install Enterprise Manager” on page 183](#).

## Task 12: Install Enterprise Manager

By installing Enterprise Manager on a single client machine, you can administer multiple object stores in a single IBM FileNet P8 domain.

### NOTES

- You can install Enterprise Manager only on a Windows machine, and only using the Windows version of the Content Engine installation media.
- If you installed Content Engine Setup on a Windows machine, you can skip [“Establish the FileNet P8 Domain and Global Configuration Data \(GCD\)” on page 187](#), as Content Engine Setup has already installed Enterprise Manager.

On a machine where you are going to install Enterprise Manager, you must first install Microsoft .NET Framework and Web Services Enhancements (WSE). Check the *IBM FileNet P8 Hardware and Software Requirements* for the latest version requirements of these two components. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#). Enterprise Manager requires no other Content Engine services or files.

You can install Enterprise Manager interactively, via Content Engine Setup, or silently, via a command-line interface.

Before doing a silent install of Enterprise Manager, you must first record a response file by running Content Engine Setup interactively, typically in a development or test setting. You then use the response file to run a silent install of Enterprise Manager in a production environment.

**CAUTION** Do not install Enterprise Manager 4.0.x on any machine running the 3.5.x version, at least until the Content Engine 4.0.x upgrade is complete. Otherwise, you will no longer be able to run Enterprise Manager 3.5.x against any remaining 3.5.x object stores.

### To install Enterprise Manager interactively or to record a response file

---

1. Log on to the Windows machine with an account that has local administrator privileges.
2. Access the Content Engine software package, and navigate to the directory `\ContentEngine\CEMP.Windows`.
3. Start Content Engine Setup, as follows.

**NOTE** If you are installing on a server where you have already applied a Content Engine 4.0 software update (Service Pack or fix pack), when you issue setup commands you must include two additional `-W` switches and associated parameters to disable version checking, as noted below.

- If you are doing an interactive (non-silent) install, run this command:

(On a server having no Content Engine 4.0 software updates applied)

```
P8CE-4.0.0-Win.exe
```

(On a server having Content Engine 4.0 software updates applied)

```
P8CE-4.0.0-Win.exe -W updateCheck.active=false -W installCheck.active=false
```

- If you are recording a response file for a subsequent silent install, run the following command (replacing C:\response.txt with the path to your response file):

(On a server having no Content Engine 4.0 software updates applied)

```
P8CE-4.0.0-Win.exe -options-record "C:\response.txt"
```

(On a server having Content Engine 4.0 software updates applied)

```
P8CE-4.0.0-Win.exe -options-record "C:\response.txt" -W updateCheck.active=false  
-W installCheck.active=false
```

The first Content Engine Setup screen may not appear for 30 seconds or more after you run the command to start it.

4. Start the Content Engine Setup program using one of the following commands, depending on whether you want the program to record a response file and whether you have applied a service pack or interim fix to the Content Engine software on this machine.

**NOTE** The first Content Engine Setup screen may not appear for 30 seconds or more after you run the command to start it.

- If you have not yet applied a service pack or interim fix and want a response file:

```
P8CE-4.0.0-Win.exe -options-record "C:\response.txt"
```

- If you have already applied a service pack or interim fix and want a response file:

```
P8CE-4.0.0-Win.exe -options-record "C:\response.txt" -W updateCheck.active=false  
-W installCheck.active=false
```

- If you have not applied a service pack or interim fix and do not want a response file:

```
P8CE-4.0.0-Win.exe
```

- If you have already applied a service pack or interim fix and do not want a response file:

```
P8CE-4.0.0-Win.exe -options -W updateCheck.active=false -W  
installCheck.active=false
```

5. Complete the Content Engine Setup wizard screens as follows:

In this screen...	Perform this action...
Welcome	Click <b>Next</b> to proceed with the Content Engine installation.  <b>NOTE</b> Click <b>Back</b> at any time while running Content Engine Setup to make changes in any previous screens. Click <b>Cancel</b> to exit Content Engine Setup.
FileNet Notice to End User	Review and accept the license agreement for IBM FileNet P8 software products, and then click <b>Next</b> .
Specify Installation Location	Accept the default (or browse to another) Content Engine installation directory, and then click <b>Next</b> .
Select Components	From the list shown on the screen, select the Administrative Tools - Enterprise Manager check box, clear all other check boxes, and then click <b>Next</b> .



In this screen...	Perform this action...
.NET Framework and WSE Requirements	<p>This screen appears only if Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0 are not both installed on the machine where you are running Content Engine Setup.</p> <p>Select the Continue? check box if you have later or compatible versions of these components, and then click <b>Next</b>. Otherwise, click <b>Cancel</b> to exit Content Engine Setup and install these two components.</p>
Specify Documentation URL	<p>If you have already installed IBM FileNet P8 Documentation, specify its URL and click <b>Test</b> to verify the connection. This setting enables access from Enterprise Manager and other IBM FileNet P8 tools and interfaces.</p> <p>Alternatively, you can specify the URL after Content Engine Setup finishes, using the General tab of the Root node properties of Enterprise Manager.</p> <p>Click <b>Next</b>.</p>
Review Installation Summary	<p>Verify your component selections and click <b>Install</b> to start installing Content Engine.</p>
Verify Installation	<p>Review the summary information about the installation and click <b>Finish</b> to exit Content Engine Setup.</p>

When Content Engine Setup exits, do one of the following:

- If you ran Content Engine Setup to create a response file for a silent install of Enterprise Manager, continue at [“To install Enterprise Manager silently” on page 185](#).
- If you ran Content Engine Setup interactively (that is, without creating a response file), do one of the following:
  - If this is the first instance of Enterprise Manager you have installed, continue at [“Establish the FileNet P8 Domain and Global Configuration Data \(GCD\)” on page 187](#).
  - If this is an additional instance of Enterprise Manager, exit this task.

### To install Enterprise Manager silently

**NOTE** Before doing a silent install of Enterprise Manager, you must have recorded a response file in [“To install Enterprise Manager interactively or to record a response file” on page 183](#).

You can manually edit the response file to suit your installation environment before you install Enterprise Manager (see [“Encrypt Passwords for Silent Installations and Upgrades” on page 642](#) for a description of this file). Once you have completed your edits, do the following steps:

1. Log on to the Windows machine where you want to install Enterprise Manager as a member of the local Administrators group.

2. Access the Content Engine software package, and navigate to the **ContentEngine** directory.
3. Silently install Enterprise Manager (replacing **C:\response.txt** with the path to your response file) using one of the following commands:

- If you have not yet applied a service pack or interim fix to the Content Engine software on this machine, run this command:

```
P8CE-4.0.0-Win.exe -options "C:\response.txt" -silent
```

- If you are installing an instance of Enterprise Manager on a machine after having applied a service pack or interim fix to the Content Engine software on that machine, run the following command (without carriage returns) with two **-W** parameters to disable version checking, as follows:

```
P8CE-4.0.0-Win.exe -options "C:\response.txt" -silent -W updateCheck.active=false  
-W installCheck.active=false
```

4. When Content Engine Setup exits, do one of the following:
  - If this is the first instance of Enterprise Manager you have installed, continue at [“Establish the FileNet P8 Domain and Global Configuration Data \(GCD\)” on page 187](#).
  - If this is an additional instance of Enterprise Manager, continue at [“Install an Additional Instance of Enterprise Manager” on page 427](#).

## Task 13: Establish the FileNet P8 Domain and Global Configuration Data (GCD)

### To create a FileNet P8 domain

---

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

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn 4.0 on the desktop, or by choosing Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn 4.0.
2. In the FileNet P8 Logon dialog box, click **Add**, to create a FileNet P8 domain configuration.  
**NOTE** For subsequent logons to Enterprise Manager, you can access an existing FileNet P8 domain by clicking **Connect**.
3. In the Add Domain Configuration dialog box, specify the following information and then click **OK**:
  - Nickname - an identifier (not part of any credentials) that can be used to connect to Content Engine
  - Username - the name of the Content Engine system user, which must be the account you specified earlier in the Setup Content Engine Bootstrap Properties screen of Content Engine Setup.
  - Remember password - Select the check box if you want to avoid typing the password each time you log on as this user. (The password will be encrypted.)
  - Use integrated - Select the check box if you want Kerberos authentication for this user
4. In the FileNet P8 Logon dialog box, click **Connect**.
5. In the Create P8 Domain screen, enter the name for a new FileNet P8 domain and click **Continue** to launch the Create a Directory Configuration wizard.
6. In the FileNet P8 - Domain Logon dialog box, type the user name and password and click **OK**.

Continue at ["To configure directory service authentication" on page 187](#).

### To configure directory service authentication

---

**NOTE** The configuration parameters required by the Create a Directory Configuration wizard are in many cases the same as those you provided to Content Engine Setup while configuring the application server's authentication provider. Refer also to the topic for your directory server within the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Directory service providers](#).

1. In the Welcome screen of the Create a Directory Configuration wizard, click **Next**.
2. In the Select Type and Name Directory Configuration screen, as follows:
  - a. Choose the same directory service from the Type drop-down list that you selected earlier while running Content Engine Setup.

- b. Type a display name (unique across all FileNet P8 domains in a forest) for the new directory configuration.
- c. Click **OK**.

If you chose Windows Active Directory, continue at [Step 3](#); if you chose any of the following directory services, continue at [Step 6](#).

- ADAM
- IBM Tivoli Directory Server
- Novell eDirectory
- Sun Java System Directory Server

3. In the Select General Directory Configuration Properties screen, specify the following information and then click **Next**:

Parameter	Value and Description
Host	The name or IP address of the host where the Windows domain controller is installed. In a forest environment, the Windows domain names must be distinct if name you specify must be unique within each.
Port	The LDAP port number (389, by default).
Directory Service User	Distinguished name of Content Engine user who will access the directory service provider. This is the same user you specified when running Content Engine Setup (see <a href="#">“Configure Authentication Provider” on page 171</a> ).
Password	Password of the directory service user.
Is SSL Enabled	True or False.
Return Name as DN	Choose whether to return names from the directory service provider in distinguished-name format.

4. In the Select User Directory Configuration Properties screen, specify the following information and then click **Next**:

Parameter	Value and Description
User Base DN	Base distinguished string to use in LDAP user searches.
User Search Filter	LDAP search filter for finding user names.
User Display Name Attribute	cn, by default.
User Short Name Attribute	samAccountName, by default.

5. In the Select Group Directory Configuration Properties screen, specify the following information, click **Next**, and continue at [Step 9](#):

Parameter	Value and Description
Group Base DN	Base distinguished string to use in LDAP group searches.
Group Search Filter	LDAP search filter for finding group names.
Group Display Name Attribute	cn, by default.
Group Short Name Attribute	cn, by default.
Search Cross Forest Group Membership	False, by default.

6. In the Select General Directory Configuration Properties screen, specify the following information and then click **Next**:

Parameter	Value and Description
Host	Name or IP address of machine where the directory server is installed.
Port	LDAP port number (389, by default).
Directory Service User	Distinguished name of Content Engine user who will access the directory service provider.
Password	Password of directory service user.
Is SSL Enabled	True or False.

7. In the Select User Directory Configuration Properties screen, specify the following information and then click **Next**:

Parameter	Value and Description
User Base DN	Base distinguished string to use in LDAP user searches.
User Search Filter	LDAP search filter for finding user names.

Parameter	Value and Description
User Display Name Attribute	<ul style="list-style-type: none"> <li>• cn, by default, for the following:                             <ul style="list-style-type: none"> <li>– ADAM</li> <li>– IBM Tivoli Directory Server</li> <li>– Novell eDirectory</li> </ul> </li> <li>• uid, for Sun Java System Directory Server</li> </ul>
User Short Name Attribute	<ul style="list-style-type: none"> <li>• cn, by default, for the following:                             <ul style="list-style-type: none"> <li>– ADAM</li> <li>– IBM Tivoli Directory Server</li> <li>– Novell eDirectory</li> </ul> </li> <li>• uid, for Sun Java System Directory Server</li> </ul>

8. In the Select Group Directory Configuration Properties screen, specify the following information, click **Next**, and continue at [Step 9](#):

Parameter	Value and Description
Group Base DN	Base distinguished string to use in LDAP group searches.
Group Search Filter	LDAP search filter for finding group names.
Group Display Name Attribute	cn, by default.
Group Short Name Attribute	cn, by default.
Group Membership Search Filter	The search filter for group membership queries.

9. In the Completing the Create a Directory Configuration Wizard screen, click **Finish**.
10. In the Configure New Domain Permissions message box, click **OK** to acknowledge that the directory configuration is complete but remains in restricted mode. The Configure New Domain Permissions wizard automatically starts. Continue at [“To configure permissions for a FileNet P8 domain” on page 191](#).

**NOTE** For multi-realm authorization, run the Directory Configuration Wizard once for each realm. Refer to [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms](#).

### To configure permissions for a FileNet P8 domain

---

1. In the Welcome screen of the Configure New Domain Permissions wizard, click **OK**.
2. In the Specify domain administrators screen,
  - a. Click **Add** to load the Select Users and Groups dialog box that lets you add users and groups to the list of GCD administrators.
  - b. (Optional) Click **Remove** to remove the Content Engine system user.
  - c. Click **Next**.
3. In the Completing the Configure New Domain Permissions Wizard screen, click **Finish**.
4. In the Configure New Domain Permissions message box, click **OK**.

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

### To set the Statement Cache Size value for the GCD database

---

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

1. Access the page containing the Statement Cache Size parameter:
  - (WebSphere) Navigate to the WebSphere administrative console page containing the field Statement Cache Size property. For example, in WebSphere 6.0.x, navigate to Resources > JDBC Providers > JDBC\_provider > Data sources > *data\_source* > WebSphere Application Server connection properties.
  - (WebLogic 8.1.x) Navigate in the tree view of WebLogic Administration Console to *FNCEDomain* > Services > JDBC > Connection Pools > *Connection\_Pool\_Name*.
  - (WebLogic 9.2.x) Navigate in the tree view of WebLogic Administration Console to *FNCEDomain* > Services > JDBC > Data Sources > *Data\_Source\_Name* > Connection Pool
2. Set the Statement Cache Size to 0 and save your change.
3. Depending on whether you ran Content Engine Setup as an upgrade to an earlier version of Content Engine, or as a new installation, continue as follows:
  - If you are upgrading, continue at procedure [“To prepare NetApp SnapLock Volumes for the upgrade” on page 493](#) in [“Upgrade Content Engine Software” on page 487](#).
  - If you are running a new installation, continue at one of the following:
    - [Task 14a “Configure Content Engine Application Server Database Connectivity \(WebSphere 5.1.x\)” on page 193](#)
    - [Task 14b “Configure Content Engine Application Server Database Connectivity \(WebSphere 6.0.x\)” on page 206](#)

- [Task 14c “Configure Content Engine Application Server Database Connectivity \(WebSphere 6.1.x\)” on page 220](#)
- [Task 14d “Configure Content Engine Application Server Database Connectivity \(WebLogic 8.1.x\)” on page 236](#)
- [Task 14e “Configure Content Engine Application Server Database Connectivity \(WebLogic 9.2.x\)” on page 241](#)
- [Task 14f “Configure Content Engine Application Server Database Connectivity \(JBoss 4.0.x\)” on page 247](#)



## Task 14a: Configure Content Engine Application Server Database Connectivity (WebSphere 5.1.x)

In this task, you will set up the data sources, connection pools, and (optionally) the distributed (XA) and non-distributed (non-XA) JDBC providers for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Perform one of the following procedures for each object store, depending on your database type:

- [“To configure WebSphere 5.1.x database connectivity \(DB2\)” on page 193](#)
- [“To configure WebSphere 5.1.x database connectivity \(MS SQL Server\)” on page 198](#)
- [“To configure WebSphere 5.1.x database connectivity \(Oracle\)” on page 201](#)

**NOTE** If you are upgrading from version 3.5.x of Content Engine, specify the 3.5.x database or tablespace information in these procedures.

### **To configure WebSphere 5.1.x database connectivity (DB2)**

---

1. Start the WebSphere administrative console (if it isn't already running).
2. For each object store, you need to create a database user alias associated with a unique database and database user, as follows:
  - a. Navigate to Security > JAAS Configuration > J2C Authentication Data. In the right-hand pane, notice that there is already an alias (created by Content Engine Setup) for the GCD tablespace user.
  - b. Click **New** to create an alias and specify the following information:
    - Alias
    - User (object store tablespace user)
    - Password (object store tablespace user password)
    - Description (optional)
  - c. Click **Apply**, and then click **Save**.
3. Navigate to Resources > JDBC Providers, click **Cell** and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both distributed (XA) and non-XA JDBC providers listed.

You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).
4. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC Providers, verify that the **Cell** option button is selected, and click **New**.
  - b. In the Configuration tab, set the value of the JDBC Providers property to *DB2 Universal JDBC Driver Provider*, and then click **Apply**.

- c. Specify the values shown in the following table:

Property	Value
Name	<User-defined non-XA JDBC provider name>
Description	(Optional)
Classpath	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/ db2jcc_license_cu.jar  \${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar
Native Library Path	(none)
Implementation Classname	com.ibm.db2.jcc.DB2ConnectionPoolDataSource

- d. Click **OK** to view the table that shows your new JDBC provider.

5. Create an XA JDBC provider as follows:

- Verify that the Cell option button is still selected, and then click **New**.
- In the Configuration tab, set the value of the JDBC Providers property to DB2 Universal JDBC Driver Provider (XA), and then click **Apply**
- Specify the values shown in the following table:

Property	Value
Name	XA JDBC provider name
Classpath	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/ db2jcc_license_cu.jar  \${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar
Native Library Path	(none)
Implementation Classname	com.ibm.db2.jcc.DB2XADataSource

- d. Click **OK** to view the table that shows your new JDBC provider.

- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

6. Create a non-XA data source for each object store, as follows:

- Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the DB2 non-XA JDBC provider that Content Engine Setup created or that you created in [Step 4](#).

- b. Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**.

Parameter	Value
Name	Data source name
JNDI Name	The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Datasource Helper Classname	DB2 Universal data store helper com.ibm.websphere.rsadapter.DB2UniversalDataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created in created in <a href="#">Step 2 on page 193</a>
Container-managed Authentication Alias	(none)
Mapping-configuration Alias	(none)

- c. Click **Apply**, click **Custom Properties** in the Additional Properties table, and specify the values shown in the following table:

Parameter	Value
databaseName	Name of DB2 database
driverType	4
serverName	Host name of machine where DB2 is installed
portNumber	50000 (by default)
resultSetHoldability	1
fullyMaterializeLobData	See Note below.

**NOTE** The default value of fullyMaterializeLobData is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location.
- Users will be storing large content elements (for example, larger than 300 MB).

Setting the value of fullyMaterializeLobData to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.

d. Click **New** and specify the the property values shown in the following table:

Name	Value	Type
webSphereDefaultIsolationLevel	2	java.lang.Integer

e. Click **Apply**.

f. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

7. Test the connection as follows:

a. Navigate to Resources > JDBC Providers and click the non-XA JDBC provider created by Content Engine Setup created or that you created in [Step 4](#)).

b. Under Additional Properties, click Data Sources, select the check box for the non-XA data source, and then click **Test Connection** to test the connection.

8. Create an XA data source for each object store, as follows:

a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the XA entry for the DB2 database type of JDBC provider.

b. Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name.
JNDI Name	The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Datastore Helper Classname	DB2 Universal data store helper com.ibm.websphere.rsadapter.DB2UniversalDataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created in created in <a href="#">Step 2 on page 193</a>
Container-managed Authentication Alias	(none)
Mapping-Configuration Alias	(none)

- c. Click **Apply**, click **Custom Properties** in the Additional Properties table, and specify the values shown in the following table:

Parameter	Value
databaseName	Name of DB2 database
driverType	4
serverName	Host name of machine where DB2 is installed
portNumber	50000 (by default)
fullyMaterializeLobData	See Note below.

**NOTE** The default value of `fullyMaterializeLobData` is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location.
- Users will be storing large content elements (for example, larger than 300 MB).

Setting the value of `fullyMaterializeLobData` to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.

- d. Click **New** and specify the the property values shown in the following table:

Name	Value	Type
webSphereDefaultIsolationLevel	2	java.lang.Integer

- e. Click **Apply**.
- f. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
9. Test the connection as follows:
- Navigate to Resources > JDBC Providers and click the XA JDBC provider created by Content Engine Setup created or that you created in step [Step 5](#)).
  - Under Additional Properties, click Data Sources, select the check box for the XA data source, and then click **Test Connection** to test the connection.
10. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebSphere 5.1.x database connectivity (MS SQL Server)**

---

1. Start the WebSphere administrative console (if it isn't already running).
2. Navigate to Security > JAAS Configuration > J2C Authentication Data. In the right-hand pane, note the alias created for the GCD database user, as you will use this alias for data source connectivity.
3. Navigate to Resources > JDBC Providers. Click **Cell** and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both XA and non-XA JDBC providers listed.

You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).

4. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC Providers, click **Cell**, and then click **New**.
  - b. In the Configuration tab, set the value of the JDBC Providers property to *WebSphere embedded ConnectJDBC driver for MS SQL Server*, and then click **Apply**.
  - c. Specify the values shown in the following table:

Property	Value
Name	<User-defined non-XA JDBC provider name>
Classpath	\${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar
Native Library Path	(none)
Implementation Classname	com.microsoft.sqlserver.jdbc.SQLServerConnectionPoolDataSource

- d. Click **OK** to view the table that shows your new JDBC provider.
5. Create an XA JDBC provider, as follows:
  - a. Verify that the Cell option button is still selected, and then click **New**.
  - b. In the Configuration tab, set the value of the JDBC Providers property to *WebSphere embedded ConnectJDBC driver for MS SQL Server (XA)*, and then click **Apply**.
  - c. Specify the values shown in the following table:

Property	Value
Name	<User-defined XA JDBC provider name>
Classpath	\${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar

Property	Value
Native Library Path	(none)
Implementation classname	com.microsoft.sqlserver.jdbc.SQLServerXADataSource

- d. Click **Apply** to view the page with the new JDBC provider settings.
  - e. Click **Save** to save your settings to the master configuration, and then click **Save** to update the master repository with your changes.
6. Create a non-XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, select the non-XA entry for the SQL Server database type of JDBC provider that Content Engine Setup created or that you created in [Step 4](#).
  - b. Under Additional Properties, click **Data Sources**, click **New**, specify the values in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name
JNDI Name	The name (at most 8 characters) to be used by the Create Object Store wizard
Statement Cache Size	0
Datasource Helper Classname	com.filenet.engine.util.DataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created by the CE installer
Container-managed Authentication Alias	(none)
Mapping-Configuration Alias	(none)

- c. Under Additional Properties, click **Connection Pool** and set the Min. Connections value to 1, and the Max. Connections value to 150, and click **OK**.

- d. Under Additional Properties, click **Custom Properties** and add the custom properties, shown in the following table. (To create a property, click **New**, specify the values, and then click **OK**):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
serverName	<Host_Name>	java.lang.String
portNumber	1433 (default)	java.lang.Integer
selectMethod	direct	java.lang.String
enable2Phase	false	java.lang.Boolean

- e. Click **Save** to apply changes to the master configuration and then click **Save** to update the master repository with your changes.
7. Create an XA data source for each object store, as follows:
- Navigate to Resources > JDBC Providers. In the table of JDBC providers, select the XA entry for the SQL Server database type of JDBC provider.
  - Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name
JNDI Name	The name (at most 8 characters) to be used by the Create Object Store wizard
Statement Cache Size	0
Datastore Helper Classname	com.filenet.engine.util.DataStoreHelper
Component-managed authentication Alias	<Node>/<alias> created by the CE installer
Container-managed Authentication Alias	(none)
Mapping-configuration Alias	(none)

- Under Additional Properties, click **Connection Pool** and set the Min. Connections value to 1, and the Max. Connections value to 150, and click **OK**.



- d. Under Additional Properties, click **Custom Properties** and add the custom properties, shown in the following table. (To create a property, click **New**, specify the values, and then click **OK**):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
serverName	<Host_Name>	java.lang.String
portNumber	1433 (default)	java.lang.Integer
selectMethod	direct	java.lang.String
enable2Phase	true	java.lang.Boolean

- e. Click **Save** to apply changes to the master configuration and then click **Save** to update the master repository with your changes.
8. Test the connection as follows:
- Navigate to Resources > JDBC Providers and click the XA JDBC provider created by Content Engine Setup created or that you created in [Step 5](#)).
  - Under Additional Properties, click Data Sources, select the check box for the XA data source, and then click **Test Connection** to test the connection.
9. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

### To configure WebSphere 5.1.x database connectivity (Oracle)

---

- Start the WebSphere administrative console (if it isn't already running).
- For each object store, you need to create a database user alias associated with a unique database and database user, as follows:
  - Navigate to Security > JAAS Configuration > J2C Authentication. In the right-hand pane, note there is already an alias created for the GCD tablespace user.
  - Click **New** to create another one for an object store with the following information:
    - Alias
    - User (object store tablespace user)
    - Password (object store tablespace user password)
    - Description (optional)
  - Click **Apply**, and then click **Save**.
- Navigate to Resources > JDBC Providers. Click **Cell** and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both XA and non-XA JDBC providers listed.

You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).

4. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC Providers, verify that the Cell option button is selected and click **New**.
  - b. In the Configuration tab, set the value of the JDBC Providers property *Oracle JDBC Driver*, and then click **Apply**.
  - c. Specify the values shown in the following table:

Property	Value
Name	Non-XA JDBC provider name
Classpath	\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Native Library Path	(none)
Implementation Classname	oracle.jdbc.pool.OracleConnectionPoolDataSource

- d. Click **OK** to view the table that shows your new JDBC provider.

5. Create an XA JDBC provider, as follows:
  - a. Verify that the Cell option button is still selected, and then click **New**.
  - b. In the Configuration tab, set the value of the JDBC Providers property to *OracleJDBC Driver (XA)*, and then click **Apply**.
  - c. Specify the values shown in the following table:

Property	Value
Name	XA JDBC provider name
Classpath	\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Native Library Path	(none)
Implementation Classname	oracle.jdbc.xa.client.OracleXADataSource

- d. Click **OK** to view the table that shows your new JDBC provider.
  - e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

6. Create a non-XA data source, as follows:

- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the Oracle non-XA JDBC provider that Content Engine Setup created or that you created in [Step 4](#).
- b. Under Additional Properties, click **Data Sources**, click **New**, specify the values in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name
JNDI Name	The name (at most 8 characters) to be used by the Create Object Store wizard
Datasource Helper Classname	com.ibm.websphere.rsadapter.ConnectJDBCDataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created in <a href="#">Step 2 on page 201</a>
Container-managed Authentication Alias	(none)
Mapping-configuration Alias	(none)

- c. Under Additional Properties, click **Connection Pool**, set the Min. Connections value to 1, and the Max. Connections value to 150, and then click **OK**.
- d. Click **Apply**, click **Custom Properties** in the Additional Properties table, and specify the values shown in the following table (click the parameter name, specify the value, and then click **OK**):

Parameter	Value
driverType	2
databaseName	Name of Oracle database
serverName	Host name of machine where Oracle is installed
portNumber	1521 (by default)

- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

7. Test the connection as follows:

- a. Navigate to Resources > JDBC Providers and click the non-XA JDBC provider created by Content Engine Setup created or that you created in [Step 4](#)).

- b. Under Additional Properties, click Data Sources, select the check box for the non-XA data source, and then click **Test Connection** to test the connection.
- 8. Create an XA data source for each object store, as follows:
  - a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the XA entry for the Oracle database type of JDBC provider.
  - b. Under Additional Properties, click **Data Sources**, click **New**, specify the values in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name
JNDI Name	The name (at most 8 characters) to be used by the Create Object Store wizard
Datasource Helper Classname	com.ibm.websphere.rsadapter.ConnectJDBCDataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created in <a href="#">Step 2 on page 201</a>
Container-managed Authentication Alias	(none)
Mapping-configuration Alias	(none)

- c. Under Additional Properties, click **Connection Pool**, set the Min. Connections value to 1, and the Max. Connections value to 150, and then click **OK**.
- d. Click Custom Properties in the Additional Properties table, and specify the values shown in the following table (click the parameter name, specify the value, and then click **OK**):

Parameter	Value
driverType	2
databaseName	Name of Oracle database
serverName	Host name of machine where Oracle is installed
portNumber	1521 (by default)

- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

9. Test the connection as follows:
  - a. Navigate to Resources > JDBC Providers and click the XA JDBC provider created by Content Engine Setup created or that you created in [Step 5](#)).
  - b. Under Additional Properties, click Data Sources, select the check box for the non-XA data source, and then click **Test Connection** to test the connection.
10. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

## Task 14b: Configure Content Engine Application Server Database Connectivity (WebSphere 6.0.x)

In this task, you will set up the data sources, connection pools, and (optionally) the distributed (XA) and non-distributed (non-XA) JDBC providers for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Perform one of the following procedures for each object store, depending on your database type:

- [“To configure WebSphere 6.0.x database connectivity \(DB2\)” on page 206](#)
- [“To configure WebSphere 6.0.x database connectivity \(MS SQL Server\)” on page 211](#)
- [“To configure WebSphere 6.0.x database connectivity \(Oracle\)” on page 215](#)

**NOTE** If you are upgrading from version 3.5.x of Content Engine, specify the 3.5.x database or tablespace information in these procedures.

### **To configure WebSphere 6.0.x database connectivity (DB2)**

---

1. Start the WebSphere administrative console (if it isn't already running).
2. For each object store, you need to create a database user alias associated with a unique database and database user, as follows:
  - a. Navigate to Security > Global Security > JAAS Configuration > J2C Authentication data. In the right-hand pane, notice that there is already an alias (created by Content Engine Setup) for the GCD tablespace user.
  - b. Click **New** to create an alias and specify the following information:
    - Alias
    - User ID (object store tablespace user)
    - Password (object store tablespace user password)
    - Description (optional)
  - c. Click **Apply**, and then click Save to apply changes to the master configuration, and then click Save to update the master repository.
3. Navigate to Resources > JDBC Providers, click **Cell** and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both distributed (XA) and non-XA JDBC providers listed.

You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).

4. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC Providers, click **Cell**, and then click **Apply**.

- b. Click **New**, set the property values shown in the following table, and then click **Next**:

Property	Value
Database type	DB2
Provider type	DB2 Universal JDBC Driver Provider
Implementation type	Connection pool data source

- c. Set the property values shown in the following table:

Property	Value
Name	<User-defined non-XA JDBC provider name>
Class path	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc_license_cu.jar \${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar
Native library path	(none)
Implementation class name	com.ibm.db2.jcc.DB2ConnectionPoolDataSource

- d. Click **OK** to view the page with your new JDBC provider.

5. Create an XA JDBC provider, as follows:

- a. Verify that the Cell option button is still selected, and then click **New**.  
 b. Click **New**, and set the property values shown in the following table, and then click **Next**:

Property	Value
Database type	DB2
Provider type	DB2 Universal JDBC Driver Provider
Implementation type	XA data source

- c. Set the property values shown in the following table:

Property	Value
Name	<User-defined XA JDBC provider name>
Class path	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc_license_cu.jar \${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar
Native library path	(none)
Implementation class name	com.ibm.db2.jcc.DB2XADataSource

- d. Click **OK** to view the page with your new JDBC provider.
- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
6. Create a non-XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the DB2 non-XA JDBC provider that Content Engine Setup created or that you created in [Step 4](#).
- b. Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**.

Parameter	Value
Name	<User-defined data source name>
JNDI Name	<User-defined JNDI name> (at most 8 characters) to be used by the Create Object Store wizard.
Data store helper class name	DB2 Universal data store helper (com.ibm.websphere.rsadapter.DB2UniversalDataStoreHelper)
Component-managed authentication alias	<Node>/<alias> created in created in <a href="#">Step 2 on page 206</a>
Container-managed authentication	(none)
Mapping-configuration alias	(none)
Database name	<DB2 database name>
Driver type	4



Parameter	Value
Server name	Host name or IP address of machine where DB2 is installed
Port number	50000 (by default)

- c. Under Additional Properties, click **Custom properties**, specify the property values shown in the following table (click the parameter name, specify the value, and then click **OK**):

Name	Value	Type
resultSetHoldability	1	java.lang.Integer
fullyMaterializeLobData	See Note below.	java.lang.Boolean

**NOTE** The default value of fullyMaterializeLobData is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location
- Users will be storing large content elements (for example, larger than 300 MB)

Setting the value of fullyMaterializeLobData to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.

- d. Click **New** and create the custom property shown in the following table:

Name	Value	Type
webSphereDefaultIsolationLevel	2	java.lang.Integer

- e. Click **OK** to verify that your new custom property is in the custom property table.
- f. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.

7. Test the connection as follows:

- Navigate to Resources > JDBC Providers and click the non-XA JDBC provider created by Content Engine Setup created or that you created in [Step 4](#)).
- Under Additional Properties, click Data Sources, select the check box for the non-XA data source, and then click **Test Connection** to test the connection.

8. Create an XA data source for each object store, as follows:
  - a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the DB2 XA JDBC provider that Content Engine Setup created or that you created in [Step 5](#).
  - b. Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**.

Parameter	Value
Name	The data source name.
JNDI Name	The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Datastore Helper Classname	DB2 Universal data store helper com.ibm.websphere.rsadapter.DB2UniversalDataStoreHelper
Component-managed Authentication Alias	<Node>/<alias> created in created in <a href="#">Step 2 on page 206</a>
Container-managed Authentication	(none)
Mapping-configuration alias	(none)
Database name	Name of DB2 database
Driver type	4
Server name	Host name of machine where DB2 is installed
Port number	50000 (by default)

- c. Under Additional Properties, click **Custom properties**, specify the property values shown in the following table, and then click **Apply**:

Name	Value	Type
fullyMaterializeLobData	See Note below.	java.lang.Boolean

**NOTE**

The default value of fullyMaterializeLobData is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location.
- Users will be storing large content elements (for example, larger than 300 MB).

Setting the value of `fullyMaterializeLobData` to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.

- d. Click **New** and create the custom property shown in the following table:

Name	Value	Type
<code>webSphereDefaultIsolationLevel</code>	2	<code>java.lang.Integer</code>

- e. Click **OK** to verify that your new custom property is in the custom property table.
- f. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
9. Test the connection as follows:
- Navigate to **Resources > JDBC Providers** and click the XA JDBC provider created by Content Engine Setup created or that you created in [Step 5](#)).
  - Under **Additional Properties**, click **Data Sources**, select the check box for the XA data source, and then click **Test Connection** to test the connection.
10. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebSphere 6.0.x database connectivity (MS SQL Server)**

---

- Start the WebSphere administrative console (if it isn't already running).
- Navigate to **Security > Global Security > JAAS Configuration > J2C Authentication Data**. In the right-hand pane, note the alias created for the GCD database user, as you will use this alias for data source connectivity.
- Navigate to **Resources > JDBC Providers**, click **Cell**, and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both distributed and non-XA JDBC providers listed.  
  
 You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).
- Create a non-XA JDBC provider, as follows:
  - Navigate to **Resources > JDBC Providers**, click **Cell**, and then click **Apply**.

- b. Click **New** and set the property values shown in the following table:

Property	Value
Database type	SQL Server
Provider type	WebSphere embedded ConnectJDBC driver for MS SQL Server
Implementation type	<i>Connection pool data source</i>

- c. Click **Next**.

Property	Value
Name	<User-defined non-XA JDBC provider name>
Class path	\${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar
Native library path	(none)
Implementation class name	com.microsoft.sqlserver.jdbc.SQLServerConnectionPoolDataSource

- d. Click **OK** to view the page with your new JDBC provider.

5. Create an XA JDBC provider, as follows:

- a. Verify that the Cell option button is still selected, and then click **New**..  
 b. Click **New** and set the property values shown in the following table:

Property	Value
Database type	SQL Server
Provider type	WebSphere embedded ConnectJDBC driver for MS SQL Server
Implementation type	<i>XA data source</i>

- c. Click **Next**.

- d. Set the property values shown in the following table:

Property	Value
Name	<User-defined XA JDBC provider>
Class path	\${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar

Property	Value
Native library path	(none)
Implementation class name	com.microsoft.sqlserver.jdbcSQLServerXADataSource

- e. Click **OK** to view the page with your new JDBC provider.
  - f. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
6. Create a non-XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the non-XA entry for the SQL Server database type of JDBC provider.
  - b. Under Additional Properties, click **Data sources**, click **New**, specify the information in the following table:

Parameter	Value
Name	<User-defined non-XA data source name>
JNDI name	<User-defined JNDI name> (at most 8 characters) to be used by the Create Object Store wizard.
User-defined data store helper	com.filenet.engine.util.DataStoreHelper
Component-managed authentication alias	<Node>/<alias> created by Content Engine Setup
Container-managed Authentication	(none)
Mapping-configuration alias	(none)

- c. Click **Apply**.

- d. Click **Custom Properties**, and create the properties shown in the following table (To create a property, click **New**, specify the value, and then click **OK**):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
serverName	Host name or IP address of machine where MS SQL Server is installed	java.lang.String
portNumber	1433	java.lang.Integer
selectMethod	direct	java.lang.String
enable2Phase	false	java.lang.Boolean

- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
- f. Navigate to Resources > JDBC Providers, click on the non-XA JDBC provider. Under Additional Properties, click **Data sources**. Click the non-XA data source.
- g. Under Additional Properties, click **WebSphere Application Server data source properties**, set the value of *Statement cache size* to 0, and click **OK**.
- h. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
7. Select the check box for the non-XA data source, and then click **Test Connection** to test the connection.
8. Create an XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the XA entry for the SQL Server database type of JDBC provider.
- b. Under Additional Properties, click **Data sources**, click **New**, specify the information in the following table:

Parameter	Value
Name	<User-defined XA data source name>
JNDI name	<User-defined JNDI name> (at most 8 characters) to be used by the Create Object Store wizard.
User-defined data store helper	com.filenet.engine.util.DataStoreHelper
Component-managed authentication alias	<Node>/<alias> created by Content Engine Setup

Parameter	Value
Container-managed Authentication	(none)
Mapping-configuration alias	(none)

- c. Click **Apply**.
- d. Click **Custom Properties**, and create the properties shown in the following table (To create a property, click **New**, specify the value, and then click **OK**):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
serverName	Host name or IP address of machine where MS SQL Server is installed	java.lang.String
portNumber	1433	java.lang.Integer
selectMethod	direct	java.lang.String
enable2Phase	true	java.lang.Boolean

- e. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
  - f. Navigate to Resources > JDBC Providers, click on the non-XA JDBC provider. Under Additional Properties, click **Data sources**. Click the non-XA data source.
  - g. Under Additional Properties, click **WebSphere Application Server data source properties**, set the value of *Statement cache size* to 0, and click **OK**.
  - h. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
9. Select the check box for the XA data source, and then click **Test Connection** to test the connection.
10. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebSphere 6.0.x database connectivity (Oracle)**

1. Start the WebSphere administrative console (if it isn't already running).
2. For each object store, you need to create a database user alias associated with a unique tablespace and tablespace user, including any temporary tablespaces.

- a. Navigate to Security > Global Security > JAAS Configuration > J2C Authentication data. In the right-hand pane, notice that there is already an alias (created by Content Engine Setup) for the GCD tablespace user.
  - b. Click **New** to create an alias and specify the following information:
    - Alias
    - User ID (object store tablespace user)
    - Password (object store tablespace user password)
    - Description (optional)
  - c. Click **Apply**, click **Save** to apply changes to the master configuration, and then click **Save** to update the master repository.
3. Navigate to Resources > JDBC Providers, click **Cell**, and then click **Apply**. If you chose to have Content Engine Setup create the JNDI names for the GCD data sources, you should see both distributed (XA) and non-XA JDBC providers listed.

You need only one pair (XA and non-XA) of JDBC providers for all your object stores. If you intend to use the JDBC providers created by Content Engine Setup, rather than create new ones, skip to [Step 6](#); otherwise, continue at [Step 4](#).

4. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC Providers and click **Cell** to create a cell-level JDBC provider and then click **Apply**.
  - b. Click **New**, set the property values shown in the following table:

Property	Value
Database type	Oracle
Provider type	Oracle JDBC Driver
Implementation type	<i>Connection pool data source</i>

- c. Click **Next**.
- d. Set the property values shown in the following table:

Property	Value
Name	<i>&lt;User-defined non-XA JDBC provider name&gt;</i>
Class path	`\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Native library path	(none)
Implementation class name	oracle.jdbc.pool.OracleConnectionPoolDataSource



- e. Click **OK** to view the page with your new JDBC provider.
5. Create an XA JDBC provider, as follows:
- a. Navigate to Resources > JDBC Providers and click **Cell** to create a cell-level JDBC provider and then click **Apply**.
  - b. Click **New**, set the property values shown in the following table:

Property	Value
Database type	Oracle
Provider type	Oracle JDBC Driver
Implementation type	<i>XA data source</i>

- c. Click **Next**.
- d. Specify the property values shown in the following table:

Property	Value
Name	<i>&lt;User-defined non-XA data source name&gt;</i>
Class path	`\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Native library path	(none)
Implementation class name	oracle.jdbc.xa.client.OracleXADataSource

- e. Click **OK** to view the page with your new JDBC provider.
6. Create a non-XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the Oracle non-XA JDBC provider that Content Engine Setup created or that you created in [Step 4](#).
  - b. Under Additional Properties, click **Data Sources**, click **New**, and specify the information in the following table:

Parameter	Value
Name	<i>&lt;User-defined non-XA data source name&gt;</i>
JNDI Name	<i>&lt;User-defined JNDI name&gt;</i> (at most 8 characters) to be used by the Create Object Store wizard.
Use this Data source in container managed persistence (CMP)	Clear the check box

Parameter	Value
Data store helper class name	com.ibm.websphere.rsadapter.Oracle10gDataStoreHelper
Component-managed authentication alias	<Node>/<alias> created in <a href="#">Step 2 on page 215</a>
Container-managed Authentication	(none)
URL	jdbc:oracle:thin:@<DB_HostName>:<Port>:<SID> Default <Port>: 1521 (by default) <DB_HostName>: host name or IP address of machine where Oracle is installed <SID>: the Oracle system identifier.

- c. Click **OK**.
  - d. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
7. Select the check box for the non-XA data source, and then click **Test Connection** to test the connection.
  8. Create an XA data source for each object store, as follows:
    - a. Navigate to Resources > JDBC Providers. In the table of JDBC providers, click the entry for the Oracle XA JDBC provider that Content Engine Setup created or that you created in [Step 5](#).
    - b. Under Additional Properties, click **Data Sources**, click **New**, and specify the information in the following table:

Parameter	Value
Name	<User-defined XA data source name>
JNDI Name	<User-defined JNDI name> (at most 8 characters) to be used by the Create Object Store wizard.
Use this Data source in container managed persistence (CMP)	Clear the check box
Data store helper class name	com.ibm.websphere.rsadapter.Oracle10gDataStoreHelper
Component-managed authentication alias	<Node>/<alias> created in <a href="#">Step 2 on page 215</a>

Parameter	Value
Authentication alias for XA recovery	Click <i>Use component-managed authentication alias</i> option button
Container-managed Authentication	(none)
URL	jdbc:oracle:thin:@<DB_HostName>:<Port>:<SID> Default <Port>: 1521 (by default) <DB_HostName>: host name or IP address of machine where Oracle is installed <SID>: Oracle system identifier

- c. Click **OK**.
  - d. Click **Save** to apply your changes to the master configuration, and then click **Save** to update the master repository with your changes.
9. Select the check box for the XA data source, and then click **Test Connection** to test the connection.
  10. If you are upgrading Content Engine, go to ["To edit the upgrader utility file" on page 499](#); otherwise, continue at ["Prepare Storage Areas for Object Stores" on page 252](#).

## Task 14c: Configure Content Engine Application Server Database Connectivity (WebSphere 6.1.x)

In this task, you will set up the data sources, connection pools, and (optionally) the distributed (XA) and non-distributed (non-XA) JDBC providers for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Perform one of the following procedures for each object store, depending on your database type:

**NOTE** If you are upgrading from version 3.5.x of Content Engine, specify the 3.5.x database or tablespace information in these procedures.

- [“To configure WebSphere 6.1.x database connectivity \(DB2\)” on page 220](#)
- [“To configure WebSphere 6.1.x database connectivity \(MS SQL Server\)” on page 226](#)
- [“To configure WebSphere 6.1.x database connectivity \(Oracle\)” on page 231](#)

### **To configure WebSphere 6.1.x database connectivity (DB2)**

---

1. Start the WebSphere administrative console (if it isn't already running).
2. For each object store, create a database user alias associated with a unique database and database user, as follows:
  - a. Navigate to Security > Secure administration, applications, and infrastructure > Java Authentication and Authorization Service > J2C authentication data. In the right-hand pane, notice that there is already an alias (created by Content Engine Setup) for the GCD tablespace user.
  - b. Click **New** to create a J2C authentication alias and specify the following information:
    - Alias
    - User ID (object store tablespace user)
    - Password (object store tablespace user password)
    - Description (optional)
  - c. Click **Apply** and then click **Save**.

Create (non-XA and XA) JDBC providers, as shown in [Step 3](#) and [Step 4](#), to serve your object stores. Or skip to [Step 5 on page 222](#) and use the existing JDBC providers Content Engine Setup created for the GCD.

3. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC > JDBC Providers, and choose Cell from the Scope drop-down list.

- b. Click **New** and set the property values shown in the following table, and then click **Next**:

Property	Value
Database type	DB2
Provider type	DB2 Universal JDBC Driver Provider
Implementation type	<i>Connection pool data source</i> to create a non-XA JDBC provider
Name	<User-defined non-XA JDBC_Provider_Name>
Description	An optional comment

- c. Specify or verify the following values, and then click **Next**:

Property	Value
Class path	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar \${DB2UNIVERSAL_JDBC_DRIVER_PATH}db2jcc_license_cu.jar
Directory location	Accept the value you specified in <a href="#">“Configure an Application Server for Content Engine (WebSphere)” on page 106</a> .
Native library path	(None)

- d. View the Summary of the new JDBC provider, click **Finish**, and then click **Save**.

4. Create an XA JDBC provider, as follows:

- a. Click **New**, set the property values shown in the following table, and then click **Next**:

Property	Value
Database type	DB2
Provider type	DB2 Universal JDBC Driver Provider
Implementation type	<i>XA data source</i> to create an XA data source
Name	<User-defined XA JDBC_Provider_Name>
Description	An optional comment

- b. Specify or verify the following values, and then click **Next**:

Property	Value
Class path	\${DB2UNIVERSAL_JDBC_DRIVER_PATH}/db2jcc.jar \${DB2UNIVERSAL_JDBC_DRIVER_PATH}db2jcc_license_cu.jar
Directory location	Accept the value you specified in <a href="#">“Configure an Application Server for Content Engine (WebSphere)”</a> on page 106.
Native library path	(None)

5. Create a non-XA data source for each object store, as follows:

- a. Navigate to Resources > JDBC > Data sources, choose Cell from the Scope drop-down list, and then click **New**.
- b. Specify the values in the following table and then click **Next**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias and XA recovery authentication	Specify the J2C authentication alias created for this object store in <a href="#">Step 2 on page 220</a> .

- c. Click the *Select an existing JDBC provider* option. Choose, from the drop-down list, one of the following, and then click **Next**.
- *JDBC Provider for GCD Datatsource* (created by Content Engine Setup)
  - The JDBC provider you created in [Step 3 on page 220](#), and then click **Next**.
- d. Specify the values in the following table, and then click **Next**:

Parameter	Value
Database name	Name of DB2 database
Driver type	4
Server name	Host name or IP address of machine where DB2 is installed

Parameter	Value
Port number	<DB_Port_Number> (default is 50000)
Use this data source in container managed persistence (CMP)	Clear the check box

- e. Review the Summary page, click **Finish**, and then click **Save**.
  - f. If you have created all your non-XA data sources, continue at [Step 6](#); otherwise, jump back to [Step b on page 222](#) to create a non-XA data source for another object store.
6. Create an XA data source for each object store, as follows:
- a. Navigate to Resources > JDBC > Data sources, and choose Cell from the Scope drop-down list.
  - b. Click **New** and specify the values in the following table and then click **Next**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias and XA recovery authentication alias	Specify the J2C authentication alias created for this object store in <a href="#">Step 2 on page 220</a> .

- c. Click the *Select an existing JDBC provider* option. Choose, from the drop-down list, one of the following, and then click **Next**.
  - *JDBC Provider for GCD Datatsource* (created by Content Engine Setup).
  - The JDBC provider you created in [Step 4 on page 221](#).
- d. Specify the values in the following table, and then click **Next**:

Parameter	Value
Database name	Name of DB2 database
Driver type	4
Server name	Host name or IP address of machine where DB2 is installed

Parameter	Value
Port number	<DB_Port_Number> (default is 50000)
Use this data source in container managed persistence (CMP)	Clear the check box

- e. Review the Summary page, click **Finish**, and then click **Save**.
  - f. If you have created all your XA data sources, continue at [Step 7](#); otherwise, jump back to [Step b on page 223](#) to create an XA data source for another object store.
7. For each non-XA data source you created in [Step 5 on page 222](#), edit its custom properties, as follows:
- a. Navigate to Resources > JDBC > Data sources.
  - b. Select a non-XA data source.
  - c. Under Additional Properties, click **Custom Properties**, click **New** and create the custom property shown in the following table (click **OK** after specifying the value):

Name	Value	Type
selectMethod	direct	java.lang.String
enable2Phase	false	java.lang.Boolean
webSphereDefaultIsolationLevel	2	java.lang.Integer
fullyMaterializeLobData	See Note below.	java.lang.Boolean

**NOTE** The default value of fullyMaterializeLobData is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location
  - Users will be storing large content elements (for example, larger than 300 MB)
- d. Setting the value of fullyMaterializeLobData to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.
  - e. Click **Save**.
  - f. If you have edited the custom properties for all your non-XA data sources, continue at [Step 8](#); otherwise, jump back to [Step b on page 224](#) to edit the custom properties for another non-XA data source.



8. For each XA data source you created in [Step 6 on page 223](#), edit its custom properties, as follows:

- a. Navigate to Resources > JDBC > Data sources.
- b. Select a XA data source that you created in [Step 6 on page 223](#).
- c. Under Additional Properties, click **Custom Properties**, (click **OK** after specifying the value):

Name	Value	Type
selectMethod	direct	java.lang.String
enable2Phase	true	java.lang.Boolean
webSphereDefaultIsolationLevel	2	java.lang.Integer
fullyMaterializeLobData	See Note below.	java.lang.Boolean

**NOTE** The default value of fullyMaterializeLobData is *true*; but you may want to set it to *false* if both of the following conditions apply:

- You will be creating an object store (see [“Create Object Stores” on page 258](#)) with a database storage area as the content storage location.
- Users will be storing large content elements (for example, larger than 300 MB).

Setting the value of fullyMaterializeLobData to *false* prevents performance degradation or memory-related errors during retrieval or indexing of large content elements. Alternatively, specify a file storage area instead of a database storage area when creating an object store.

- d. Click **Save**.
- e. If you have edited the custom properties for all your XA data sources, continue at [Step 9](#); otherwise, jump back to [Step b on page 225](#) to edit the custom properties for another XA data source.

9. Edit the connection pool settings for the (non-XA and XA) data sources, as follows:

- a. Navigate to Resources > JDBC > Data sources.
- b. For each data source in the data source list, do the following:
  - i. Click on the data source.
  - ii. Under Additional Properties, click **Connection pool properties**, and set the parameter values shown in the following table:

Parameter	Value
Maximum connections	[7 * (# of object stores)] + (# of execution threads) + 2
Minimum connections	(# of execution threads) + 10

- iii. Save your changes.
10. Test data source connectivity, as follows:
    - a. Navigate to Resources > JDBC > Data sources.
    - b. Select the check boxes for the non-XA and XA data sources you created in this procedure.
    - c. Click **Test connection** and verify that all test connection operations were successful.
  11. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

### **To configure WebSphere 6.1.x database connectivity (MS SQL Server)**

---

1. Start the WebSphere administrative console (if it isn't already running).
2. Navigate to Security > Secure administration, applications, and infrastructure > Java Authentication and Authorization Service > J2C Authentication Data. In the right-hand pane note the alias of the GCD database user.

If you will be using the same MS SQL Server user to connect to all your databases, use this alias to configure database connectivity and skip to [Step 3](#); otherwise, create a J2C authentication alias for each database, as follows:

- a. Click **New** to create an alias and specify the following information:
  - Alias
  - User (object store database user)
  - Password (object store database user password)
  - Description (optional)
- b. Save your changes.
- c. If you have created all your J2C authentication aliases, continue at [Step 3](#); otherwise, jump back to [Step a](#) to create another J2C authentication alias.

Create (non-XA and XA) JDBC providers, as shown in [Step 3](#) and [Step 4](#), to serve your object stores. Or skip to [Step 5 on page 227](#) and use the existing JDBC providers Content Engine Setup created for the GCD.

3. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC > JDBC Providers, and select Cell = `<profileNodeCell>` from the scope drop-down list, where *profile* is the name of the profile you created for this instance of Content Engine.

- b. Click **New** and set the property values shown in the following table for the non-XA data source:

Property	Value
Database type	User-defined
Implementation class name	com.microsoft.sqlserver.jdbc.SQLServerConnectionPoolDataSource
Name	<non-XA User-defined JDBC_Provider_Name>
Description	An optional comment

- c. Click **Next** and then specify the value for Class path:

`${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar`

- d. Click **Next**, and then click **Finish** and **Save**.

- 4. Create an XA JDBC provider, as follows:

- a. Navigate to Resources > JDBC > JDBC Providers, and select Cell = <profileNodeCell> from the scope drop-down list.
- b. Click **New** and set the property values shown in the following table for the XA data source:

Property	Value
Database type	User-defined
Implementation class name	com.microsoft.sqlserver.jdbc.SQLServerXADataSource
Name	<XA User-defined JDBC_Provider_Name>
Description	An optional comment

- c. Click **Next** and then specify the value for Class path:

`${MSSQLSERVER_JDBC_DRIVER_PATH}/sqljdbc.jar`

- d. Click **Next**, and then click **Finish** and **Save**.

- 5. Create a non-XA data source for each object store, as follows:

- a. Click the non-XA JDBC provider you created in [Step 3 on page 226](#).
- b. Click **Data sources** under Additional Properties and then click **New**.

- c. Specify the information in the following table, and then click **Apply**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias and XA recovery authentication alias	<Node>/<alias>  If all databases will be accessed by only one MS SQL Server user, specify the alias created by Content Engine Setup.  If each database will be accessed by its own MS SQL Server user, specify an alias defined in <a href="#">Step 2 on page 226</a> .

- d. Click **Next** and then specify the information in the following table:

Parameter	Value
Data store helper class name	com.filenet.engine.util.DataStoreHelper
Use this data source in container managed persistence (CMP)	Clear the check box

- e. Click **Next**. Review the Summary page, click **Finish**, and then click **Save** to save the changes to the master configuration.
- f. If you have created all your non-XA data sources, continue at [Step 6](#); otherwise, jump back to [Step b](#) to create a non-XA data source for another object store.
6. Create an XA data source for each object store, as follows:
- Navigate to Resources > JDBC > JDBC Providers. In the table of JDBC providers, click the XA User-Defined JDBC provider.

- b. Under Additional Properties, click **Data Sources**, click **New**, specify the information in the following table, and then click **Apply**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias	<Node>/<alias> If all databases will be accessed by only one MS SQL Server user, specify the alias created by Content Engine Setup. If each database will be accessed by its own MS SQL Server user, specify an alias defined in <a href="#">Step 2 on page 226</a> .

- c. Click **Next**, and then specify the information in the following table:

Parameter	Value
Data store helper class name	com.filenet.engine.util.DataStoreHelper
Use this data source in container managed persistence (CMP)	Clear the check box

- d. Click **Next**. Review the Summary page, click **Finish** and then click **Save** to save the changes to the master configuration.
- e. If you have created all your XA data sources, continue at [Step 7](#); otherwise, jump back to [Step b](#) to create a non-XA data source for another object store.
7. For each non-XA data source you created in [Step 5 on page 227](#), edit its custom properties, as follows:
- a. Navigate to Resources > JDBC > Data sources.

- b. Under Additional Properties, click **Custom Properties**, and select and specify values for each of the properties in the following table (click **OK** after specifying each value):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
portNumber	<DB_Port_Number> (default is 1433)	java.lang.Integer
selectMethod	direct	java.lang.Integer
serverName	<DB_HostName> or IP address	java.lang.String

- c. Click **New** and create the custom property shown in the following table:

Name	Value	Type
enable2Phase	false	java.lang.Boolean

- d. Click **Save**.
- e. If you have edited the custom properties for all your non-XA data sources, continue at [Step 8](#); otherwise, jump back to [Step b on page 230](#) to edit the custom properties for another non-XA data source.
8. For each XA data source you created in [Step 6 on page 228](#), edit its custom properties, as follows:
- Navigate to Resources > JDBC > Data sources.
  - Select the XA data source that you created in [Step 6 on page 228](#), edit its custom properties
  - Under Additional Properties, click **Custom Properties**, and select and specify values for each of the properties in the following table (click **OK** after specifying each value):

Name	Value	Type
databaseName	<DB_Name>	java.lang.String
portNumber	<DB_Port_Number> (default is 1433)	java.lang.Integer
selectMethod	direct	java.lang.Integer
serverName	<DB_HostName> or IP address	java.lang.String

- d. Click **New** and create the custom property shown in the following table (click **OK** after specifying the value):

Name	Value	Type
enable2Phase	true	java.lang.Boolean

- e. Click **Save**.
  - f. If you have edited the custom properties for all your XA data sources, continue at [Step 9](#); otherwise, jump back to [Step b on page 230](#) to edit the custom properties for another XA data source.
9. Edit the connection pool settings for the (non-XA and XA) data sources, as follows:
- a. Navigate to Resources > JDBC > Data sources.
  - b. For each data source in the data source list, do the following:
    - i. Click on the data source.
    - ii. Under Additional Properties, click **Connection pool properties**, and set the parameter values shown in the following table:

Parameter	Value
Maximum connections	$[7 * (\text{\# of object stores})] + (\text{\# of execution threads}) + 2$
Minimum connections	$(\text{\# of execution threads}) + 10$

- iii. Save your changes.
10. Test data source connectivity, as follows:
- a. Navigate to Resources > JDBC > Data sources.
  - b. Select the check boxes for the non-XA and XA data sources you created in this procedure.
  - c. Click **Test connection** and verify that all test connection operations were successful.
11. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebSphere 6.1.x database connectivity (Oracle)**

---

- 1. Start the WebSphere administrative console (if it isn't already running).
- 2. For each object store, create a database user alias associated with a unique tablespace and tablespace user, including any temporary tablespaces, as follows:
  - a. Navigate to Security > Secure administration, applications, and infrastructure > Global Security > JAAS Configuration > J2C Authentication. In the right-hand pane, notice that there is already an alias (created by Content Engine Setup) for the GCD tablespace user.

- b. Click **New** to create an alias and specify the following information:
  - Alias
  - User (object store tablespace user)
  - Password (object store tablespace user password)
  - Description (optional)
- c. Click **Apply** and then click **Save**.

Create (non-XA and XA) JDBC providers, as shown in [Step 3](#) and [Step 4](#), to serve your object stores. Or skip to [Step 5 on page 233](#) and use the existing JDBC providers content Engine Setup created for the GCD.

3. Create a non-XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC > JDBC Providers, and choose **Cell** from the Scope drop-down list.
  - b. Click **New**, set the property values as shown in the following table and then click **Next**:

Property	Value
Database type	Oracle
Provider type	Oracle JDBC Driver
Implementation type	Connection pool data source
Name	<User-defined non-XA JDBC_Provider_Name>
Description	An optional comment

- c. Specify or verify the values in the following table, and then click **Next**:

Property	Value
Class path	\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Directory location	Accept the value you specified in <a href="#">“Configure an Application Server for Content Engine (WebSphere)” on page 106</a> .

- d. View the Summary of the new JDBC provider, click **Finish**, and then click **Save**.
4. Create an XA JDBC provider, as follows:
  - a. Navigate to Resources > JDBC > JDBC Providers and click **Cell** from the drop-down list.



- b. Click **New**, set the property values shown in the following table, and then click **Next**:

Property	Value
Database type	Oracle
Provider type	Oracle JDBC Driver
Implementation type	XA data source
Name	<User-defined XA JDBC_Provider_Name>
Description	An optional comment

- c. Specify or verify the values in the following table, and then click **Next**:

Property	Value
Class path	\${ORACLE_JDBC_DRIVER_PATH}/ojdbc14.jar
Directory location	Accept the value you specified in <a href="#">“Configure an Application Server for Content Engine (WebSphere)”</a> on page 106.

- d. View the Summary of the new JDBC provider, click **Finish**, and then click **Save**.

5. Create a non-XA data source for each object store, as follows:

- a. Navigate to Resources > JDBC > Data sources, and choose Cell from the Scope drop-down list.
- b. Click **New** and specify the values in the following table and then click **Next**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias and XA recovery authentication	Specify the J2C authentication alias created for this object store in <a href="#">Step 2 on page 231</a> .

- c. Click the *Select an existing JDBC provider* option and then choose, from the drop-down list, one of the following:
- *JDBC provider for GCD Datasource*, which Content Engine Setup created,

- The JDBC provider you created in [Step 3 on page 232](#)
- d. Specify the values in the following table and then click **Next**:

Parameter	Value
URL	jdbc:oracle:thin:@<DB_HostName>:<Port>:<SID> <b>NOTE</b> <Port> is 1521 by default. <SID> is the Oracle system identifier.
Datastore Helper Classname	Depending on your Oracle version, choose one: <ul style="list-style-type: none"> <li>• <i>Oracle9i and prior data store helper</i></li> <li>• <i>Oracle10g data store helper</i></li> </ul>
Use this Data source in container managed persistence (CMP)	Clear the check box

- e. Review the Summary page, click **Finish**, and then click **Save**.
- f. If you have created all your non-XA data sources, continue at [Step 6](#); otherwise, jump back to [Step b on page 233](#).
6. Create an XA data source for each store, as follows:
- a. Navigate to Resources > JDBC > Data sources, and choose Cell from the Scope drop-down list.
  - b. Click **New** and specify the values in the following table and then click **Next**:

Parameter	Value
Data source name	<User-defined_data_source_name>
JNDI name	<User-defined_JNDI_name> The name of the data source (at most 8 characters) to be used by the Create Object Store wizard.
Component-managed authentication alias and XA recovery authentication	Specify the J2C authentication alias created for this object store in <a href="#">Step 2 on page 231</a> .

- c. Click the *Select an existing JDBC provider* option and then choose, from the drop-down list, one of the following:
- *JDBC provider for GCD Datasource*, which Content Engine Setup created,
  - The JDBC provider you created in [Step 4 on page 232](#)

d. Specify the values in the following table and then click **Next**:

Parameter	Value
URL	jdbc:oracle:thin:@<DB_HostName>:<Port>:<SID> <b>NOTE</b> <Port> is 1521 by default. <SID> is the Oracle system identifier.
Datastore Helper Classname	Depending on your Oracle version, choose one: <ul style="list-style-type: none"> <li>Oracle9i and prior data store helper</li> <li>Oracle10g data store helper</li> </ul>
Use this Data source in container managed persistence (CMP)	Clear the check box

e. Review the Summary page, click **Finish**, and then click **Save**.

f. If you have created all your XA data sources, continue at [Step 7](#); otherwise, jump back to [Step b on page 234](#).

7. Edit the connection pool settings for the (non-XA and XA) data sources, as follows:

a. Navigate to Resources > JDBC > Data sources.

b. For each data source in the data source list, do the following:

i. Click on the data source.

ii. Under Additional Properties, click **Connection pool properties**, and set the parameter values shown in the following table:

Parameter	Value
Maximum connections	$[7 * (\text{\# of object stores})] + (\text{\# of execution threads}) + 2$
Minimum connections	$(\text{\# of execution threads}) + 10$

iii. Save your changes.

8. Test data source connectivity, as follows:

a. Navigate to Resources > JDBC > Data sources.

b. Select the check boxes for the non-XA and XA data sources you created in this procedure.

c. Click **Test connection** and verify that all test connection operations were successful.

9. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

## Task 14d: Configure Content Engine Application Server Database Connectivity (WebLogic 8.1.x)

In this task, you will set up the data sources and connection pools needed for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Do the procedures in one of the following sections, depending on your database type.

The procedures in this topic use the terminology in the following table (where the abbreviation *XA* means *distributed*, and *non-XA* means *non-distributed*). For each term in the table, you will substitute a value appropriate to your configuration when you create the connection pools and data sources.

Term	Meaning
AdminServer	The default name of the administration server for a WebLogic Server domain
<ConnectionPoolName>	The names of the XA and non-XA connection pools
<DB_HostName>	The host name or IP address of the database machine
<DB>	The meaning depends on the database type: <ul style="list-style-type: none"> <li>(DB2) The name of the database containing the object store tablespace</li> <li>(MS SQL Server) The name of the object store database</li> <li>(Oracle) The system identifier (SID) of the database containing the object store tablespace</li> </ul>
<DB_UserName> (MS SQL Server)	The account Content Engine uses to access <DB>
<DB_Password> (MS SQL Server)	The password for <DB_UserName>
<TS> (DB2 and Oracle)	The name of the object store tablespace
<DataSourceName>	The names of the XA and non-XA data sources
<DataSourceJNDI_Name>	The JNDI name of the data source <b>NOTE</b> It is strongly recommended that <DataSourceJNDI_Name> be identical to <DataSourceName>
<Port>	The port number used by the database engine: <ul style="list-style-type: none"> <li>50000 (default for DB2)</li> <li>1433 (default for MS SQL Server)</li> <li>1521 (default port for Oracle)</li> </ul>

As part of configuring database connectivity, you have the opportunity to customize the maximum number of connection pools available for object-store data sources. Here are some guidelines for choosing this value.

- In general, the value should be proportional to the maximum number of concurrent users of the object store.
- A nominal value of 25 should be adequate for object stores.

**CAUTION** Due to a known problem in all versions of WebLogic supported by IBM FileNet P8, if you intend to create an object store in the same database (MS SQL Server) or tablespace (Oracle or DB2) as the GCD, use the data sources and connection pools you set up for the GCD. Do *not* create separate data sources and connection pools. Be aware that the GCD can share its database (tablespace) with at most one object store.

Do the procedures in one of the following sections, depending on your database type.

- [“To configure WebLogic 8.1.x database connectivity \(DB2\)” on page 237](#)
- [“To configure WebLogic 8.1.x database connectivity \(MS SQL Server\)” on page 238](#)
- [“To configure WebLogic 8.1.x database connectivity \(Oracle\)” on page 239](#)

**NOTE** If you are upgrading from version 3.5.x of Content Engine, specify the 3.5.x database or tablespace information in these procedures.

**To configure WebLogic 8.1.x database connectivity (DB2)**

---

Do this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two connection pools—one for XA transactions and one for non-XA transactions—with the parameter values shown in the following table:

Parameter	Value
Database Type	DB2
Database Driver	Other
Name	<ConnectionPoolName>
Driver Class Name (XA)	com.ibm.db2.jcc.DB2XADataSource
Driver Class Name (non-XA)	com.ibm.db2.jcc.DB2Driver
URL of database	jdbc:db2://localhost:<Port>/<DB_Name>
Tablespace User Name	<TS_UserName>
Password (DB2 and Oracle only)	<TS_Password>
Host Name	<DB_HostName>

Parameter	Value
Port	<Port>
driverType	4

3. Create two data sources—one for XA and one for non-XA transactions—with the parameter values shown in the following table:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDI_Name>
Honor Global Transactions	<ul style="list-style-type: none"> <li>• Select for (XA transactions)</li> <li>• Clear (for non-XA transactions)</li> </ul>
Emulate Two-Phase Commit	Clear
Pool Name	<ConnectionPoolName>

4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebLogic 8.1.x database connectivity (MS SQL Server)**

Perform this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two connection pools—one for XA and one for non-XA transactions—with the parameter values shown in the following table:

Parameter	Value
Database Type	MS SQL Server
Database Driver	Other
Name	<ConnectionPoolName>
XA Driver Classname	com.microsoft.sqlserver.jdbc.SQLServerXADataSource
Non-XA Driver Classname	com.microsoft.sqlserver.jdbc.SQLServerDriver

Parameter	Value
URL of database	jdbc:sqlserver://<DatabaseHostName>:<Port>
Database username	<DB_UserName>
Properties	user=<DB_UserName> portNumber=<Port> databaseName=<DB_Name> serverName=<DB_HostName>
Statement Cache Size	0
Enable XA Transaction Timeout	Select

3. Create two data sources—one for XA and one for non-XA transactions—with the parameter values shown in the following table:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDIName>
Honor Global Transactions	<ul style="list-style-type: none"> <li>• Select (for XA transactions)</li> <li>• Clear (for non-XA transactions)</li> </ul>
Emulate Two-Phase Commit	Clear
Pool Name	<ConnectionPoolName>

4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

#### To configure WebLogic 8.1.x database connectivity (Oracle)

Perform this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two connection pools—one for distributed transactions (XA) and one for non-distributed (non-XA) transactions—with the parameter values shown in the following table:

Parameter	Value
Database Type	Oracle
Database Driver	Oracle
Name	<ConnectionPoolName>

Parameter	Value
Driver Class Name (XA)	oracle.jdbc.xa.client.OracleXADataSource
Driver Class Name (non-XA)	oracle.jdbc.OracleDriver
URL of database	jdbc:oracle:thin:@<DB_HostName>:<Port>:<DB_Name
Tablespace username	<TablespaceUserName>
Password	<TablespacePassword>
Database Name	<DB>
Host Name	<DB_HostName>
Port	<Port>
Enable XA Transaction Timeout (XA)	Select
<b>NOTE</b> To view this parameter, select Advanced Options and then click <b>Show</b> .	

3. Create two data sources—one for XA and one for non-XA transactions—with the parameter values shown in the following table:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDIName>
Honor Global Transactions	<ul style="list-style-type: none"> <li>• Select (for XA transactions)</li> <li>• Clear (for non-XA transactions)</li> </ul>
Emulate Two-Phase Commit	Clear
Pool Name	<ConnectionPoolName>

4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).



## Task 14e: Configure Content Engine Application Server Database Connectivity (WebLogic 9.2.x)

In this task, you will set up the data sources and connection pools needed for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Do the procedures in one of the following sections, depending on your database type.

The procedures in this topic use the terminology in the following table (where the abbreviation *XA* means *distributed*, and *non-XA* means *non-distributed*). For each term in the table, you will substitute a value appropriate to your configuration when you create the connection pools and data sources.

Term	Meaning
AdminServer	The default name of the administration server for a WebLogic Server domain
<ConnectionPoolName>	The names of the XA and non-XA connection pools
<DB_HostName>	The host name or IP address of the database machine
<DB>	The meaning depends on the database type: <ul style="list-style-type: none"> <li>• (DB2) The name of the database containing the object store tablespace</li> <li>• (MS SQL Server) The name of the object store database</li> <li>• (Oracle) The system identifier (SID) of the database containing the object store tablespace</li> </ul>
<DB_UserName>	The account Content Engine uses to access <DB>
<DB_Password>	The password for <DB_UserName>
<TS>	The name of the object store tablespace (DB2 and Oracle)
<TS_UserName>	The account Content Engine uses to access <TS>
<TS_Password>	The password for <TS_Name>
<DataSourceName>	The names of the XA and non-XA data sources

Term	Meaning
<DataSourceJNDI_Name>	The JNDI name of the data source <b>NOTE</b> It is strongly recommended that <DataSourceJNDI_Name> be identical to <DataSourceName>
<Port>	The port number used by the database engine: <ul style="list-style-type: none"> <li>• 50000 (default for DB2)</li> <li>• 1433 (default for SQL Server)</li> <li>• 1521 (default port for Oracle)</li> </ul>

As part of configuring database connectivity, you have the opportunity to customize the maximum number of connection pools available for object-store data sources. Here are some guidelines for choosing this value.

- In general, the value should be proportional to the maximum number of concurrent users of the object store.
- A nominal value of 25 should be adequate for object stores.

**CAUTION** Due to a known problem in all versions of WebLogic supported by IBM FileNet P8, if you intend to create an object store in the same database as the GCD, use the data sources and connection pools you set up for the GCD. Do *not* create separate data sources and connection pools.

Do the procedures in one of the following sections, depending on your database type.

- [“To configure WebLogic 9.2.x database connectivity \(DB2\)” on page 242](#)
- [“To configure WebLogic 9.2.x database connectivity \(MS SQL Server\)” on page 244](#)
- [“To configure WebLogic 9.2.x database connectivity \(Oracle\)” on page 245](#)

**NOTE** If you are upgrading from version 3.5.x of Content Engine, specify the 3.5.x database or tablespace information in these procedures.

### To configure WebLogic 9.2.x database connectivity (DB2)

Perform this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two data sources—one for distributed transactions (XA) and one for non-distributed (non-XA) transactions—with the parameter values shown in the following table:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDIName>

Parameter	Value
Database Type	DB2
Database Driver	Other
Support Global Transactions	<ul style="list-style-type: none"> <li>Select (for XA)</li> <li>Clear (for non-XA)</li> </ul>
Emulate Two-Phase Commit	Clear (XA only)
Database Name	<DB_Name>
Host Name	<DB_HostName>
Port	<Port>
Database User Name	<DB_DB2_Login>
Password	<DB_DB2_Password>
Set XA Transaction Timeout	Select

In the Test Database Connection screen, specify the parameter values in the following table:

Parameter	Value
Driver Class Name (XA)	com.ibm.db2.jcc.DB2XADataSource
Driver Class Name (non-XA)	com.ibm.db2.jcc.DB2Driver
URL	jdbc:db2://<DB_HostName>:<port>/<DB_Name>
Database User Name	<DB_UserName>
Password	<DB_Password>
Properties	user=<DB_UserName> databaseName=<DB_Name> portnumber=<Port> serverName=<DB_HostName> selectMode=default (XA only)
Test Table Name	SQL select count(*) from dual
Select Targets	AdminServer

- In the Test Database Connection screen, click **Test Configuration** to verify the database connection.
- If you are upgrading Content Engine, go to ["To edit the upgrader utility file" on page 499](#); otherwise, continue at ["Prepare Storage Areas for Object Stores" on page 252](#).

**To configure WebLogic 9.2.x database connectivity (MS SQL Server)**

Perform this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two data sources—one for distributed transactions (XA) and one for non-distributed (non-XA) transactions—with the parameter values shown in the following tables:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDIName>
Database Type	MS SQL Server
Database Driver	Other
Support Global Transactions	<ul style="list-style-type: none"> <li>• Select (for distributed transactions)</li> <li>• Clear (for non-distributed transactions)</li> </ul>
Emulate TwoPhase Commit (XA only)	Clear
Database Name	<DB_Name>
Host Name	<DB_HostName>
Port	<Port>
Database User Name	<DB_UserName>
Password	<DB_Password>
Driver Classname (XA)	com.microsoft.sqlserver.jdbc.SQLServerXADataSource
Driver Classname (Non-XA)	com.microsoft.sqlserver.jdbc.SQLServerDriver
URL of database	jdbc:sqlserver://<DatabaseHostName>:<PortNumber>
Properties	user=<DB_UserName> portNumber=<Port> databaseName=<DB_Name> serverName=<DB_HostName>
Statement Cache Size	0
Test Table Name	SQL select count(*) from sysusers
Select Targets	AdminServer
Set XA Transaction Timeout	Select

3. In the Test Database Connection screen, click **Test Configuration** to verify the database connection.
4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

**To configure WebLogic 9.2.x database connectivity (Oracle)**

---

Perform this procedure on the WebLogic machine where Content Engine is installed.

1. Log on to WebLogic Server Administration Console.
2. Create two data sources—one for distributed transactions (XA) and one for non-distributed (non-XA) transactions—with the parameter values shown in the following table:

Parameter	Value
Name	<DataSourceName>
JNDI Name	<DataSourceJNDIName>
Database Type	Oracle
Database Driver	<ul style="list-style-type: none"> <li>• Oracle's Driver (XA Thin) (for XA)</li> <li>• Oracle's Driver (Thin) (for non-XA)</li> </ul>
Support Global Transactions	<ul style="list-style-type: none"> <li>• Select (for XA)</li> <li>• Clear (for non-XA)</li> </ul>
Emulate Two-Phase Commit	Clear (XA only)
Database Name	<DB>
Host Name	<DB_HostName>
Port	<Port>
Database User Name	<DB_UserName>
Password	<DB_Password>
Set XA Transaction Timeout	Select

In the Test Database Connection screen, specify the parameter values in the following table:

Parameter	Value
Driver Class Name (XA)	oracle.jdbc.client.OracleXADataSource
Driver Class Name (non-XA)	oracle.jdbc.OracleDriver
URL	jdbc:oracle:thin:@<DB_HostName>:<port>:<DB_Name>

Parameter	Value
Database User Name	<DB_UserName>
Password	<DB_Password>
Properties	user=<DB_UserName> databaseName=<DB_Name> portnumber=<Port> serverName=<DB_HostName> selectMode=default (XA only)
Test Table Name	SQL select count(*) from dual
Select Targets	AdminServer

3. In the Test Database Connection screen, click **Test Configuration** to verify the database connection.
4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

## Task 14f: Configure Content Engine Application Server Database Connectivity (JBoss 4.0.x)

In this task, you will set up the XA (distributed) and non-XA data sources needed for Content Engine to communicate with the databases or tablespaces associated with the object stores you will create in a later task.

Do the procedures in one of the following sections, depending on your database type. These procedures assume the following environment for your JBoss installation (substitute your own environment in its place where applicable):

- **JBOSS\_DIST** is the home directory of your JBoss installation.
- **server1** is the JBoss server instance on which you will install and deploy Content Engine.
- **os1-ds.xml** and **os1-xa-ds.xml** are non-XA and XA data sources for object store os1.

### To configure JBoss 4.0.x database connectivity (DB2)

---

1. To create data sources for an object store, do the following:
  - a. Copy files **FNGCD-ds.xml** and **FNGCD-xa-ds.xml** from directory **JBOSS\_DIST/server/server1/deploy** to new files, **os1-ds.xml** and **os1-xa-ds.xml**, respectively, in this directory.  
**NOTE** The name of each XML file you create for your data sources must end in **-ds.xml**.
  - b. Edit **os1-ds.xml** as follows:
    - i. Remove the <local-tx-datasource> XML elements for MS SQL Server and Oracle.
    - ii. Change the content of the <jndi-name> element from `FNGCDDS` to `os1ds`.
    - iii. Set the content of the <connection-url> element to:

```
jdbc:db2://<JBoss_host_name>:<DB_port>/<dbname>
```

where *<dbname>* is the name of the database that will contain the object store data.
    - iv. Change the content of the <user-name> and <password> elements to that of the user and password of the tablespace you intend to use for your object store.
  - c. Edit **os1-xa-ds.xml** as follows:
    - i. Remove the <xa-datasource> XML elements for MS SQL Server and Oracle.
    - ii. Change the content of the <jndi-name> element from `FNGCDSXA` to `os1dsxa`.
    - iii. Set the content of the <xa-datasource-property name="ServerName"> element to the fully qualified domain name of the machine where DB2 is installed.
    - iv. Set the content of the <xa-datasource-property name="DatabaseName"> element to the name of the database that will contain the object store.
    - v. Set the content of the <xa-datasource-property name="PortNumber"> element to the port used by DB2.

- vi. Change the content of the <user-name> and <password> elements to that of the user and password of the tablespace to be used by your object store.
- d. Save your edits to **os1-ds.xml** and **os1-ds**.
2. Repeat [Step 1 on page 247](#) for each additional object store you intend to create (substituting different names for the XML files and XML elements).
3. (Optional) Encrypt the passwords for your XA and non-XA data sources, as shown in [“To encrypt data source passwords” on page 250](#).
4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

### To configure JBoss 4.0.x database connectivity (MS SQL Server)

---

1. To create data sources for an object store, do the following:
  - a. Copy files **FNGCD-ds.xml** and **FNGCD-xa-ds.xml** from directory **JBOSS\_DIST/server/server1/deploy** to new files, **os1-ds.xml** and **os1-xa-ds.xml**, respectively, in this directory.  
**NOTE** The name of each XML file you create for your data sources must end in **-ds.xml**.
  - b. Edit **os1-ds.xml** as follows:
    - i. Remove the <local-tx-datasource> XML elements for DB2 and Oracle.
    - ii. Change the content of the <jndi-name> element from `FNGCDDS` to `os1ds`.
    - iii. Set the content of <connection-url> element to:  

```
jdbc:sqlserver://<dbserver>:<db_port>;DatabaseName=<dbname>
```

where *<dbname>* is the name of the database that will contain the object store data.
    - iv. Change the content of the <user-name> and <password> elements to that of the user and password of the database you intend to use for your object store.
  - c. Edit **os1-xa-ds.xml** as follows:
    - i. Remove the <xa-datasource> XML elements for DB2 and Oracle.
    - ii. Change the content of the <jndi-name> element from `FNGCDSXA` to `os1dsxa`.
    - iii. Set the content of the <xa-datasource-property name="ServerName"> element to the fully qualified domain name of the machine where MS SQL Server is installed.
    - iv. Set the content of the <xa-datasource-property name="DatabaseName"> element to the name of the database that will contain the object store.
    - v. If MS SQL Server is not using the default port number (1433), add this XML element:  

```
<xa-datasource-property name="PortNumber">dbport</xa-datasource-property>
```

where *dbport* is the port number used by the MS SQL Server database instance.
    - vi. Change the content of the <user-name> and <password> elements to that of the user and password of the database to be used by your object store.



- d. Save your edits to **os1-ds.xml** and **os1-ds**.
2. Repeat [Step 1 on page 248](#) for each additional object store you intend to create (substituting different names for the XML files and XML elements).
3. (Optional) Encrypt the passwords for your XA and non-XA data sources, as shown in [“To encrypt data source passwords” on page 250](#).
4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

### To configure JBoss 4.0.x database connectivity (Oracle)

---

1. To create data sources for an object store, do the following:
  - a. Copy files **FNGCD-ds.xml** and **FNGCD-xa-ds.xml** from directory **JBOSS\_DIST/server/server1/deploy** to new files, **os1-ds.xml** and **os1-xa-ds.xml**, respectively, in this directory.
 

**NOTE** The name of each XML file you create for your data sources must end in **-ds.xml**.
  - b. Edit **os1-ds.xml** as follows:
    - i. Remove the `<local-tx-datasource>` XML elements for DB2 and MS SQL Server.
    - ii. Change the content of the `<jndi-name>` element from `FNGCDDS` to `os1ds`.
    - iii. Set the content of `<connection-url>` element to:
 

```
jdbc:oracle:thin:@<dbserver>:<db_port>:<Oracle_service_name>
```

 where *dbserver* is the name of the database machine and *port* is the port used by the database that will contain the GCD.
    - iv. Change the content of the `<user-name>` and `<password>` elements to that of the user and password of the tablespace you intend to use for your object store.
  - c. Edit **os1-xa-ds.xml** as follows:
    - i. Remove the `<xa-datasource>` XML elements for DB2 and MS SQL Server.
    - ii. Change the content of the `<jndi-name>` element from `FNGCDDSXA` to `os1dsxa`.
    - iii. Set the content of the `<xa-datasource-property name="URL">` element to:
 

```
jdbc:oracle:thin:@<dbserver>:<db_port>:<Oracle_service_name>
```

 where *dbserver* is the name of the database machine and *port* is the port used by the database that will contain the GCD.
    - iv. Change the content of the `<xa-datasource-property name="User">` and `<xa-datasource-property name="Password">` elements to that of the user and password of the tablespace to be used for the object store.
  - d. Save your edits to **os1-ds.xml** and **os1-ds**.
2. Repeat [Step 1 on page 249](#) for each additional object store you intend to create (substituting different names for the XML files and XML elements).

3. (Optional) Encrypt the passwords for your XA and non-XA data sources, as shown in [“To encrypt data source passwords” on page 250](#).
4. If you are upgrading Content Engine, go to [“To edit the upgrader utility file” on page 499](#); otherwise, continue at [“Prepare Storage Areas for Object Stores” on page 252](#).

### To encrypt data source passwords

---

Use the following procedure to generate an encrypted data source password.

#### NOTES

- Substitute your own value (at least eight characters) for the Salt attribute in place of the value `twsalt12` used in this procedure.
  - Choose an obscuring password in place of `Magic`, which is used in this procedure.
  - Choose a real LDAP password in place of `Secret`, which is used in this procedure.
  - Assume the path to the JBoss **conf** directory is `...\\server\\myserver\\conf`, where `myserver` is the name of the JBoss server.
1. Edit the file **jboss-service.xml** in the **conf** directory by adding the following text at the end of the Security section of the file, just before the Transactions section.

```
<mbean code="org.jboss.security.plugins.JaasSecurityDomain"
  name="jboss.security:service=JaasSecurityDomain, domain=ServerMasterPassword">
  <constructor>
    <arg type="java.lang.String" value="ServerMasterPassword"/>
  </constructor>
  <!-- The opaque master password file used to decrypt the encrypted database
  password key -->
  <attribute name="KeyStorePass">
    {CLASS}org.jboss.security.plugins.FilePassword:${jboss.server.home.dir}/
    conf/server.password
  </attribute>
  <attribute name="Salt">twsalt12</attribute>
  <attribute name="IterationCount">13</attribute>
</mbean>
```

2. Create the **server.password** file, as follows:

- a. Navigate to the **conf** directory:

```
cd %JBOSS_HOME%\server\myserver\conf
```

- b. If the **server.password** file already exists, run the following command to delete it:

```
del server.password
```

- c. Run the following command (with no carriage return) to create the **server.password** file:

```
java -cp ...\\lib\\jbossx.jar org.jboss.security.plugins.FilePassword twsalt12 13
"Magic" server.password
```

3. Use the obscuring password `Magic` to encrypt your real LDAP password `Secret`, as follows:

- a. Navigate to the **conf** directory:

```
cd %JBOSS_HOME%\server\myserver\conf
```

- b. Run the following command to encrypt your real LDAP password:

```
java -cp ..\lib\jbossx.jar org.jboss.security.plugins.PBEUtils twsalt12 13  
"Magic" "Secret"
```

- c. Note the command output, for example Encoded password: Fm/hKKlyXZj, which you will use in step 4.
4. Edit **login-config.xml** in the **conf** directory by pasting the encrypted password for LdapExtLoginModule, as shown below:

```
<module-option name="bindCredential">Fm/hKKlyXZj</module-option>  
<module-option  
  name="jaasSecurityDomain">jboss.security:service=JaasSecurityDomain,  
  domain=ServerMasterPassword  
</module-option>
```

## Task 15: Prepare Storage Areas for Object Stores

This task has two purposes:

- To prepare locations of initial storage areas for the object stores you will be creating, via the Create Object Store wizard ([“Create Object Stores” on page 258](#)).
- To prepare locations of additional file storage areas (see [“Create Additional File Storage Areas” on page 428](#)) for object stores you already created.

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

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

For information about file storage areas, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content storage > File storage areas](#).

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

For information about fixed storage areas, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content storage > Fixed storage areas](#).

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

### NOTES

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

By default, the Create Object Store wizard creates a database storage area. If your object stores will use database storage areas only, you can skip the rest of this task and continue at [“Create Object Stores” on page 258](#), provided that one of the following conditions is met:

- Your database type is non-DB2.
- Your database type is DB2 and either of the following conditions is met:
  - Your database storage areas will not contain large content elements (larger than 300 MB, for example).
  - You have set the value of the custom property `fullyMaterializeLobData` to `false` for the data sources associated with the database (see [“Configure Content Engine Application Server Database Connectivity \(WebSphere 6.0.x\)” on page 206](#)).

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

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

If the content of all your object stores will be in fixed storage areas only, create your fixed content devices now, skip the rest of this topic, and continue at [“Create Object Stores” on page 258](#).

To create a fixed content device, refer to the the procedures in IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content storage > Fixed storage areas](#).

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

## Configure File Servers for File Storage Areas

In this section you will configure file servers for the initial file storage areas of the object stores to be created, and for additional file storage areas of existing object stores.

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for currently supported operating systems for file servers. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

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

1. Create or designate an existing top-level directory on the file server where file storage areas will reside.
2. Secure the directory so only Content Engine Server and Content Search Engine can access it.
3. Expose the directory via the remote file access protocol that applies to the operating system of the file server.
4. (Recommended) Under the top-level directory, create a subdirectory for each file storage area you intend to create.

If you decide to put a file storage area directly within a top-level directory, rather than in a subdirectory, and you later decide to create an additional file storage area on this file server, you will have to create another top-level directory for it, as you will not be able to use the previously created top-level directory.

## Remote File Access Protocols

The supported remote file access protocols between Content Engine and a file server are as follows:

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

- Distributed File System (DFS)

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

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

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

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

### Users and Groups

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

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

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

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

Abbreviation	Meaning
<CE_OS_User>	Content Engine operating system user
<CE_OS_Group>	Content Engine operating system group
<fsa1>	Directory where content will be stored

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

### To configure a UNIX-based file server

---

1. Log on to the UNIX file server as a user with read/write access to the device where you want to create a storage area.

2. Create or designate a directory <fsa1> where content will be stored. For example:

```
$ mkdir /opt/ibmnet/file_stores/<fsa1>
```

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

```
chown <CE_OS_User>:<CE_OS_Group> <fsa1>
```

**NOTE** The UID (user ID) for <CE\_OS\_User> and the GID (group ID) for <CE\_OS\_Group> on the file server must match the UID and GID for the same user and group on the machine where Content Engine Server and Content Search Engine are running. This will normally be true if all machines use the same directory service, but they may be different.

4. Change the permissions on <fsa1> so that <CE\_OS\_User> and <CE\_OS\_Group> both have read/write/execute privileges and all other users have no privileges:

```
chmod 0770 <fsa1>
```

5. Via NFS, export <fsa1>. Alternatively, if the file server will host more than one file storage area, export the parent directory.

In the latter case, for example, export **/opt/ibmnet/file\_stores**, rather than **/opt/ibmnet/file\_stores/<fsa1>**, and then create a separate subdirectory to serve as the root of each file storage area.

**NOTE** IBM recommends that trusted hosts be restricted to just those on which an instance of Content Engine Server or Content Search Engine is executing. Root access should also be restricted. Refer to the UNIX administrator manual for details on exporting files in NFS.

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

---

1. Log on to the Windows file server as <CE\_OS\_User>.

2. Create (or designate) a directory *<fsa1>* where content will be stored. For example:
 

```
C:\> md c:\filenet\file_stores\<fsa1>
```
3. Navigate in Windows Explorer to *<fsa1>*, right-click the file icon, and choose Properties.
4. In the Security tab, click **Advanced**.
5. In the Advanced Security Settings dialog box,
  - a. Grant Full Control to *<CE\_OS\_User>* and *<CE\_OS\_Group>*, and select *This Folder, subfolders, and files* from the *Apply onto* drop-down list.
  - b. Remove all other users and groups in the *Permission entries* table.
  - c. Click **OK**.
6. In the Sharing tab, do the following:
  - a. Click **Share this folder** and click **Permissions**.
  - b. Grant Full control to *<CE\_OS\_User>* and *<CE\_OS\_Group>*.
  - c. Remove all other users and groups in the *Permission entries* table.
  - d. Click **OK**.

---

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

1. Do all the steps in [“To configure a Windows-based file server for a Windows client using CIFS” on page 255](#).
2. Use the procedures in Microsoft documentation to configure Windows Services for NFS to expose *<fsa1>*.

#### NOTES

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

### Configure the Remote Access Protocol on the Client Machine

When configuring the remote file access protocol (NFS or CIFS), the “client machine” is the one where Content Engine Server and/or Content Search Engine are running.

Configuring the remote access protocol (NFS or CIFS) means designating a directory (where content is be stored) so that it appears to be on the a local file system of the client machine.



---

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

---

1. Log on to the Content Engine machine as the user who launched the application server where Content Engine has been deployed.
2. Mount the exported NFS file system (from [Step 5 of “To configure a UNIX-based file server” on page 255](#)) onto a local directory on the Content Engine machine. The mount point must be in the same location on all machines where Content Engine Server and Content Search Server are installed the local file system. For example,

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

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

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

3. If you haven't already done so (in [“Install and Deploy Content Engine” on page 160](#)), set the default file-creation permissions mask for the JVM instance that will host Content Engine Server so that the owner (the user running JVM) and the members of the owner's group have read/write/execute access permissions, and all others have no access, as follows:
  - a. Stop Content Engine Server.
  - b. Use the UNIX utility program *umask* to set the default file-creation permissions mask for the JVM instance:

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

This ensures that the access permissions on files and directories created by Content Engine Server are identical to those you specified for the root directory in [Step 4 of “To configure a UNIX-based file server” on page 255](#).

- c. Start Content Engine Server.

---

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

---

If both Content Engine Server and the file server are in the same Windows domain, no action is required.

## Task 16: Create Object Stores

### CAUTION

- Do the procedure in this section *only if* you are installing version 4.0.0 of Content Engine as a new application. If you are upgrading from version 3.5.x of Content Engine, skip to [“Verify the Content Engine Installation” on page 263](#).
- Before creating object stores, be sure you have installed the latest Content Engine service pack, as indicated in [“Install Content Engine Software Updates” on page 177](#).

In this task you will create object stores. If you haven't already done so, set up the initial storage areas of the object stores (see [“Prepare Storage Areas for Object Stores” on page 252](#)).

### NOTES

#### Object Store Notes

- To create an object store, you need to know the following:
  - Display name for new object store
  - Object store description (optional)
  - JNDI names associated with the data sources for your database
  - Database engine (DB2, SQL Server, or Oracle)
  - (SQL Server) Type of authentication: Windows Authentication or Database Engine Authentication
  - (Oracle) User name and password for the tablespace corresponding to the object store
  - Database alias
    - (DB2) Tablespace name
    - (SQL Server) Database name
    - (Oracle) Net service name
  - Initial administrative and user groups that can access the object store

#### Deployment Notes

- You must create at least one object store in your FileNet P8 domain.
- The name you assign to an object store has the following characteristics:
  - Can be no more than 64 alphanumeric characters.
  - Can be different from the name of the associated database (SQL Server) or tablespace (DB2 or Oracle).
  - Must be different from any other object store name within the FileNet P8 domain.
  - There may be other restrictions imposed by the underlying database programs. Consult your database documentation or administrator for any restrictions.

### Authentication Notes

- Each object store can have a different set of default users and groups. IBM FileNet P8 does not support distribution groups.
- Once an object store is in production and contains many objects, you must use the Security Script wizard of Enterprise Manager to add new users and groups so that they have access to those existing objects. Therefore, you should be careful to add all the default users and groups to the object store before putting it into production.

### Database Notes

- All the object stores you create for Content Engine must be based on the same database type. For example, you cannot have one object store based on Oracle and another based on SQL Server.
- You cannot create an object store with an Oracle tablespace user name that contains a space or has mixed-case characters.
- Before creating an object store, you must create the following:
  - an associated DB2 tablespace, Oracle tablespace, or SQL Server database.
  - a pair of connection pools—one for distributed transactions (XA) and one for non-distributed (non-XA) transactions.
  - a pair of data sources—one for XA transactions and one for non-XA transactions.
- The Create Object Store wizard will fail if you try to assign a new object store to a database/tablespace that is not completely empty.
- Once an object store is created, you can refine its definition and add content to it.

### To create an object store

---

1. Log on as the Content Engine operating system user to a machine where Enterprise Manager is installed.
2. Launch the Enterprise Manager (double-click the desktop shortcut) and log on.  
**NOTE** You will be required to log on when the Enterprise Manager is started. (Windows Integrated authentication will not require a password.)
3. In the FileNet P8 Logon screen, select the FileNet P8 domain in which you will create an object store, and then click **Connect**.
4. In the FileNet P8 Domain Logon screen, log on as a GCD administrator, and then click **OK** to start Enterprise Manager.
5. In the tree view, right-click the **Object Stores** container and choose **New Object Store**.
6. Work through the initial Create Object Store wizard screens as follows:

In this screen...	Perform this action...
Welcome	Click <b>Next</b> to proceed with creating an object store.

In this screen...	Perform this action...
Name and Describe the Object Store	<p>Type the display name for the object store. The symbolic name and description are entered automatically as you type the display name, but you can edit them before leaving this screen. Click <b>Next</b>.</p> <p><b>NOTE</b> The symbolic name, used for internal programmatic purposes, must contain only ASCII characters and must begin with an alphabetic character.</p>
Specify the Data Sources	<p>Specify the JNDI names associated with the data sources for your database, as follows:</p> <ul style="list-style-type: none"> <li>• JNDI Data Source Name: The data source that binds the non-XA connection pool to the database/tablespace.</li> <li>• JNDIXA Data Source Name: The data source that binds the XA connection pool to the database/tablespace.</li> </ul> <p>Click <b>Next</b>.</p>
Specify the default Content Store	<p>Click one of the following to specify the default location for the content of the object store, and then click <b>Next</b>. If you clicked</p> <ul style="list-style-type: none"> <li>• Database Storage Area – continue at <a href="#">“Specify object store administrators” on page 261</a></li> <li>• File Storage Area – continue at <a href="#">“Specify File Storage Area Directory” on page 260</a></li> <li>• Fixed Storage Area – continue at <a href="#">“Fixed Storage Area” on page 260</a></li> </ul>
Specify File Storage Area Directory	<p>Type, or browse (Windows only) to, the pathname of the directory to contain the file storage area</p> <p>Specify the UNC of the network share to contain the file storage area.</p> <p><b>NOTE</b> Refer to <a href="#">“Users and Groups” on page 254</a> for information on the permissions required on the directory where your file storage area will be created.</p> <p>Click <b>Next</b> and continue at <a href="#">“Specify object store administrators” on page 261</a>.</p>
Fixed Storage Area	<p>If you specified Fixed Storage Area as your default Content Store, choose a fixed content device from the drop-down list, and specify the staging area path. Click <b>Next</b> and continue at <a href="#">“Specify object store administrators” on page 261</a>.</p>

In this screen...	Perform this action...
Specify object store administrators	<p>An object store administrator can log on to Enterprise Manager, has administrative (add/delete/change) access to the object store, and is in security lists on objects. The initial object-store administrators list includes the account that creates the object store.</p> <p>To add object store administrators to the list, do the following:</p> <ol style="list-style-type: none"> <li>Click <b>Add...</b></li> <li>In the Select Users and Groups screen, specify object type, realm, and search criteria as needed, and then click <b>Find</b> to display a list from which to select object store administrators.</li> <li>Select a single item, or press and hold the CTRL (or SHIFT) key, and use the left mouse button to select multiple items (or the first and last of a range of items).</li> <li>Release the CTRL (or SHIFT) key and click <b>OK</b>.</li> </ol> <p>To remove object store administrators from the list, do the following:</p> <ol style="list-style-type: none"> <li>Select a single item, or press and hold the CTRL (or SHIFT) key, and use the left mouse button to select multiple items (or the first and last of a range of items).</li> <li>Release the CTRL (or SHIFT) key and click <b>OK</b>.</li> </ol> <p>When your list is complete, click <b>Next</b> and continue at <a href="#">“Specify Initial User Groups” on page 262</a>.</p> <p><b>NOTE</b> If you specify an empty list, the wizard automatically adds #AUTHENTICATED-USERS, which gives all network users in the authentication realm administrative access to the object store (for example, under Windows authentication, all accounts in Domain Users).</p>

In this screen...	Perform this action...
Specify Initial User Groups	<p>User groups have non-administrative (browse directories and read documents) access to the object store. Specify the initial list of groups, as follows:</p> <p>To add user groups to the list, do the following:</p> <ol style="list-style-type: none"> <li>Click <b>Add....</b></li> <li>In the Select Users and Groups screen, specify object type, realm, and search criteria as needed, and then click <b>Find</b> to display a list from which to select user groups.</li> <li>Select a single item, or press and hold the CTRL (or SHIFT) key, and use the left mouse button to select multiple items (or the first and last of a range of items).</li> <li>Release the CTRL (or SHIFT) key and click <b>OK</b>.</li> </ol> <p>To remove user groups from the list, do the following:</p> <ol style="list-style-type: none"> <li>Select a single item, or press and hold the CTRL (or SHIFT) key, and use the left mouse button to select multiple items (or the first and last of a range of items).</li> <li>Release the CTRL (or SHIFT) key and click <b>OK</b>.</li> </ol> <p>When your list is complete, click <b>Next</b>.</p> <p><b>NOTE</b> If you specify an empty list, the wizard automatically adds #AUTHENTICATED-USERS, which gives non-administrative access to all network users in the authentication realm.</p>
Completing the Create an Object Store Wizard	<p>Review your selections and click <b>Finish</b> to create the object store. When the status displayed in the Object Store Create Status window shows that the object store has been successfully created, click <b>OK</b>.</p>

**NOTE** If, as an object store administrator, you need to add more users or groups to the object store, refer to the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to... > Update object store with new users and groups](#).

## Task 17: Verify the Content Engine Installation

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

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

**CAUTION** Do the procedure in this section *only if* you are installing version 4.0.0 of Content Engine as a new application. If you are upgrading from version 3.5.x of Content Engine, continue at [“Upgrade Content Engine Data” on page 498](#).

### To verify the Content Engine installation

---

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

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

- d. Click **Check In**.
- e. Right-click the document object you just checked in and click the Versions tab of the Properties dialog box.
- f. You should now see the Major Version field change once the document is checked in.

#### **To enable Enterprise Manager to display file storage area status**

---

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

To enable a Windows logon user to see an online status for a file storage area and to add content to it, you must add the user to the directory security of the file storage area with the following permissions: Modify, Read & Execute, List Folder Contents, Read, and Write. Alternatively, as a member of the Local Administrators group you can add the Windows logon user to a group account, as explained in the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Storage area security](#).



## Task 18: Install Content Search Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Content Search Engine.

### **To install the Content Search Engine software updates**

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. Content Search Engine 4.0.1 Service Pack
  - b. Any subsequent fix pack (P8CSE-4.0.1-001 or later)
  - c. Any subsequent interim fixes (typically optional)

## Task 19a: Install Process Engine (Windows)

### NOTES

- Ensure that all applicable tasks listed in the [“Installation Planning Considerations”](#) on page 24 have been completed before beginning this task.
- If Process Engine will be installed by a domain user rather than a local user, see [“Specify IBM FileNet P8 Accounts”](#) on page 60 for details on creating required users and groups.
- Be sure to have the correct database information. See [“Database Considerations”](#) on page 32 for the specific requirements. Entering incorrect information during Process Engine installation could cause the installation to fail. Recovery could mean uninstalling and re-installing the Process Engine software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to install and configure Process Engine but can be re-enabled after the installation is complete. See [“To turn off Oracle Password Complexity Verification”](#) on page 269 and [“To re-enable Oracle Password Complexity Verification”](#) on page 283 for details.
- Determine whether Process Engine will be installed by running the setup program interactively or in silent mode. If this will be a silent installation, complete all steps in this task, up to [“To install the Process Engine software interactively”](#) on page 269. After completing silent installation steps, return to this task, at [“To Complete Additional Configuration”](#) on page 281 and proceed through the remainder of this task.
- Determine when to execute pre-installation SQL scripts for SQL Server and Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle or the sa password for SQL Server.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts”](#) on page 647 for detailed information on the scripts and modes of execution.

### To verify the database connection (Oracle)

---

The database that Process Engine will use must be running and fully configured for IBM FileNet use prior to starting Process Engine Setup. If the database will be remote, you must also install the database client software on the Process Engine server and make sure it is running before you start Process Engine Setup. See the [“Verify that Oracle Server Is Installed for IBM FileNet P8”](#) on page 85 task for installation and configuration procedures.

Take the following steps to verify that the Oracle database instance used by Process Engine is accessible. This will verify that you can connect to the Oracle database in the same way Process

**To create the Process Engine ODBC data source and test the connection(SQL Server only)**

Engine Setup will. How you log on to sqlplus to begin this verification will vary, depending upon how you will choose to execute the SQL scripts. Make whatever corrections are necessary before proceeding with the Process Engine installation.

1. Execute the following at a command prompt:

```
sqlplus
```

2. Enter a command at the SQLPlus prompt, as follows:

- a. If the SQL scripts will be run from Process Engine Setup by prompting for the sys password, type the following command:

```
sys/<password> as sysdba
```

- b. If the SQL scripts will be run from Process Engine Setup 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
9.2.0.7.0
```

In this example:

- the instance, ORACLE\_SID is p8dbshr
- the database server name is hqvwbuCS
- the Oracle server is at the 9.2.0.7 version

**To create the Process Engine ODBC data source and test the connection(SQL Server only)**

The ODBC data source is required for both local and remote databases.

1. Start **Program > Administrator Tools > Data Source (ODBC)**.
2. Click **Add** on the **System DSN** tab.
3. Select **SQL Server** as the driver to use for the new data source and click **Finish**.

4. Enter a name and description for the data source. The name will be required input for the Process Engine Setup program when configuring for a SQL Server database.
5. Choose the SQL Server to connect to from the dropdown list of servers and click **Next**.  
**NOTE** If only a server name appears in the list, the connection will be with the default instance. If there are named instances in the database, the name will appear as *<server>/<instance name>*.
6. Do the following, and then click **Next**:
  - a. Choose SQL Server authentication.
  - b. Select the option to get default settings for additional configuration options by connecting to the SQL Server.
  - c. Indicate the Login ID and Password to connect to the database.  
**NOTE** This database login ID information needd not be for an administrator and it is only used to connect to the database to get the default values for the remaining settings required to configure the data source.
7. Change the default database to be the Process Engine database created earlier in [“Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81](#).
8. Turn off **Use ANSI null, paddings, and warning** and turn on **Use ANSI quoted identifiers**.
9. Turn on **Perform translation for character data** and click **Finish**.
10. Verify the settings for the data source configuration and click **Test Data Source**. If the test is successful click **OK**. Otherwise resolve the problem before continuing.
11. Double-click **SQL Server** on the **Connection Pooling** tab.
12. Select **Don't pool connection to this driver** and click **OK**.
13. Click **OK** on the ODBC Data Source Administrator window to finish configuration of the data source.

On the summary screen click **Test Data Source**. If error messages display, resolve them before proceeding with Process Engine software installation.

### To verify the database connection (DB2)

---

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

1. At a command prompt, enter:  

```
db2cmd
```
2. At the DOS prompt, start the DB2 Command Line Processor by typing:  

```
db2
```
3. At the DB2 prompt, enter the following command:  

```
connect to <database alias> user <f_sw> using <f_sw password>
```

where `<f_sw>` is either the default `f_sw` user or the assigned alias

DB2 will display the database connection information.

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

```
db2 => connect to pedbwin user f_sw using f_sw
```

```
Database Connection Information
```

```
Database server      = DB2/AIX64 8.2.0
```

```
SQL authorization ID = F_SW
```

```
Local database alias = PEDBWIN
```

In this example, the database alias is `pedbwin` and the `f_sw` user password is `f_sw`.

---

### To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the installation process. Turn off this Oracle feature and do not re-enable it until you have installed Process Engine and used the `set_f_maint_pw` utility to change the `f_sw` (or its alias).

See [“To set the `f\_maint` and `f\_sw` passwords \(Oracle and SQL Server\)” on page 282](#) for procedures to change the `f_sw` password.

---

### To install the Process Engine software interactively

---

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

1. Log on as a member of the local Administrators group or a user with equivalent permissions. The user you log on as must also be a database administrator. See [“Specify IBM FileNet P8 Accounts” on page 60](#) for information on requirements for logging on as a Windows domain user for Process Engine installation.

**NOTE** This user does not need to be a database administrator unless you will be executing the SQL scripts from the installer.

**WARNING** Setup will run a number of steps but to complete the installation the computer must be restarted. If you are installing on a Windows 2003 R2 server, you must create an empty file named `fn_newinit` in the `c:\temp` directory before restarting the computer.

2. Access the Process Engine software package, and start the **P8PE-4.0.0-Win.exe** Setup program.

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

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

```
\\server08\Software\InstallationDisks\FileNet\Release P8 4.0.0\ProcessEngine\Windows\IMSInst
```

When expressed in 8.3 format might be:

**\\server08\Software\INSTAL~1\FileNet\RELEAS~1.0\PROCES~1\Windows\MSInst**

This is 73 characters long.

3. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement.
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation is installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation web site. You can use multi-part root folders (for example, <b>/docs/ecm_help</b>) if your application server supports them.</p>
Specify Installation Location for Common Files	Choose the destination directory for configuration files that will be shared with other IBM FileNet P8 components. Accept the default location or click <b>Browse</b> to change the location.
Specify Installation Location for Program Files	Indicate the installation location for the executable files. This is the drive where the <b>\FNSW</b> directory will be created. Select a local drive in a cluster configuration.
Specify Installation Location for Data Files	Indicate the installation location for the configuration and data files. This is the drive where the <b>\FNSW_LOC</b> directory will be created. Select a shared drive in a cluster configuration.

In this screen...	Perform this action...
Specify Network Clearinghouse Domain Name	<p>Enter <i>&lt;domain name&gt;:&lt;organization&gt;</i>, where:</p> <ul style="list-style-type: none"> <li>The maximum length of your <i>&lt;domain name&gt;</i> entry does not exceed 19 characters.</li> <li>The maximum length of your <i>&lt;organization&gt;</i> entry does not exceed 19 characters.</li> <li>Both your <i>&lt;domain name&gt;</i> and <i>&lt;organization&gt;</i> entries contain only alphanumeric characters and underscores.</li> </ul> <p>A typical convention is to enter <i>&lt;your PE machine name&gt;:&lt;your company name&gt;</i>. If the machine name or company name include hyphens, replace them with underscores in your entry.</p>
Specify the Database Location	Indicate whether the database will be local or remote.
Specify the Database Type	<p>Indicate whether the database will be:</p> <ul style="list-style-type: none"> <li>Oracle</li> <li>SQL Server</li> <li>DB2</li> </ul>

4. Depending on the type and location of the database you selected above, continue at one of the following procedures:
- [“To complete local or remote SQL Server database screens” on page 272](#)
  - [“To complete remote Oracle database screens” on page 273](#)
  - [“To complete local Oracle database screens” on page 275.](#)
  - [“To complete remote DB2 database screens” on page 277](#)

**NOTE** Only remote DB2 databases are supported.

**To complete local or remote SQL Server database screens**

---

Complete the Process Engine Setup screens that are specific to local or remote SQL Server databases, as follows:

In this screen...	Perform this action...
Specify Execution Mode for SQL Server Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the SQL Server Administrator “sa” user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually</li> <li>• I want to run the SQL scripts with a prompted password</li> <li>• I want Setup to run the SQL scripts silently using operating system authentication</li> </ul> <p>Validation of the SQL Server connection will also occur if you choose to run the scripts now, either with a prompted password or using operating system authentication. No validation of the database connection will be done if you indicate that you have already run the SQL scripts manually.</p>
Specify the SQL Server System Administrator Password  <b>NOTE</b> This prompt only appears if you indicated that you want to run the SQL scripts with a prompted password.	Enter the SQL Server sa password.



In this screen...	Perform this action...
Specify SQL Server Configuration Parameters	<p>Enter the following information:</p> <ul style="list-style-type: none"> <li>• ODBC data source name</li> <li>• Database name</li> <li>• Filegroup name</li> </ul> <p><b>NOTE</b> You might see a slight delay before seeing the next screen in Setup if you chose to run the SQL Scripts from the setup program. Validation of the database connection will be done after clicking <b>Next</b> on this screen. If validation is successful you will see the next screen. If validation fails, an error will be returned and you must resolve the problem before proceeding. No validation of the database connection will be done if you indicate that you have already run the SQL scripts manually.</p>
Specify SQL Server Version	Indicate whether you will be using SQL Server 2000 or SQL Server 2005 for the Process Engine database.

**To complete remote Oracle database screens**

Complete the Process Engine Setup screens that are specific to remote Oracle databases, as follows:

In this screen...	Perform this action...
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually</li> <li>• I want to run the SQL scripts with a prompted password</li> <li>• I want Setup to run the SQL scripts silently using operating system authentication</li> </ul>

In this screen...	Perform this action...
Specify the Oracle SYS Password  <b>NOTE</b> This prompt only appears if you indicated that you want to run the scripts with a prompted password.	Enter the Oracle SYS password.
Specify Oracle Configuration Parameters	Enter the appropriate values for the Oracle Home Directory.  The Oracle Home path you enter refers to the local Oracle installation directory.  To find the Oracle Home directory, open the command prompt and type <code>set</code> . Look for the Oracle_Home variable.
Specify Remote Oracle Configuration Parameters	Enter the following Oracle database configuration parameters. <ul style="list-style-type: none"> <li>• Global Database Name (as identified in the <b>tnsnames.ora</b> file)</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.  All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85</a> .
Specify Oracle Version	Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.  <b>NOTE</b> The Oracle versions must be the same on the client and server.

**To complete local Oracle database screens**

---

Complete the Process Engine Setup screens that are specific to local Oracle databases, as follows:

In this screen...	Perform this action...
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts need to be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually</li> <li>• I want to run the SQL scripts with a prompted password</li> <li>• I want Setup to run the SQL scripts silently using operating system authentication</li> </ul>
Specify the Oracle SYS Password	<p>Enter the Oracle SYS password.</p> <p><b>NOTE</b> This prompt only appears if you indicated that you want to run the SQL scripts with a prompted password.</p>
Specify Oracle Configuration Parameters	<p>Enter the appropriate values for the Oracle Home Directory.</p> <p>The Oracle Home path you enter refers to the local Oracle installation directory.</p> <p>To find the Oracle Home directory, open the a command prompt and type <code>set</code>. Look for the Oracle_Home variable.</p>

In this screen...	Perform this action...
Specify Local Oracle Configuration Parameters	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Oracle SID            To find the Oracle SID, open a cmd prompt and type <code>set</code>. Look for the Oracle_SID variable</li> <li>• Temporary Tablespace Name            Default value is VWTEMP_TS</li> <li>• Data Tablespace Name            Default value is VWDATA_TS</li> <li>• Index Tablespace Name            Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85</a>.</p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>

**To complete remote DB2 database screens**

In this screen...	Perform this action...
Database Alias Name	Enter the database alias name as assigned in <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> .
Data Tablespace name	defaults to VWDATA_TS
DB2 Passwords	Enter the password for the Operating System user <b>f_sw</b> or assigned alias.  Enter the password set for this user in either <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a> if database authentication is SERVER or SERVER_ENCRYPT, or <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> if authentication is CLIENT.
DB2 Passwords Although the screen title doesn't change the question in the screen does.	Enter the password for the Operating System user <b>f_maint</b> or assigned alias.  Enter the password set for this user in either <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a> if database authentication is SERVER or SERVER_ENCRYPT, or <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> if authentication is CLIENT.

**To complete final Process Engine Setup screens**

Complete the screens to finalize the Process Engine installation, as follows:

In this screen...	Perform this action...
Determine Administrative User and Group Aliasing Method	Determine whether FileNet default operating system and database users and groups will be used, or if you want to define aliases for these users and groups. <ul style="list-style-type: none"> <li>• Yes, configure aliases</li> <li>• No, use actual account names for aliases</li> </ul>

In this screen...	Perform this action...
<p>Specify Administrative User and Group Aliases</p> <p><b>NOTE</b> This screen is only presented if you chose to configure aliases.</p>	<p>Indicate the alias you want to create for each of the following users and groups.</p> <ul style="list-style-type: none"> <li>• FNADMIN OS group</li> <li>• FNUSR OS group</li> <li>• FNOP OS group</li> <li>• FNSW OS user</li> <li>• f_sw DB user</li> <li>• f_maint DB user</li> </ul> <p><b>NOTE</b> If you ran SQL scripts manually before starting Process Engine Setup, the aliases here for f_sw and f_maint must match the users specified as runtime and maintenance users.</p> <p>If you are logged on as a domain user for this installation the fnadmin, fnop and fnuser groups and the fnsw user, or aliases for them have already been defined on this server. Assign the same alias here as was defined earlier.</p> <p>The f_sw and f_maint users should be dedicated users for IBM FileNet use.</p> <p>The default password for the fnsw user (or its alias) will be set and must not be changed until Process Engine installation is complete. <a href="#">“To Complete Additional Configuration” on page 281</a> for information on changing the fnsw password and associated changes to Windows services.</p>
<p>Choose an Application Server</p>	<p>Select an application server and version from the drop-down list. You must use the same application server type and version as Content Engine.</p>

In this screen...	Perform this action...
Content Engine API Configuration	<p>Configure the Content Engine API, as follows:</p> <ul style="list-style-type: none"> <li>Transport Method Select WSI from the drop-down list.</li> <li>Content Engine Client Software URL: Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example: for WebLogic <code>http://hqcomp1:7001/wsi/FNCEWS40DIME</code> for WebSphere <code>http://hqcomp2:9080/wsi/FNCEWS40DIME</code> for JBoss: <code>http://hqcomp3:8080/wsi/FNCEWS40DIME</code> Do not modify the remainder of the string from the default values. <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.</li> </ul>
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.
Completing the Setup wizard	Click <b>Finish</b> to complete the Process Engine installation.

- Setup will run a number of steps but to complete the installation the computer must be restarted. If you are installing on a Windows 2003 R2 server, you must create an empty file named **fn\_newinit** in the **c:\temp** directory before restarting the computer. When prompted, select **Yes, restart my computer** and click **Finish**.
- When the system restarts, log on using the same account you used in step 1. After you log on, Process Engine Setup will continue. Click **Next** to continue the installation.

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

```
setup.exe -postboot
```

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

7. When the dialog box informs you that Process Engine has been successfully installed, click **Finish**.
8. Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE400_setup.log
IS mini-installer logs	<Windir>/mini_installer.log, Windows Event logs, and log files under <b>IFNSW_LOC\logs</b>
Output files from SQL Script validation (only if the SQL scripts were executed from Process Engine Setup).	C:\Program Files\FileNet\PE

9. Proceed to [“To Complete Additional Configuration” on page 281](#).

### To Install the Process Engine software silently

---

Complete all preceding tasks in this topic, up to [“To install the Process Engine software interactively” on page 269](#). Then take the following steps to silently install Process Engine on a Windows machine.

1. Access the Process Engine software package, and copy to a temporary directory on the local disk.
2. Setup will run a number of steps but to complete the installation the computer must be restarted. If you are installing on a Windows 2003 R2 server, you must create an empty file named **fn\_newinit** in the **c:\temp** directory before restarting the computer.
3. Edit the **PE\_silent\_input.txt** file to reflect the appropriate responses for your installation. See [“Encrypt Passwords for Silent Installations and Upgrades” on page 642](#) for information on use of the password encryption tool.
4. Save the edited response file to your temporary directory.
5. Log on as a member of the local Administrators group or a user with equivalent permissions. The user you log on as must also be a database administrator. See [“Specify IBM FileNet P8 Accounts” on page 60](#) for information on requirements for logging on as a Windows domain user for Process Engine installation.



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

```
P8PE-4.0.0-Win.exe -silent -options PE_silent_input.txt
```

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

7. Proceed to ["To Complete Additional Configuration" on page 281.](#)

### **To Complete Additional Configuration**

Process Engine Setup automatically creates several internally required local users. These users have either the default names or the aliases you defined while running Setup. Because these users are created with default passwords, IBM strongly recommends you reset their passwords to maintain system security on each Process Engine.

User Name	User Type	Description	How to modify
f_maint	Database (Oracle, SQL)	Has DBA privileges. Used for RDBMS maintenance. Also referred to as the database maintenance user.	Execute set_f_maint_pw. See steps below.
f_sw	Database (Oracle, SQL)	Has privileges only for database objects created by Process Engine. Used for RDBMS maintenance. Also referred to as the database runtime user.	Execute set_f_maint_pw. See steps below.

User Name	User Type	Description	How to modify
fnswh or alias	Operating System	Primary Process Engine user. Used to execute Process Engine software and services.	Execute <b>Windows Control Panel &gt; Administrative Tools &gt; Computer Management &gt; System Tools &gt; Local Users and Groups</b> .  After you change the password for the fnswh user, you must also use the Windows Services tool to update the Log On tab for the IMS ControlService because the fnswh user is used to start that service. Process Engine Services Manager is started as the Local System account accordingly.  <b>NOTE</b> Process Engine activity will cease if the fnswh user's password expires.

**To set the f\_maint and f\_sw passwords (Oracle and SQL Server)**

Use the following procedure to change the passwords for the f\_maint and f\_sw users after installing the Process Engine software. These procedures will also verify the ability to connect to the Process Engine database.

**NOTE** You must also use the set\_f\_maint\_pw utility to change the f\_maint and f\_sw passwords later, such as when the passwords are due to expire. Changing these passwords any other way will cause the Process Engine software to fail.

1. Log on as the fnswh (or its alias).
2. At a command prompt, enter:

```
set_f_maint_pw <user_name>
```

where the <user\_name> is f\_maint, f\_sw or the aliases defined for them in Process Engine Setup.

3. Press **Return** at the prompt for the current password.

**NOTE** Pressing Return will only work the first time the password is changed, when the default password is still valid. After the password is changed from the default, this prompt requires you to enter the current password.

4. Enter a new password for the user when prompted. The password must be between 6 and 30 characters. If the database is SQL Server, the first character cannot be numeric, per SQL Server naming conventions.
5. Confirm the new password when prompted.

### **To re-enable Oracle Password Complexity Verification**

---

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

### **To verify the /etc/services file settings**

---

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

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

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

### **To verify TCP/IP parameter settings**

---

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

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

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

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

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

### **To redirect log messages to the Image Services error log**

---

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

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

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

(Optional; Oracle database only) To remove fnsw and oracle users from the ORA\_DBA group

**(Optional; Oracle database only) To remove fnsw and oracle users from the ORA\_DBA group**

Process Engine Setup creates several users that are no longer required after installation is complete. Remove the fnsw (or its alias) and oracle users from the ORA\_DBA group.

Proceed to [“To install the Process Engine software updates” on page 338](#).

## Task 19b: Install Process Engine (Solaris)

### NOTES

- Before beginning this task, ensure that all applicable tasks listed in the [“Installation Planning Considerations” on page 24](#) have been completed.
- Process Engine requires the presence of several partitions. Before installing PE verify that your Operating System is set up with a correctly configured volume manager.
- You will find references to logging on as the root and fnsw users within the following procedures.
  - The root user must run in the Bourne or Korn shell.
  - The root user who will be doing a silent installation must run in the Korn shell.
  - The fnsw user must run in the Korn shell.
- Before starting Process Engine installation, be sure to have the correct database information including server names, database and instance names, and tablespace names. Entering incorrect information during installation could cause the installation to fail. Recovery could mean uninstalling and re-installing the Process Engine software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to install and configure Process Engine but can be re-enabled after the installation is complete. See [“To turn off Oracle Password Complexity Verification” on page 291](#) and [“To re-enable Oracle Password Complexity Verification” on page 301](#) for details.
- Complete all steps in this task, up to [“To install the Process Engine software interactively” on page 291](#) if Process Engine software will be installed silently. After completing silent installation steps, return to this task at [“To reset administrative user passwords” on page 300](#) and proceed through the remainder of this task.
- Determine when to execute pre-installation SQL scripts for Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle in an xterm window.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts” on page 647](#) for detailed information on the scripts and modes of execution.

### To increase the operating system kernel limits

---

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

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

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

```
vi /etc/system
```

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

**NOTE** The `shminfo_shmmin` value must be set to 0 or the Process Engine installation will fail. While you can change this value after installation, IBM recommends setting this value no higher than 1.

```
set semsys:seminfo_semmap=50
set semsys:seminfo_semmni=2000
set semsys:seminfo_semmns=2000
set semsys:seminfo_semmnu=500
set semsys:seminfo_semmnl=512
set semsys:seminfo_semopm=256
set semsys:seminfo_semume=500
set semsys:seminfo_sevmx=32767
set semsys:seminfo_semaem=16384
set shmsys:shminfo_shmmax=20971520
set shmsys:shminfo_shmmin=0
set shmsys:shminfo_shmmni=2000
set shmsys:shminfo_shmseg=100
set msgsys:msginfo_msgmni=2048
set max_nprocs=1000
set fnsod:sod_Debug=0
set rlim_fd_max=1024
set rlim_fd_cur=256
```

4. Save your changes.
5. Reboot the server.

### To create or modify internal Process Engine users and groups

---

Using Solaris' `admintool`, you must create or modify specific local groups and users required internally by Process Engine prior to running Process Engine Setup. Optionally, aliases for these users and groups can be defined. Process Engine Setup will ask whether to use the default users and groups or identify the aliases that have been defined.

**NOTE** Throughout the remainder of this document, references will be made to the `fnsu` user or the `fnsr`, `fnop`, and `fnadmin` groups. Unless otherwise noted, the references apply to both the default user/group names and to the assigned aliases.

These required users and groups are detailed in the following tables. After doing what is noted in the Action column, log out and log back in to pick up the changes. (Solaris uses the directory in

the root file system called **/home** as the users' home directory. If you want to use **/home**, you will have to make a file system for it.)

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

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

### To verify national language character set and time settings

---

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

```
locale -k t_fmt
```

The result might appear similar to this:

```
t_fmt=%r
```

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

```
LC_TIME=C
export LC_TIME
```

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

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

### To create partitions

---

Process Engine requires the presence of the several partitions. You can use third-party software such as Veritas Volume Management to create these partitions, or software provided with the operating system, such as Solstice DiskSuite, to create soft partitions. Regardless of the tool you use, each partition has requirements for the mount point, size, user, group, and mode. These requirements are specified in the following table. There are no requirements for the partition names; the names listed in the tables are examples.



**WARNING** Always use a volume manager on UNIX-based operating systems.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fns	<b>/fns</b> file system	1200MB	fns	fnusr	775
local	<b>/fns/local</b> file system	250MB	fns	fnusr	775
fn_sec_db0 (raw)	n/a logical volume	64MB	fns	fnusr	664
fn_sec_rl0 (raw)	n/a logical volume	64MB	fns	fnusr	664

### To verify the database connection (Oracle)

The database that Process Engine will use must be running and fully configured for IBM FileNet use before you run Process Engine Setup. In addition, if the database is remote, you must also install the database client software on Process Engine server and make sure it is running before you start Process Engine Setup. See [Task 3b on page 85](#) task for database installation and configuration procedures.

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

1. Execute the following at a command prompt:

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

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

- a. If the SQL scripts will be run from Process Engine Setup by prompting for the sys password, type the following command:

```
sys/<password> as sysdba
```

- b. If the SQL scripts will be run from Process Engine Setup 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
9.2.0.7.0
```

In this example:

- the instance, ORACLE\_SID is p8dbshr
- the database server name is hqvwbucs
- the Oracle server is at the 9.2.0.7 version

### Verify the database connection (DB2)

---

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

1. Log on as the client instance owner.
2. Execute the following at a command prompt:

```
db2
```

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

```
connect to <database alias> user <f_sw> using <f_sw password>
```

where <f\_sw> is either the default f\_sw user or the assigned alias

DB2 will display the database connection information.

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

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

```
Database Connection Information
```

```
Database server      = DB2/AIX64 8.2.0
```

```
SQL authorization ID = F_SW
```

```
Local database alias = PEDBINST
```

In this example, the database alias is pedbinst and the f\_sw user password is f\_sw.

## To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the installation process. Turn off this Oracle feature and do not re-enable it until you have installed Process Engine and used the `set_f_maint_pw` utility to change the `f_sw` (or its alias).

See ["To set the `f\_maint` and `f\_sw` passwords \(Oracle only\)" on page 300](#) for procedures to change the `f_sw` password.

## To install the Process Engine software interactively

---

1. Log on to the server as the *root* user.  
**NOTE** This user does not need to be a database administrator unless you will be executing the SQL scripts from the installer.
2. Access the Process Engine software package.
3. From the console, launch the **P8PE-4.0.0-Sol.bin** Setup program.
4. Wait for files to finish unpacking.
5. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement, then click <b>Next</b> .
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation is installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation web site. You can use multi-part root folders (for example, <code>/docs/ecm_help</code>) if your application server supports them.</p>
Specify Installation Location for Common Files	Choose the destination directory for configuration files that will be shared with other IBM FileNet P8 components.

In this screen...	Perform this action...
Specify Network Clearinghouse Domain Name	<p>Enter <i>&lt;domain name&gt;:&lt;organization&gt;</i>, where:</p> <ul style="list-style-type: none"> <li>The maximum length of your <i>&lt;domain name&gt;</i> entry does not exceed 19 characters.</li> <li>The maximum length of your <i>&lt;organization&gt;</i> entry does not exceed 19 characters.</li> <li>Both your <i>&lt;domain name&gt;</i> and <i>&lt;organization&gt;</i> entries contain only alphanumeric characters and underscores.</li> </ul> <p>A typical convention is to enter <i>&lt;your PE machine name&gt;:&lt;your company name&gt;</i>. If the machine name or company name include hyphens, replace them with underscores in your entry.</p>
Specify the Database Location	Indicate whether the database will be local or remote.
Specify the Database Type	Indicate whether the database will be:
<p><b>NOTE</b> This screen only displays if a remote database was selected.</p>	<ul style="list-style-type: none"> <li>Oracle</li> <li>DB2</li> </ul> <p><b>NOTE</b> Only remote DB2 databases are supported.</p>

6. Depending upon your selection here you will be presented with a number of additional screens appropriate to database location and database software. Proceed as appropriate to:
- “To complete remote DB2 database screens” on page 293
  - “To complete remote Oracle database screens” on page 293
  - “To complete local Oracle database screens” on page 295.

**To complete remote DB2 database screens**

In this screen...	Perform this action...
Database Alias Name	Enter the database alias name as assigned in <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> .
Local DB2 instance owner name	Enter the DB2 client instance owner’s name. Case sensitive.
FileNet Data Tablespace name	defaults to VWDATA_TS
DB2 Passwords	Enter the passwords for the Operating System <b>f_sw</b> and <b>f_maint</b> users or their assigned aliases.  Enter the passwords set for these users in either <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a> if database authentication is SERVER or SERVER_ENCRYPT, or <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> if authentication is CLIENT.

**To complete remote Oracle database screens**

Complete the Process Engine Setup screens that are specific to remote Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	Enter the appropriate values for the Oracle Home Directory:  The Oracle Home path you enter refers to the local Oracle installation directory.  The Oracle User Name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 286</a> .  The Oracle DBA OS group name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 286</a> .

In this screen...	Perform this action...
Specify Remote Oracle Configuration Parameters	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Global Database Name (as identified in the tnsnames.ora file)</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85.</a></p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>

**To complete local Oracle database screens**

---

Complete the Process Engine Setup screens that are specific to local Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	<p>Enter the appropriate values for the Oracle Home Directory:</p> <p>The Oracle Home path you enter refers to the local Oracle installation directory.</p> <p>To find the Oracle Home directory, open the a command prompt and type <code>set</code>. Look for the Oracle_Home variable.</p> <p>The Oracle User Name modified in <a href="#">“To create or modify internal Process Engine users and groups” on page 286</a>.</p> <p>The Oracle DBA OS group name modified in <a href="#">“To create or modify internal Process Engine users and groups” on page 286</a>.</p>
Oracle Database Information	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Oracle SID</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85</a>.</p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>

In this screen...	Perform this action...
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>



**To complete final Process Engine Setup screens**

---

Complete the screens to finalize the Process Engine installation, as follows:

In this screen...	Perform this action...
Determine Administrative User and Group Aliasing Method	Determine whether FileNet default operating system and database users and groups will be used, or if you want to define aliases for these users and groups. <ul style="list-style-type: none"> <li>• Yes, configure aliases</li> <li>• No, use actual account names for aliases</li> </ul>
Specify Administrative User and Group Aliases  <b>NOTE</b> This screen is only presented if you chose to configure aliases.	Indicate the alias you want to create for each of the following users and groups. <ul style="list-style-type: none"> <li>• FNADMIN OS group</li> <li>• FNUSR OS group</li> <li>• FNOP OS group</li> <li>• FNSW OS user</li> <li>• f_sw DB user</li> <li>• f_maint DB user</li> </ul> <b>NOTE</b> If you ran SQL scripts manually before starting Process Engine Setup, the aliases here for f_sw and f_maint must match the users specified as runtime and maintenance users.
Specify Device File Names	Enter the full pathname for the device files for the fn_SEC_DB0 and fn_SEC_RL0 volumes.  For example, if Solstice DiskSuite software was used, and fn_SEC_DB0 was on the d100 raw device, the entry would be:  <b>/dev/md/rdsk/d100</b>
Choose an Application Server	Select an application server and version from the drop down boxes. You must use the same application server type and version as Content Engine.  Click <b>Next</b> .

In this screen...	Perform this action...
Content Engine API Configuration	<p>Configure the Content Engine API, as follows:</p> <ul style="list-style-type: none"> <li>• Transport Method Select WSI from the drop-down list.</li> <li>• Content Engine Client Software URL: Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example:  for WebLogic <code>http://hqcomp1:7001/wsi/FNCEWS40DIME</code>  for WebSphere <code>http://hqcomp2:9080/wsi/FNCEWS40DIME</code>  for JBoss: <code>http://hqcomp3:8080/wsi/FNCEWS40DIME</code>  Do not modify the remainder of the string from the default values.  <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.  Click <b>Next</b>. </li></ul>
Please Read the information below	Review the information on the installation process. Note that you can check the progress of the installation in the <b>/fnsw/local/logs/wizard</b> file. Click <b>Next</b> when you are ready to proceed.
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.
xterm Window	Enter the SYS password for Oracle.  <b>NOTE</b> This window displays if you chose to run the SQL scripts in an xterm window.

In this screen...	Perform this action...
Completing the Setup wizard	Click <b>Finish</b> to complete the Process Engine installation.
The wizard requires that you log out and log back in.	Click <b>Finish</b> , log off and log back in as fnsw or the alias you defined in <a href="#">“To create or modify internal Process Engine users and groups”</a> on page 286.

- Setup will run for several minutes. Progress is displayed and at some points might not advance for an extended period of time.
- Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	<b>/fnsw/local/logs</b> and all its subdirectories
IS mini-installer logs	<b>/fnsw/tmp_installer/ MINI_INSTALLER_ERROR.LOG</b>

- Log off as the root user and log in as fnsw (or the alias).
- Proceed to [“To reset administrative user passwords”](#) on page 300.

### To install the Process Engine software silently

Complete all procedures in this topic, up to [“To install the Process Engine software interactively”](#) on page 291. Then take the following steps to silently install Process Engine.

- Access the Process Engine software package and copy the content to a temporary directory on the local disk.
- Edit the **PE\_silent\_input.txt** file to reflect the appropriate responses for your installation. See [“Encrypt Passwords for Silent Installations and Upgrades”](#) on page 642 for information on use of the password encryption tool.
- Save the edited response file to your temporary directory.
- Log on as the root user in the Korn shell.
- Navigate to the temporary directory on the local disk.
- Open a command prompt and execute:
- `P8PE-4.0.0-Sol.bin -silent -options PE_silent_input.txt`
- Proceed to [“To reset administrative user passwords”](#) on page 300.

**To reset administrative user passwords**

Process Engine Setup automatically creates several internally required, local administrative users. These local users have no corresponding LDAP accounts. Because these users are created with default passwords, IBM strongly recommends resetting the passwords for these users to maintain system security, as shown in the following table. See the IBM FileNet Image Services documentation for instructions on using the referenced tools.

User Name	User Type	Description	How to modify
f_maint	Database (Oracle)	Has DBA privileges. Used for RDBMS maintenance.	Execute set_f_maint_pw. See steps below.
SysAdmin	SEC (internal Process Engine security software)	Primary administrator user for IBM FileNet software tools.	Execute Xapex -> Security Administration. Log on as SysAdmin.
FieldService	SEC	Used internally by Process Engine software.	
Operator	SEC	Used internally by Process Engine software.	

**To set the f\_maint and f\_sw passwords (Oracle only)**

Use the following procedures to change the passwords for these users after installing the Process Engine software. These procedures will also verify the ability to connect to the Process Engine database.

**NOTE** You must also use the set\_f\_maint\_pw utility to change the f\_maint and f\_sw passwords later, such as when the passwords are due to expire. Changing these passwords any other way will cause the Process Engine software to fail.

1. Log on as the fns user.
2. At a command prompt, enter:  

```
set_f_maint_pw <user name>
```

 where <user name> is f\_maint, f\_sw or the aliases defined for them in Process Engine Setup.
3. Click **Return** at the prompt for the current password.  
**Note** Pressing Return will only work the first time the password is changed, when the default password is still valid. After the password is changed from the default, this prompt requires you to enter the current password.
4. Enter a new password for the user when prompted. The password must be between 6 and 30 characters.
5. Confirm the new password when prompted.

---

## To re-enable Oracle Password Complexity Verification

---

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

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

---

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

---

## To enable ports

---

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons. By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet uses several ports higher than 32768. See ["IBM FileNet P8 Port Numbers" on page 643](#) for details on which ports IBM FileNet uses.

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

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

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

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

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

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

Enter the following line:

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

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

Enter the following line:

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

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

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

---

## To restore any custom modifications for root and fnsw users.

---

Process Engine Setup 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 in are `<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 Process Engine Setup:

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

#### To verify the /etc/services file settings

---

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

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

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

#### To clean up before starting Process Engine

---

1. Log in as fns.

2. Execute:

```
killfns -DAyS
```

3. Execute:

```
ipcs -a
```

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

### To edit the /etc/inittab file

---

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.

For example, change:

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

to

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

Proceed to [“To install the Process Engine software updates”](#) on page 338.

## Task 19c: Install Process Engine (AIX)

### NOTES

- Before beginning this task, ensure that all applicable tasks listed in the [“Installation Planning Considerations” on page 24](#) have been completed. Process Engine requires the presence of several partitions. Before installing PE verify that your Operating System is set up with a correctly configured volume manager.
- You will find references to logging on as the root and fnsu users within the following procedures. The root user must run in the Bourne or Korn shell and the fnsu user must run in the Korn shell.
- Before starting Process Engine installation, be sure to have the correct database information including server names, database and instance names, and tablespace names. Entering incorrect information during installation could cause the installation to fail. Recovery could mean uninstalling and re-installing the Process Engine software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to install and configure Process Engine but can be re-enabled after the installation is complete. See [“To turn off Oracle Password Complexity Verification” on page 309](#) and [“To re-enable Oracle Password Complexity Verification” on page 318](#) for details. If this will be a silent installation, complete all steps in this task, up to [“To install the Process Engine software interactively” on page 309](#). After completing silent installation steps, return to this task, at [“To reset administrative user passwords” on page 317](#) and proceed through the remainder of this task.
- Determine when to execute pre-installation SQL scripts for Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle in an xterm window.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts” on page 647](#) for detailed information on the scripts and modes of execution.

### To set time zone parameters

---

Review and change the time zone parameters if necessary. In SMIT, choose System Environments > Change/Show Date and Time > Change Time Zone Using System Defined Values. Choose the Daylight Savings Time option if applicable. At the CUT Time Zone menu, choose the option associated with your site. For example, in California, the time zone needs to be set to the Pacific time zone (PST8PDT) Pacific U.S.; Yukon (cut -8).



**To modify /etc/rc.dt and /etc/tunables/nextboot for AIX 5.2 and 5.3**

---

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

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

2. Add the following statements at the beginning for **/etc/rc.dt** file:

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

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

4. Check the values by executing:

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

**To create or modify internal Process Engine users and groups**

---

You must create or modify specific local groups and users required internally by Process Engine prior to running Process Engine Setup. Optionally, aliases for these users and groups can be defined. Process Engine Setup will ask whether to use the default users and groups or identify the aliases that have been defined.

**NOTE** Throughout the remainder of this document, references will be made to the fnsu user or the fnsu, fnop, and fnadmin groups. Unless otherwise noted, the references apply to both the default user/group names and to the assigned aliases.

These required users and groups are detailed in the following tables. After doing what is noted in the Action column, log out and log back in to pick up the changes.

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

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

- Create the fnsw or alias user with fnusr as its primary group.
- Create the fnadmin group with fnsw and root as its members.

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

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

## To create partitions

---

Process Engine requires the presence of several partitions. The following table documents the requirements. All volume names must be spelled exactly as specified here.

**WARNING** Always use a volume manager on UNIX-based operating systems.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fns	<b>/fns</b> file system	1200MB	fns	fnusr	775
local	<b>/fns/local</b> file system	250MB	fns	fnusr	775
fn_sec_db0 (raw)	n/a logical volume	64MB	fns	fnusr	664
fn_sec_rl0 (raw)	n/a logical volume	64MB	fns	fnusr	664

## To verify the database connection (Oracle)

---

The database that Process Engine will use must be running and fully configured for IBM FileNet use before you run Process Engine Setup. If the database is remote, you must also install the database client software on the Process Engine server and make sure it is running before you start Process Engine Setup. See [“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85](#) task for database installation and configuration procedures.

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

1. Execute the following at a command prompt:

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

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

- a. If the SQL scripts will be run from Process Engine Setup by prompting for the sys password, type the following command:

```
sys/<password> as sysdba
```

- b. If the SQL scripts will be run from Process Engine Setup 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
9.2.0.7.0
```

In this example:

- the instance, ORACLE\_SID is p8dbshr
- the database server name is hqvwbucs
- the Oracle server is at the 9.2.0.7 version

### **To verify the database connection (DB2)**

---

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

1. Log on as the client instance owner.
2. Execute the following at a command prompt:

```
db2
```

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

```
connect to <database alias> user <f_sw> using <f_sw password>
```

where <f\_sw> is either the default f\_sw user or the assigned alias

DB2 will display the database connection information.

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

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

Database Connection Information

Database server      = DB2/AIX64 8.2.0
SQL authorization ID = F_SW
Local database alias = PEDBINST
```

In this example, the database alias is pedbinst and the f\_sw user password is f\_sw.

### To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the installation process. Turn off this Oracle feature and do not re-enable it until you have installed Process Engine and used the set\_f\_maint\_pw utility to change the f\_sw (or its alias).

See ["To set the f\\_maint and f\\_sw passwords \(Oracle only\)" on page 317](#) for procedures to change the f\_sw password.

### To install the Process Engine software interactively

---

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

**NOTE** This user does not need to be a database administrator unless you will be executing the SQL scripts from the installer.
2. Access the Process Engine software package.
3. From the console, launch the **P8PE-4.0.0-AIX.bin** Setup program.
4. Wait for files to finish unpacking.
5. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement.

In this screen...	Perform this action...
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation is installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation web site. You can use multi-part root folders (for example, <b>/docs/ecm_help</b>) if your application server supports them.</p>
Specify Installation Location for Common Files	<p>Choose the destination directory for configuration files that will be shared with other IBM FileNet P8 components.</p>
Specify Network Clearinghouse Domain Name	<p>Enter <code>&lt;domain name&gt;:&lt;organization&gt;</code>, where:</p> <ul style="list-style-type: none"> <li>• The maximum length of your <code>&lt;domain name&gt;</code> entry does not exceed 19 characters.</li> <li>• The maximum length of your <code>&lt;organization&gt;</code> entry does not exceed 19 characters.</li> <li>• Both your <code>&lt;domain name&gt;</code> and <code>&lt;organization&gt;</code> entries contain only alphanumeric characters and underscores.</li> </ul> <p>A typical convention is to enter <code>&lt;your PE machine name&gt;:&lt;your company name&gt;</code>. If the machine name or company name include hyphens, replace them with underscores in your entry.</p>
Specify the Database Location	<p>Indicate whether the database will be local or remote.</p>
Specify the Database Type	<p>Indicate whether the database will be:</p> <ul style="list-style-type: none"> <li>• Oracle</li> <li>• DB2</li> </ul> <p><b>NOTE</b> Only remote DB2 databases are supported.</p>
<b>NOTE</b> This screen only displays if a remote database was selected.	

6. Depending upon your selection here you will be presented with a number of additional screens appropriate to database location and database software. Proceed as appropriate to:
- [“To complete remote DB2 database screens” on page 311](#)
  - [“To complete remote Oracle database screens” on page 311](#)
  - [“To complete local Oracle database screens” on page 313.](#)

**To complete remote DB2 database screens**

In this screen...	Perform this action...
Database Alias Name	Enter the database alias name as assigned in <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> .
Local DB2 instance owner name	Enter the DB2 client instance owner’s name. Case sensitive.
Data Tablespace name	defaults to VWDATA_TS
DB2 Passwords	Enter the passwords for the Operating System <b>f_sw</b> and <b>f_maint</b> users or their assigned aliases.  Enter the passwords set for these users in either <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a> if database authentication is SERVER or SERVER_ENCRYPT, or <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> if authentication is CLIENT.

**To complete remote Oracle database screens**

Complete the Process Engine Setup screens that are specific to remote Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	Enter the appropriate values for the Oracle Home Directory:  The Oracle Home path you enter refers to the local Oracle installation directory.  The Oracle User Name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 305</a> .  The Oracle DBA OS group name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 305</a> .

In this screen...	Perform this action...
Specify Remote Oracle Configuration Parameters	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Global Database Name (as identified in the tnsnames.ora file)</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85.</a></p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>



**To complete local Oracle database screens**

---

Complete the Process Engine Setup screens that are specific to local Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	<p>Enter the appropriate values for the Oracle Home Directory:</p> <p>The Oracle Home path you enter refers to the local Oracle installation directory.</p> <p>To find the Oracle Home directory, open the a command prompt and type "set". Look for the Oracle_Home variable.</p> <p>The Oracle User Name modified in <a href="#">"To create or modify internal Process Engine users and groups" on page 305.</a></p> <p>The Oracle DBA OS group name modified in <a href="#">"To create or modify internal Process Engine users and groups" on page 305.</a></p>
Oracle Database Information	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Oracle SID</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">"Verify that Oracle Server Is Installed for IBM FileNet P8" on page 85.</a></p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>

In this screen...	Perform this action...
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts need to be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>

### To complete final Process Engine Setup screens

Complete the screens to finalize the Process Engine installation, as follows:

In this screen...	Perform this action...
Determine Administrative User and Group Aliasing Method	<p>Determine whether FileNet default operating system and database users and groups will be used, or if you want to define aliases for these users and groups.</p> <ul style="list-style-type: none"> <li>• Yes, configure aliases</li> <li>• No, use actual account names for aliases</li> </ul>
Specify Administrative User and Group Aliases	<p>Indicate the alias you want to create for each of the following users and groups.</p> <ul style="list-style-type: none"> <li>• FNADMIN OS group</li> <li>• FNUSR OS group</li> <li>• FNOP OS group</li> <li>• FNSW OS user</li> <li>• f_sw DB user</li> <li>• f_maint DB user</li> </ul> <p><b>NOTE</b> If you ran SQL scripts manually before starting Process Engine Setup, the aliases here for f_sw and f_maint must match the users specified as runtime and maintenance users.</p>

In this screen...	Perform this action...
Choose an Application Server	Select an application server and version from the drop down boxes. You must use the same application server type and version as Content Engine.  Click <b>Next</b> .
Content Engine API Configuration	Configure the Content Engine API, as follows: <ul style="list-style-type: none"> <li>• Transport Method                              Select WSI from the drop-down list.</li> <li>• Content Engine Client Software URL:                              Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example:                               for WebLogic                              http://hqcomp1:7001/wsi/FNCEWS40DIME                               for WebSphere                              http://hqcomp2:9080/wsi/FNCEWS40DIME                               for JBoss:                              http://hqcomp3:8080/wsi/FNCEWS40DIME                               Do not modify the remainder of the string from the default values.   <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.                         </li> </ul>
Please Read the information below	Review the information on the installation process. Note that you can check the progress of the installation in the <b>/fnsw/local/logs/wizard</b> file. Click <b>Next</b> when you are ready to proceed.
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.

In this screen...	Perform this action...
xterm Window  <b>NOTE</b> This window displays if you chose to run the SQL scripts in an xterm window.	Enter the SYS password for Oracle.
Completing the Setup wizard	Click <b>Finish</b> to complete the Process Engine installation.
The wizard requires that you log out and log back in.	Click <b>Finish</b> , log off and log back in as fnsw or the alias you defined in <a href="#">“To create or modify internal Process Engine users and groups”</a> on page 305.

- Setup will run for several minutes. Progress is displayed and at some points might not advance for an extended period of time.
- Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	<b>/fnsw/local/logs</b> and all its subdirectories
IS mini-installer logs	<b>/fnsw/tmp_installer/ MINI_INSTALLER_ERROR.LOG</b>

- Log off as the root user and log in as fnsw (or the alias).
- Proceed to [“To reset administrative user passwords”](#) on page 317.

**To install the Process Engine software silently**

Complete all procedures in this topic, up to [“To install the Process Engine software interactively”](#) on page 309. Then take the following steps to silently install Process Engine .

- Access the Process Engine software package, and copy the contents to a local temporary directory on the local disk.
- Edit the **PE\_silent\_input.txt** file to reflect the appropriate responses for your installation. See [“Encrypt Passwords for Silent Installations and Upgrades”](#) on page 642 for information on use of the password encryption tool.
- Save the edited response file to your temporary directory.
- Log on as the root user in the Korn shell.
- Navigate to the temporary directory on the local disk.

- Open a command prompt and execute:

```
P8PE-4.0.0-AIX.bin -silent -options PE_silent_input.txt
```

- Proceed to [“To reset administrative user passwords” on page 317.](#)

### To reset administrative user passwords

During Process Engine software installation, several local users required internally by Process Engine are automatically created. To maintain security, IBM strongly recommends resetting the passwords for these users. The following table lists the users created, the level of system access each user has, and the tool used to change the password. See the IBM FileNet Image Services documentation for instructions on using the referenced tools

User Name	User Type	Description	How to modify
f_maint	Database (Oracle)	Has DBA privileges. Used for RDBMS maintenance.	Execute set_f_maint_pw. See steps below.
SysAdmin	SEC (internal Process Engine security software)	Primary administrator user for IBM FileNet software tools.	Execute Xapex -> Security Administration. Log on as SysAdmin.
FieldService	SEC	Used internally by Process Engine software.	
Operator	SEC	Used internally by Process Engine software.	

### To set the f\_maint and f\_sw passwords (Oracle only)

Use the following procedures to change the passwords for these users after installing the Process Engine software. These procedures will also verify the ability to connect to the Process Engine database.

**NOTE** You must also use the set\_f\_maint\_pw utility to change the f\_maint and f\_sw passwords later, such as when the passwords are due to expire. Changing these passwords any other way will cause the Process Engine software to fail.

- Log on as the fns user.
- At a command prompt, enter:

```
set_f_maint_pw <user name>
```

where <user name> is f\_maint, f\_sw or the aliases defined for them in Process Engine Setup.

- Click **Return** at the prompt for the current password.

**Note** Pressing Return will only work the first time the password is changed, when the default password is still valid. After the password is changed from the default, this prompt requires you to enter the current password.

4. Enter a new password for the user when prompted. The password must be between 6 and 30 characters.
5. Confirm the new password when prompted.

---

**To re-enable Oracle Password Complexity Verification**

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

---

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

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

---

**To restore any custom modifications for root and fnsw users.**

Process Engine Setup 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 in are *<.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 Process Engine Setup:

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

---

**To verify the /etc/services file settings**

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

```
smux      199/tcp          # snmpd smux port
tms       32768/tcp
cor       32769/tcp
nch       32770/udp
fn_trapd  35225/udp
```

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

### To clean up before starting Process Engine

---

1. Log in as fnsw.
2. Execute:  

```
killfnsw -DAY5
```
3. Execute:  

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

### To edit the `/etc/inittab` file

---

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.

For example, change:

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

to

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

Proceed to [“Install Process Engine Software Updates”](#) on page 338.

## Task 19d: Install Process Engine (HP-UX)

These procedures apply to Process Engine on both the HP PA-RISC and Integrity platforms.

### NOTES

- Before beginning this task, ensure that all applicable tasks listed in the [“Installation Planning Considerations” on page 24](#) have been completed.
- Process Engine requires the presence of several partitions. Before installing PE verify that your Operating System is set up with a correctly configured volume manager.
- You will find references to logging on as the root and fnsu users within the following procedures. The root user must run in the Bourne or Korn shell and the fnsu user must run in the Korn shell.
- Before starting Process Engine installation, be sure to have the correct database information including server names, database and instance names, and tablespace names. Entering incorrect information during installation could cause the installation to fail. Recovery could mean uninstalling and re-installing the Process Engine software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to install and configure Process Engine but can be re-enabled after the installation is complete. See [“To turn off Oracle Password Complexity Verification” on page 327](#) and [“To re-enable Oracle Password Complexity Verification” on page 336](#) for details.
- If this will be a silent installation, complete all steps in this task, up to [“To install the Process Engine software interactively” on page 327](#). After completing silent installation steps, return to this task, at [“To reset administrative user passwords” on page 335](#) and proceed through the remainder of this task.
- Determine when to execute pre-installation SQL scripts for Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle in an xterm window.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts” on page 647](#) for detailed information on the scripts and modes of execution.



## To perform symbolic links for X11 libraries

---

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

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

## To check and optionally modify the timezone settings

---

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

1. As the root user, enter:

```
sam
```

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

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

4. If the Pending Value for the `timezone` parameter is correct, proceed to step 6. To change the value, continue with Step 5.
5. To change the `timezone` kernel parameter value:
  - a. Select the `timezone` parameter by pressing the spacebar and then press Tab to go to Actions menu.
  - b. Select the Modify Configurable Parameter option from the Actions menu and press Return.
  - c. In the popup window that displays, the Specify New Formula/Value option should already be selected.
  - d. Tab to the Formula/Value field and type the new value.
  - e. Tab to OK and press Return. When the popup window disappears, you should see the new value in the Pending Value column.

- f. Rebuild the kernel to make your change take effect:
  - i. Press the F4 key to access the menu bar.
  - ii. From the Action menu, select the Create a New Kernel option using the Arrow keys and press Return.
  - iii. Answer Yes when prompted about creating the kernel now.
  - iv. On the next screen, make sure the Move Kernel into Place and Shutdown/Reboot the System Now option is selected, tab to OK and press the Return key to reboot the system and make the new changes take affect.

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

```
echo $TZ
```

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

```
/sbin/set_parms timezone
```

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

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

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

### To configure kernel parameters

Ensure that the following parameters are set to at least the values shown unless otherwise noted. The values are appropriate for both HP PA-RISC and Integrity operating systems unless noted otherwise. These values are sufficient to install and initialize the software but system tuning will be required, specifically for the nfiles and maxfiles parameters. See the following HP web site for configuration details: [http://www.hp.com/products1/unix/java/infoclibrary/prog\\_guide/java1/configuration.html](http://www.hp.com/products1/unix/java/infoclibrary/prog_guide/java1/configuration.html)

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

Kernel Parameter	Minimum Setting
semgni	2000
semmap	2002 (not applicable to Integrity)
shmmax	0x20000000 or 536870912 (512MB)
shmseg	120
semmnu	1000
semume	500
msgmni	2048
msgseg	16384
msgtql	6640
msgmap	msgtql + 2
bufpages	0 (required by Oracle) (not applicable to Integrity)
dbc_max_pct	1 to 30 (required by Oracle, the value can not be greater than 30)
dbc_min_pct	5 (required by Oracle)
fs_async	0 (required by Oracle)
maxfiles_lim	1024
npty	60
o_sync_is_o_dsync	0
timezone	Set appropriately For details, see <a href="#">"To check and optionally modify the timezone settings"</a> on page 321.

### To create or modify internal Process Engine users and groups

Using SAM, you must create or modify specific local groups and users required internally by Process Engine prior to running Process Engine Setup. Optionally, aliases for these users and groups can be defined. Process Engine Setup will ask whether to use the default users and groups or identify the aliases that have been defined.

**NOTE** Throughout the remainder of this document, references will be made to the fnsu user or the fnusr, fnop, and fnadmin groups. Unless otherwise noted, the references apply to both the default user/group names and to the assigned aliases.

Create these groups and users in the following order:

- Create groups (Without members)
- Create users

- Modify users (attributes and/or secondary group memberships)

These required users and groups are detailed in the following tables. After doing what is noted in the Action column, log out and log back in to pick up the changes.

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

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

## To create partitions

---

Process Engine requires the presence of several partitions. The following table documents the requirements.

**WARNING** Always use a volume manager on UNIX-based operating systems.

Volume Name	Mount Point	Minimum Size	User	Group	Mode
fns	<b>/fns</b> file system	1200MB	fns	fnusr	775
local	<b>/fns/local</b> file system	250MB	fns	fnusr	775
fn_sec_db0 (raw)	n/a logical volume	64 MB	fns	fnusr	664
fn_sec_rl0 (raw)	n/a logical volume	64MB	fns	fnusr	664

## To verify the database connection (Oracle)

---

The database that Process Engine will use must be running and fully configured for IBM FileNet use before you run Process Engine Setup. In addition, if the database is remote, you must also install the database client software on Process Engine server and make sure it is running before you start Process Engine Setup. See [“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85](#) task for database installation and configuration procedures.

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

1. Execute the following at a command prompt:

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

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

- a. If the SQL scripts will be run from Process Engine Setup by prompting for the sys password, type the following command:

```
sys/<password> as sysdba
```

- b. If the SQL scripts will be run from Process Engine Setup 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
9.2.0.7.0
```

In this example:

- the instance, ORACLE\_SID is p8dbshr
- the database server name is hqvwbucs
- the Oracle server is at the 9.2.0.7 version

### **Verify the database connection (DB2)**

---

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

1. Log on as the client instance owner.
2. Execute the following at a command prompt:

```
db2
```

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

```
connect to <database alias> user <f_sw> using <f_sw password>
```

where <f\_sw> is either the default f\_sw user or the assigned alias

DB2 will display the database connection information.

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

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

Database Connection Information

Database server      = DB2/AIX64 8.2.0
SQL authorization ID = F_SW
Local database alias = PEDBINST
```

In this example, the database alias is pedbinst and the f\_sw user password is f\_sw.

### To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the installation process. Turn off this Oracle feature and do not re-enable it until you have installed Process Engine and used the set\_f\_maint\_pw utility to change the f\_sw (or its alias).

See ["To set the f\\_maint and f\\_sw passwords \(Oracle only\)" on page 335](#) for procedures to change the f\_sw password.

### To install the Process Engine software interactively

---

1. Log in as root.

**NOTE** This user does not need to be a database administrator unless you will be executing the SQL scripts from the installer.

2. Access the Process Engine software package and execute **P8PE-4.0.0-HPUX.bin** for PA-RISC, or **P8PE-4.0.0-HPUXi.bin** for HP Integrity.
3. Wait for files to finish unpacking.
4. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement.

In this screen...	Perform this action...
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation is installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation web site. You can use multi-part root folders (for example, <code>/docs/ecm_help</code>) if your application server supports them.</p>
Specify Installation Location for Common Files	<p>Choose the destination directory for configuration files that will be shared with other IBM FileNet P8 components.</p>
Specify Network Clearinghouse Domain Name	<p>Enter <code>&lt;domain name&gt;:&lt;organization&gt;</code>, where:</p> <ul style="list-style-type: none"> <li>• The maximum length of your <code>&lt;domain name&gt;</code> entry does not exceed 19 characters.</li> <li>• The maximum length of your <code>&lt;organization&gt;</code> entry does not exceed 19 characters.</li> <li>• Both your <code>&lt;domain name&gt;</code> and <code>&lt;organization&gt;</code> entries contain only alphanumeric characters and underscores.</li> </ul> <p>A typical convention is to enter <code>&lt;your PE machine name&gt;:&lt;your company name&gt;</code>. If the machine name or company name include hyphens, replace them with underscores in your entry.</p>
Specify the Database Location	<p>Indicate whether the database will be local or remote.</p>
Specify the Database Type	<p>Indicate whether the database will be:</p> <ul style="list-style-type: none"> <li>• Oracle</li> <li>• DB2</li> </ul> <p><b>NOTE</b> Only remote DB2 databases are supported.</p>
<b>NOTE</b> This screen only displays if a remote database was selected.	

Depending upon your selection here you will be presented with a number of additional screens appropriate to database location and database software. Proceed as appropriate to:

- [“To complete remote DB2 database screens” on page 329](#)
- [“To complete remote Oracle database screens” on page 329](#)



- [“To complete local Oracle database screens” on page 331.](#)

**To complete remote DB2 database screens**

---

In this screen...	Perform this action...
Database Alias Name	Enter the database alias name as assigned in <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103.</a>
Local DB2 instance owner name	Enter the DB2 client instance owner’s name. Case sensitive.
FileNet Data Tablespace name	defaults to VWDATA_TS
DB2 Passwords	<p>Enter the passwords for the Operating System <b>f_sw</b> and <b>f_maint</b> users or their assigned aliases.</p> <p>Enter the passwords set for these users in either <a href="#">“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92</a> if database authentication is SERVER or SERVER_ENCRYPT, or <a href="#">“Verify that DB2 Client Is Installed for IBM FileNet P8” on page 103</a> if authentication is CLIENT.</p>

**To complete remote Oracle database screens**

---

Complete the Process Engine Setup screens that are specific to remote Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	<p>Enter the appropriate values for the Oracle Home Directory:</p> <p>The Oracle Home path you enter refers to the local Oracle installation directory.</p> <p>The Oracle User Name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 323.</a></p> <p>The Oracle DBA OS group name created in <a href="#">“To create or modify internal Process Engine users and groups” on page 323.</a></p>

In this screen...	Perform this action...
Specify Remote Oracle Configuration Parameters	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Global Database Name (as identified in the tnsnames.ora file)</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85.</a></p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts need to be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>

**To complete local Oracle database screens**

---

Complete the Process Engine Setup screens that are specific to local Oracle databases, as follows:

In this screen...	Perform this action...
Specify Oracle Configuration Parameters	<p>Enter the appropriate values for the Oracle Home Directory:</p> <p>The Oracle Home path you enter refers to the local Oracle installation directory.</p> <p>To find the Oracle Home directory, open the a command prompt and type "set". Look for the Oracle_Home variable.</p> <p>The Oracle User Name modified in <a href="#">"To create or modify internal Process Engine users and groups" on page 323.</a></p> <p>The Oracle DBA OS group name modified in <a href="#">"To create or modify internal Process Engine users and groups" on page 323.</a></p>
Oracle Database Information	<p>Enter the following Oracle database configuration parameters.</p> <ul style="list-style-type: none"> <li>• Oracle SID</li> <li>• Temporary Tablespace Name Default value is VWTEMP_TS</li> <li>• Data Tablespace Name Default value is VWDATA_TS</li> <li>• Index Tablespace Name Default value is VWINDEX_TS</li> </ul> <p>Enter the optional tablespace to be used by Process Engine for indexes. The data tablespace will be used if no index tablespace is designated.</p> <p>All values for tablespace names must match those you used when tablespaces were created in <a href="#">"Verify that Oracle Server Is Installed for IBM FileNet P8" on page 85.</a></p>
Specify Oracle Version	<p>Both Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and server.</p>

In this screen...	Perform this action...
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts need to be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>

**To complete final Process Engine Setup screens**

Complete the screens to finalize the Process Engine installation, as follows:

In this screen...	Perform this action...
Determine Administrative User and Group Aliasing Method	<p>Determine whether FileNet default operating system and database users and groups will be used, or if you want to define aliases for these users and groups.</p> <ul style="list-style-type: none"> <li>• Yes, configure aliases</li> <li>• No, use actual account names for aliases</li> </ul>
Specify Administrative User and Group Aliases  <b>NOTE</b> This screen is only presented if you chose to configure aliases.	<p>Indicate the alias you want to create for each of the following users and groups.</p> <ul style="list-style-type: none"> <li>• FNADMIN OS group</li> <li>• FNUSR OS group</li> <li>• FNOP OS group</li> <li>• FNSW OS user</li> <li>• f_sw DB user</li> <li>• f_maint DB user</li> </ul> <p><b>NOTE</b> If you ran SQL scripts manually before starting Process Engine Setup, the aliases here for f_sw and f_maint must match the users specified as runtime and maintenance users.</p>
Specify Device File Names	<p>Enter the full pathname for the device files for the fn_SEC_DB0 and fn_SEC_RL0 volumes.</p>

In this screen...	Perform this action...
Choose an Application Server	Select an application server and version from the drop down boxes. You must use the same application server type and version as Content Engine.  Click <b>Next</b> .
Content Engine API Configuration	Configure the Content Engine API, as follows: <ul style="list-style-type: none"> <li>• Transport Method                              Select WSI from the drop-down list.</li> <li>• Content Engine Client Software URL                              Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example:                               for WebLogic                              http://hqcomp1:7001/wsi/FNCEWS40DIME                               for WebSphere                              http://hqcomp2:9080/wsi/FNCEWS40DIME                               for JBoss                              http://hqcomp3:8080/wsi/FNCEWS40DIME                               Do not modify the remainder of the string from the default values.   <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.                               Click <b>Next</b>.                         </li> </ul>
Please Read the information below	Review the information on the installation process. Note that you can check the progress of the installation in the <b>/fnsw/local/logs/wizard</b> file. Click <b>Next</b> when you are ready to proceed.
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.

In this screen...	Perform this action...
xterm Window  <b>NOTE</b> This window displays if you chose to run the SQL scripts in an xterm window.	Enter the SYS password for Oracle.
Completing the Setup wizard	Click <b>Finish</b> to complete the Process Engine installation.
The wizard requires that you log out and log back in.	Click <b>Finish</b> , log off and log back in as fnsw or the alias you defined in <a href="#">“To create or modify internal Process Engine users and groups”</a> on page 323.

- Setup will run for several minutes. Progress is displayed and at some points might not advance for an extended period of time.
- Check the following log files and correct any errors or failures indicated before proceeding to the next step:

Log	Location
Process Engine logs	<b>/fnsw/local/logs</b> and all its subdirectories
IS mini-installer logs	<b>/fnsw/tmp_installer/ MINI_INSTALLER_ERROR.LOG</b>

- Log off as the root user and log in as fnsw (or the alias).
- Proceed to [“To reset administrative user passwords”](#) on page 335.

### To install the Process Engine software silently

Complete all procedures in this topic, up to [“To install the Process Engine software interactively”](#) on page 327. Then take the following steps to silently install Process Engine .

- Access the Process Engine installation software, and copy the contents to a local temporary directory on the local disk.
- Edit the **PE\_silent\_input.txt** file to reflect the appropriate responses for your installation. See [“Encrypt Passwords for Silent Installations and Upgrades”](#) on page 642 for information on use of the password encryption tool.
- Save the edited response file to your temporary directory.
- Log on as the root user in the Korn shell.
- Navigate to the temporary directory on the local disk.

- Open a command prompt and execute:

```
P8PE-4.0.0-HPUX.bin -silent -options PE_silent_input.txt
```

or

```
P8PE-4.0.0-HPUXi.bin -silent -options PE_silent_input.txt
```

- Proceed to [“To reset administrative user passwords” on page 335](#).

### To reset administrative user passwords

During Process Engine software installation, several local users required internally by Process Engine are automatically created. To maintain security, IBM strongly recommends resetting the passwords for these users. The following table lists the users created, the level of system access each user has, and the tool used to change the password. See the IBM FileNet Image Services documentation for instructions on using the referenced tools.

User Name	User Type	Description	How to modify
f_maint	Database (Oracle)	Has DBA privileges. Used for RDBMS maintenance.	Execute set_f_maint_pw. See steps below.
SysAdmin	SEC (internal Process Engine security software)	Primary administrator user for IBM FileNet software tools.	Execute Xapex -> Security Administration. Log on as SysAdmin.
FieldService	SEC	Used internally by Process Engine software.	
Operator	SEC	Used internally by Process Engine software.	

### To set the f\_maint and f\_sw passwords (Oracle only)

Use the following procedures to change the passwords for these users after installing the Process Engine software. These procedures will also verify the ability to connect to the Process Engine database.

**NOTE** You must also use the set\_f\_maint\_pw utility to change the f\_maint and f\_sw passwords later, such as when the passwords are due to expire. Changing these passwords any other way will cause the Process Engine software to fail.

- Log on as the fsw user.
- At a command prompt, enter:

```
set_f_maint_pw <user name>
```

where <user name> is f\_maint, f\_sw or the aliases defined for them in Process Engine Setup.

- Click **Return** at the prompt for the current password.

**Note** Pressing Return will only work the first time the password is changed, when the default password is still valid. After the password is changed from the default, this prompt requires you to enter the current password.

4. Enter a new password for the user when prompted. The password must be between 6 and 30 characters.
5. Confirm the new password when prompted.

---

### To re-enable Oracle Password Complexity Verification

---

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

---

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

---

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

---

### To restore any custom modifications for root and fnsw users.

---

Process Engine Setup 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 in are *<.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 Process Engine Setup:

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



### To verify the /etc/services file settings

---

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

```
snmp      161/udp  snmpd      # Simple Network Management Protocol Agent
snmp-trap 162/udp  trapd      # Simple Network Management Protocol Traps
tms       32768/tcp
cor       32769/tcp
nch       32770/udp
fn_trapd  35225/udp
```

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

### To clean up before starting Process Engine

---

1. Log in as fnsf.

2. Execute:

```
killfnsf -DAyS
```

3. Execute:

```
ipcs -a
```

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

### To edit the /etc/inittab file

---

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.

For example, change:

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

to

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

### To edit the ims\_start file

---

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

Change:

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

to

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

Proceed to [“Install Process Engine Software Updates”](#) on page 338.

## Task 20: Install Process Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Process Engine.

### To install the Process Engine software updates

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. Process Engine 4.0.2 Service Pack
  - b. Any subsequent fix pack (P8PE-4.0.2-001 or later)
  - c. Any subsequent interim fixes (typically optional)

Proceed to [“Install the Latest Content Engine Client Files on Process Engine Servers”](#) on page 339.

## Task 21: Install the Latest Content Engine Client Files on Process Engine Servers

Install any Content Engine client file updates that are available.

### **To install the Content Engine client files**

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. Open the readme for the P8CE-4.0.0-002 (or later) Fix Pack and perform the installation procedures provided to install the Content Engine Java Client files:
  - a. Content Engine 4.0.1 Client Updater
  - b. Any subsequent fix pack (P8CE-4.0.1-001 or later)
  - c. Any subsequent interim fixes (typically optional)

Proceed to [“Configure Process Task Manager”](#) on page 340.

## Task 22: Configure Process Task Manager

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

**CAUTION** If you are installing with a DB2 database, do not configure Process Engine or start the software until you have installed Process Engine 4.0.2 or later.

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

---

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

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

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

2. Start the Process Task Manager, as follows:

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

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

3. Right-click your Process Engine server in the feature pane.
4. If the Process Engine software is not already running, start it by choosing Start from the Action menu.
5. Select the Process Engine in the feature pane and click the **Security** tab to configure the General settings.

Provide the service username defined in [“To create other Process Engine accounts” on page 78](#). See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Configure the Process Engine > Security](#) for details on the user and groups.

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

Once complete, click **Apply**.

6. Click **Yes** to restart the Process Task Manager.
7. Click the **Region Passwords** tab and the **Add** icon on the right-most side of the Task Manager screen to add an isolated region and password.

**NOTE** The password must match the password that is entered when creating a Process Engine Region in [“Create a Process Engine Isolated Region” on page 408](#).

8. Click **Yes** to restart the Process Task Manager.

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

On both the General and Advanced tabs, enter the appropriate value for each property. Most properties do not have a default value because they are system specific; not entering values or entering incorrect values will disrupt runtime activity. For property descriptions, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Configure the Process Engine > Security](#).

10. After all parameters have been entered, click **Apply** and restart the Process Service when prompted. Receiving the message to restart the Process Service serves as verification of successful connection to the Content Engine.

### To verify connection to the Process Engine database

---

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

```
vwcomp -l
```

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

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

Proceed to [“Complete Post-Install Process Engine Configuration”](#) on page 342.

## Task 23: Complete Post-Install Process Engine Configuration

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

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

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

1. Start vwtool at a command prompt. Log on using the Service Username you provided when completing the steps in ["Configure Process Task Manager" on page 340](#).
2. Use the processmap command to find the largest contiguous free memory area, as in:

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

vwtool returns the following:

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

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

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

Address	Attrib	Size	Owner
=====	=====	=====	=====
00000000	Free	65536	
00010000	Private	12288	
00013000	Free	53248	
00020000	Private	4096	

.....(pages of memory addresses omitted here)

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

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

```
Largest FREE block found : 453873664 bytes at address 0x4B577000  

Rounded up to a 64K boundary, free block address 0x4B580000
```

In this example, 0x4B580000 is the address we want. In some cases you might only see the line referencing the largest free block because the value is already at a 64K boundary.

```
<vwtool:1>
```

3. Run the Windows regedit command to create a DWORD value for IS StartShmAddress, using the address found in step 2. Run regedit on the Process Engine server and go to:

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

4. Create a new DWORD value. Name it:

`StartShmAddress`

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

Value name = `StartShmAddress`

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

From the example above the value will be `4B580000`.

Base is hexadecimal.

6. Click **OK**.
7. Exit from regedit.
8. Restart the Process Engine software.
9. Verify the setting by executing the following at a command prompt:

`ipc_tool -A`

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

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

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

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

Shared Memory Address Manager Information

Address	Shm id	Creator
---------	--------	---------

Enter <space> to continue, 'q' to quit:

```
0 0x4b580000  FNSHM_464d0000  Shared address manager
```

**NOTE** The First shared memory address is `0x4B580000`, the value from this example.

```
1 0x4c580000  FNSHM_464a0000  FileNet server software
```

Total Image Services shared memory allocated: 16 MB

(This does not include segment #0)

10. Exit `ipc_tool`. If the shared memory address is correct, proceed to the next installation task. If the value is not correct, verify steps 1 - 8 above before proceeding.

## Task 24: Install Application Engine

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

### NOTES

- If you plan to install and use the IBM FileNet Workplace XT product, installing Application Engine is not required.
- Before installing Application Engine, check the latest version of the IBM *FileNet P8 4.0.x Release Notes* for known issues that might impact this software installation. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)
- Before you install Application Engine, read the following topics in this guide:
  - [“Installation Planning Considerations” on page 24](#)
  - [“Operating System Considerations” on page 25](#)
  - [“Security Considerations” on page 28](#)

Also, make sure your installation location meets the requirements specific for Application Engine outlined in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)

- The installer creates the folder structure and files needed for Application Engine.
- Switching from WAR file to EAR file deployment.

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, to add the required files to your setup. If this change of deployment type seems likely for your setup we suggest that you install Application Engine as if you would deploy an EAR file, which will create both WAR and EAR files, and then use the WAR file to deploy your web application.

- After installing, you must configure and start Application Engine application. See [“Configure Application Engine.” on page 355.](#)
- (Highly Available installations) To install Application Engine in a web farm or clustered environment, follow the instructions in the *IBM FileNet P8 Platform High Availability Technical Notice*. The document outlines the required HA install procedure and references this guide for detailed installation and deployment instructions. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)
- To ensure proper functionality and performance only install one instance of Application Engine per application server (or virtual machine or WebSphere LPAR). You can, however, deploy multiple instances of a single Application Engine version per application server, see [“Deploy Multiple Application Engine Instances” on page 443.](#)



- Before logging on to Workplace for the first time, at least one object store must exist on the Content Engine to hold the site preferences. See [“Create Object Stores” on page 258](#) for more information.
- You must have installed and configured a supported application server for Application Engine. Refer to the appropriate topic for your server type:
  - [“Configure an Application Server for Application Engine \(WebSphere\)” on page 123](#)
  - [“Configure an Application Server for Application Engine \(WebLogic\)” on page 124](#)
  - [“Configure an Application Server for Application Engine \(JBoss\)” on page 126](#)
- If you run the installer to upgrade Application Engine, the installer verifies that the currently installed version of Application Engine can be upgraded. See [“Upgrade Application Engine” on page 551](#) for more information.

### To install the Application Engine software

---

1. Log on to the application server:

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

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

2. Start the installation process.

- To install Application Engine interactively:
  - i. Access the IBM FileNet Application Engine 4.0.1 Service Pack installation software.
  - ii. Launch the appropriate Setup program (P8AE-4.0.1-<operating\_system>.bin/.exe) and continue with [Step 3 on page 346](#) below.
- To install Application Engine silently:
  - i. Access the IBM FileNet Application Engine 4.0.1 installation software package, and copy the appropriate **AE\_silent\_input.txt** or **AE\_silent\_input\_UNIX.txt** file to a local directory.
  - ii. Follow the instructions in the silent input file to edit the file to reflect the appropriate responses for your installation.

**CAUTION** If you are modifying the silent input file to perform an upgrade from AE 3.5.x to AE 4.0.1 you must modify all instances of *<AE\_install\_path>* in the script as follows:

- UNIX
 

```
Change ../FileNet/AE to ../FileNet
```
- Windows
 

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

iii. From a command prompt, navigate to, and execute the installer, then continue with [Step 4 “\(UNIX only\) Log out and log in again to set the environment variables.” on page 354.](#)

– For UNIX:

```
./P8AE-4.0.1-<operating system>.bin -silent -options  
<path_to_edited_input_file>/AE_silent_input_UNIX.txt
```

– For Windows:

```
P8AE-4.0.1-Win.exe -silent -options  
<path_to_edited_input_file>\AE_silent_input.txt
```

3. Complete the Setup screens as follows:

In this screen...	Perform this action...
License Agreement	Review and accept the license agreement.
<b>NOTE</b> This screen appears only when you use the installer to perform a fresh install.	
Specify Installation Location	<p>For the Directory Name field, enter or browse to the location where you want to install the Application Engine software <b>&lt;AE_install_path&gt;</b>, or accept the default location:</p> <ul style="list-style-type: none"> <li>• UNIX - <b>/opt/FileNet/AE/</b></li> <li>• Windows - <b>C:\Program Files\FileNet\AE\</b></li> </ul> <p>The installation program installs the Application Engine software in this directory.</p> <p><b>NOTES</b></p> <ul style="list-style-type: none"> <li>• The installer will use the <b>&lt;AE_install_path&gt;</b> to place a number of other files in default locations. See:                             <ul style="list-style-type: none"> <li>– <a href="#">“Specify Configuration Directory” on page 351</a></li> <li>– <a href="#">“Logfile Location” on page 352.</a></li> <li>– <a href="#">“Configure User Token Security” on page 353</a></li> </ul> </li> <li>• On an UPGRADE the default <b>&lt;AE_install_path&gt;</b> will be:                             <ul style="list-style-type: none"> <li>– UNIX - <b>/opt/FileNet/</b></li> <li>– Windows - <b>C:\Program Files\FileNet\</b></li> </ul> </li> <li>• If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the <b>/FileNet/AE</b> part of the path.</li> </ul>

In this screen...	Perform this action...
Choose the Installation Type	<p>Select the installation type.</p> <ul style="list-style-type: none"> <li>• Select Typical to install these components:               <ul style="list-style-type: none"> <li>– Workplace web application</li> </ul> </li> <li>• Select Custom to select the individual components to install:               <ul style="list-style-type: none"> <li>– Workplace web application</li> <li>– Workplace Source Code (for custom application development)</li> </ul> </li> </ul>
Select Components <b>NOTE</b> This screen is displayed only if you selected Custom as the installation type.	<p>Select the components you want to install.</p> <ul style="list-style-type: none"> <li>• Web Application</li> <li>• Workplace Source Code</li> </ul>
Verify Upgrade <b>NOTE</b> This screen appears only when you use the installer to perform an upgrade.	<p>The Setup program detects supported older versions of Application Engine.</p> <p>If the installer reports that no supported version of Application Engine exists on your server or if you don't want to upgrade your Application Engine at this time, click <b>Cancel</b> to exit Setup.</p>
Choose an Application Server	Select an application server and version from the drop-down lists.

In this screen...	Perform this action...
Content Engine API Configuration	<p>Configure the Content Engine API, as follows:</p> <ol style="list-style-type: none"> <li>a. Transport Method                             <p>Select EJB (Enterprise Java Beans) from the drop down list.</p> </li> <li>b. Content Engine Client Software URL:                             <p>Replace the sample server name and port number (<i>CEserver:2809</i>) with the Content Engine server name and port number for your Content Engine server. For information, see <a href="#">"IBM FileNet P8 Port Numbers" on page 643</a>.</p> <p><b>NOTE</b> To verify the correct port to use, navigate to the ports section on the application server where Content Engine is deployed and check the BOOTSTRAP_ADDRESS port.</p> <p><b>NOTE</b> To change the Content Engine name later, or to connect to a different Content Engine, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.</p> </li> <li>c. Content Engine Upload URL                             <p>Replace the sample server name and port number (<i>CEserver:2809</i>) with the Content Engine server name and port number to use when uploading document content to the Content Engine server.</p> </li> <li>d. Content Engine Download URL                             <p>Replace the sample server name and port number (<i>CEserver:2809</i>) with the Content Engine server name and port number for your Content Engine server from which to download document content.</p> </li> </ol>
Component Manager URL	Enter the URL for the Component Manager on Application Engine.

In this screen...	Perform this action...
Choose Workplace Deployment Type	<p>Select the type of deployment supported by your application server.</p> <ul style="list-style-type: none"> <li>• Select Deploy as WAR file if your server can deploy a WAR file.</li> <li>• Select Deploy as EAR file if your server can deploy an EAR file. Enter an <b>Application Name</b> for the deployed application. The default is Workplace.</li> </ul> <p><b>NOTE</b> If you decide to deploy Application Engine as a WAR file and later decide to redeploy as an EAR file you must uninstall Application Engine and then reinstall the application, selecting EAR file deployment. For information, see <a href="#">“Switching from WAR file to EAR file deployment.”</a> on page 344.</p>
Choose Authentication Method	<p>Select the authentication method for use at your site.</p> <ul style="list-style-type: none"> <li>• Application-Managed Authentication uses authentication specific to the application and does not share credentials.</li> <li>• Container-Managed Authentication provides the ability to use single sign-on (SSO) capabilities to share credentials between Application Engine and custom applications.</li> </ul> <p>When you select Container-Managed Authentication, Setup installs a sample log-in application, and modifies the <b>web.xml</b> file to support SSO. You will need to perform additional configuration for SSO after Setup is finished.</p>

In this screen...	Perform this action...
Specify Documentation URL	<p>For the documentation URL, enter the Documentation Server URL, which is where the IBM FileNet P8 Platform Documentation is installed, then click <b>Next</b>.</p> <p>Your entry must be in the following format:</p> <pre>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;/</pre> <p>where:</p> <p><i>docserver</i> is the name of the Java application server.</p> <p><i>port#</i> is the port number.</p> <p><i>ecm_help</i> is the root folder of the documentation website. You can use multi-part root folders (for example, <b>/docs/ecm_help</b>) if your application server supports them.</p> <p>See <a href="#">“Install IBM FileNet P8 Platform Documentation (WebSphere)” on page 130</a>, <a href="#">“Install IBM FileNet P8 Platform Documentation (WebLogic)” on page 135</a>, or <a href="#">“Install IBM FileNet P8 Platform Documentation (JBoss)” on page 141</a> for more information.</p> <p><b>NOTE</b> For information on how to reconfigure the Documentation URL after installation is completed, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs &gt; Bootstrap properties</a>.</p>

In this screen...	Perform this action...
Specify Configuration Directory	<p>Accept the default location or browse to the location where you want to store the configuration files.</p> <p>The default location for the configuration files is in a separate <b>Config/AE</b> directory one level up from the <b>&lt;AE_install_path&gt;</b> directory selected earlier.</p> <p>Default location of the configuration directory:</p> <ul style="list-style-type: none"> <li>• UNIX - <b>/opt/FileNet/Config/AE/</b></li> <li>• Windows - <b>C:\Program Files\FileNet\Config\AE\</b></li> </ul> <p><b>CAUTION</b> A UNC admin share (for example, <b>\\server\C\$</b>) for a shared location is not supported. You can use an ordinary file share.</p> <p><b>NOTES</b></p> <ul style="list-style-type: none"> <li>• If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the <b>/FileNet/Config/AE</b> part of the path.</li> <li>• The configuration files for an EAR file deployment, a web farm, or a clustered environment must be located in a shared folder that is accesible by all copies of the Workplace application. For more information, see the <i>IBM FileNet P8 Platform High Availability Technical Notice</i>. To download this guide from the IBM support page, see "<a href="#">Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs</a>" on page 21.</li> </ul>
Specify Upload Location	<p>Select the Upload Directory.</p> <p>The Upload Directory is the directory used by Workplace to store temporary copies of files uploaded to Workplace.</p> <p>Accept the default option or browse for a directory to hold the temporary upload files.</p> <p><b>CAUTION</b> A UNC admin share (for example, <b>\\server\C\$</b>) for a shared upload directory location is not supported. You can use an ordinary file share.</p>

In this screen...	Perform this action...
Specify Download Location	<p>Select the Download Directory.</p> <p>The Download Directory is the directory used by Workplace to store temporary copies of files downloaded from Workplace.</p> <p>Accept the default option or browse for a directory to hold the temporary download files.</p> <p><b>CAUTION</b> A UNC admin share (for example, <code>\\server\C\$</code>) for a shared download directory is not supported. You can use an ordinary file share.</p>
Logfile Location	<p>Select the Log files directory.</p> <p>The Log files is the directory used by the installer to store the <code>app_engine_install_log_401.txt</code> log file.</p> <p>Accept the default option or browse for a directory.</p> <p>The default location for the log files is in a separate <b>Logs</b> directory one level up from the <code>&lt;AE_install_path&gt;</code> directory selected earlier.</p> <p>Default location of the logs directory:</p> <ul style="list-style-type: none"> <li>• UNIX - <code>/opt/FileNet/Logs/</code></li> <li>• Windows - <code>C:\Program Files\FileNet\Logs\</code></li> </ul> <p><b>NOTE</b> If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the <code>/FileNet/Logs</code> part of the path.</p>



In this screen...	Perform this action...
Configure User Token Security	<p>Configure user token security.</p> <ol style="list-style-type: none"> <li>If needed, select the check box to create maximum strength (448-bit) keys. By default Setup creates limited strength (128-bit) keys.</li> <li>Enter the number of keys to use. <b>NOTE</b> Security generally increases with the number of keys used.</li> <li>Make a note of the user token crypto key path. The <b>UTCryptoKeyFile.properties</b> file contains the user token cryptography key used by IBM FileNet P8 applications to launch into each other without the need for additional login. The default location for the User Token Crypto Key file is in a separate <b>Authentication</b> directory one level up from the <b>&lt;AE_install_path&gt;</b> directory selected earlier. Default location of the Authentication directory:           <ul style="list-style-type: none"> <li>UNIX - <b>/opt/FileNet/Authentication/</b></li> <li>Windows - <b>C:\Program Files\FileNet\Authentication\</b></li> </ul> <b>NOTE</b> If you select a custom install location it is recommended to follow the same directory structure as seen in a typical install and retain the <b>/FileNet/Authentication</b> part of the path.   <b>CAUTION</b> For multiple applications to pass user tokens to each other, each participating application <b>must</b> use the same encryption key file. Copy the <b>UTCryptoKeyFile.properties</b> file installed with Application Engine to all servers that are hosting a token-sharing application.             For information, see the IBM FileNet P8 Developer Help topic <a href="#">Developer Help &gt; Workplace Integration and Customization Introduction &gt; User Tokens &gt; Configuring Applications to Use Tokens</a>.         </li> </ol>
User Selections	<p>Review the list of selections you have made.</p> <p><b>NOTE</b> More than one panel may be needed to display your selections.</p>
Please Read the Summary Information Below	<p>Verify your selections, and click <b>Install</b>.</p>

In this screen...	Perform this action...
Completing Application Engine Setup	Click <b>Finish</b> to complete the installation.

4. (UNIX only) Log out and log in again to set the environment variables.
5. View the **app\_engine\_install\_log\_4\_0\_1.txt** file located, located in **<AE\_install\_path>/Logs**.  
 Verify that no errors or failures were logged. Correct any errors before you proceed.

6. (Solaris only) Set the anon ports.

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

When Solaris first starts up, it takes the first several ports, called anonports, for its communication daemons. By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet uses several ports higher than 32768. See ["IBM FileNet P8 Port Numbers" on page 643](#) for details on which ports IBM FileNet uses.

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

- a. Determine the current `tcp_smallest_anon_port` setting.

From a command prompt, enter the following:

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

- b. Enable port 32777.

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

- Solaris 9

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

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

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

- Solaris10

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

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

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

- c. Reboot the Application Engine server.

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

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

7. (UNIX only) Verify that the P8TASKMAN\_HOME system environment variable is set to the following:

```
P8TASKMAN_HOME=<AE_install_path>/CommonFiles
```

8. Configure Application Engine.

Follow the instructions for your application server to configure Application Engine:

- [Task 25a “Configure Application Engine \(WebSphere\)” on page 356](#)
- [Task 25b “Configure Application Engine \(WebLogic\)” on page 373](#)
- [Task 25c “Configure Application Engine \(JBoss\)” on page 381](#)

## Task 25a: Configure Application Engine (WebSphere)

This topic covers the configuration of your Application Engine web application (Workplace) on WebSphere. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using WebSphere with container-managed authentication or SSO, edit **web.xml**. See [“To edit web.xml for container-managed authentication or SSO” on page 356](#).
2. If you are using SSO, edit **web.xml**. See [“\(SSO Only\) To edit web.xml for SSO \(optional\)” on page 359](#).
3. Configure the Application Engine. See [“To configure Application Engine \(WebSphere 5.1.x\)” on page 361](#) or [“To configure Application Engine \(WebSphere 6.1\)” on page 368](#).
4. Configure the server ports. See [“To configure the server ports” on page 371](#).

### To edit web.xml for container-managed authentication or SSO

---

**NOTE** Perform this procedure only if your site uses WebSphere with container-managed authentication or Single Sign-On (SSO). If you are using SSO, you must perform additional configuration steps as directed at the end of this procedure.

1. Make a back-up copy of **web.xml**.

**<AE\_install\_path>/Workplace/WEB-INF/web.xml**

2. Edit **web.xml**.

**NOTE** Text in **bold** in the examples below indicates changes made to the original web.xml file.

- a. Search for the parameter challengeProxyEnabled and set it to false.

```
<param-name>challengeProxyEnabled</param-name>
<param-value>false</param-value>
```

- b. Search for the first instance of <web-resource-collection>, and uncomment the url-pattern as noted in the file comments below.

```
<web-resource-collection>
  <web-resource-name>action</web-resource-name>
  <description>Define the container secured resource</description>
  <url-pattern>/containerSecured/*</url-pattern>
```

```
<!--
```

```
    Uncomment this section if all resources that require credentials
    must be secured in order to obtain a secured Thread. If using WebSphere,
    this section must be
        uncommented.
```

```
--> Move this commenting tag here from just before the </web-resource-
collection> closing tag below.
```

```
<url-pattern>/containerSecured/*</url-pattern>
<url-pattern>/</url-pattern>
<url-pattern>/author/*</url-pattern>
<url-pattern>/Browse.jsp</url-pattern>
<url-pattern>/eprocess/*</url-pattern>
<url-pattern>/Favorites.jsp</url-pattern>
<url-pattern>/GetPortalSitePreferences.jsp</url-pattern>
```

```

<url-pattern>/GetTokenSignIn.jsp</url-pattern>
<url-pattern>/GetUserInformation.jsp</url-pattern>
<url-pattern>/GetUserToken.jsp</url-pattern>
<url-pattern>/HomePage.jsp</url-pattern>
<url-pattern>/IntegrationWebBasedHelp.jsp</url-pattern>
<url-pattern>/is/*</url-pattern>
<url-pattern>/operations/*</url-pattern>
<url-pattern>/portlets/Author/edit.jsp</url-pattern>
<url-pattern>/portlets/Author/portlet.jsp</url-pattern>
<url-pattern>/portlets/Browse/edit.jsp</url-pattern>
<url-pattern>/portlets/Browse/portlet.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/edit.jsp</url-pattern>
<url-pattern>/portlets/ExternalUrl/portlet.jsp</url-pattern>
<url-pattern>/portlets/GroupPageDesign.jsp</url-pattern>
<url-pattern>/portlets/GroupPageSettings.jsp</url-pattern>
<url-pattern>/portlets/Inbox/edit.jsp</url-pattern>
<url-pattern>/portlets/Inbox/portlet.jsp</url-pattern>
<url-pattern>/portlets/MultiPagesDesign.jsp</url-pattern>
<url-pattern>/portlets/OrganizePages.jsp</url-pattern>
<url-pattern>/portlets/PortalPageDesign.jsp</url-pattern>
<url-pattern>/portlets/PortalPageInfo.jsp</url-pattern>
<url-pattern>/portlets/PortletAlias.jsp</url-pattern>
<url-pattern>/portlets/PortletSettings.jsp</url-pattern>
<url-pattern>/portlets/PreviewAndSetup.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/edit.jsp</url-pattern>
<url-pattern>/portlets/PublicQueue/portlet.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/edit.jsp</url-pattern>
<url-pattern>/portlets/QuickSearch/portlet.jsp</url-pattern>
<url-pattern>/portlets/Workflows/edit.jsp</url-pattern>
<url-pattern>/portlets/Workflows/portlet.jsp</url-pattern>
<url-pattern>/properties/*</url-pattern>
<url-pattern>/redirect/*</url-pattern>
<url-pattern>/regions/*</url-pattern>
<url-pattern>/Search.jsp</url-pattern>
<url-pattern>/select/*</url-pattern>
<url-pattern>/SelectReturn.jsp</url-pattern>
<url-pattern>/Tasks.jsp</url-pattern>
<url-pattern>/UI-INF/*</url-pattern>
<url-pattern>/utils/*</url-pattern>
<url-pattern>/WcmAdmin.jsp</url-pattern>
<url-pattern>/WcmAuthor.jsp</url-pattern>
<url-pattern>/WcmBootstrap.jsp</url-pattern>
<url-pattern>/WcmCloseWindow.jsp</url-pattern>
<url-pattern>/WcmDefault.jsp</url-pattern>
<url-pattern>/WcmError.jsp</url-pattern>
<url-pattern>/WcmJavaViewer.jsp</url-pattern>
<url-pattern>/WcmObjectBookmark.jsp</url-pattern>
<url-pattern>/WcmPortletHelp.jsp</url-pattern>
<url-pattern>/WcmPortletSearch.jsp</url-pattern>
<url-pattern>/WcmQueueBookmark.jsp</url-pattern>
<url-pattern>/WcmSignIn.jsp</url-pattern>
<url-pattern>/WcmSitePreferences.jsp</url-pattern>
<url-pattern>/WcmUserPreferences.jsp</url-pattern>
<url-pattern>/WcmWorkflowsBookmark.jsp</url-pattern>
<url-pattern>/wizards/*</url-pattern>
<url-pattern>/Author/*</url-pattern>
<url-pattern>/axis/*.jws</url-pattern>
<url-pattern>/Browse/*</url-pattern>
<url-pattern>/ceTunnel</url-pattern>
<url-pattern>/CheckoutList/*</url-pattern>
<url-pattern>/downloadMultiTransferElement/*</url-pattern>
<url-pattern>/ExternalUrl/*</url-pattern>
<url-pattern>/findRecordTarget</url-pattern>
<url-pattern>/formCallback/*</url-pattern>
<url-pattern>/getAnnotSecurity/*</url-pattern>
<url-pattern>/getCEAnnotations/*</url-pattern>
    
```

```

<url-pattern>/getContent/*</url-pattern>
<url-pattern>/getForm/*</url-pattern>
<url-pattern>/getISAnnotations/*</url-pattern>
<url-pattern>/getISAnnotSecurity/*</url-pattern>
<url-pattern>/getISContent/*</url-pattern>
<url-pattern>/getMultiContent/*</url-pattern>
<url-pattern>/getPreview</url-pattern>
<url-pattern>/getProcessor/*</url-pattern>
<url-pattern>/getRealms/*</url-pattern>
<url-pattern>/getUsersGroups/*</url-pattern>
<url-pattern>/Inbox/*</url-pattern>
<url-pattern>/integrationCommandProxy</url-pattern>
<url-pattern>/integrationResponse</url-pattern>
<url-pattern>/integrationResponseProxy</url-pattern>
<url-pattern>/integrationWebBasedCommand</url-pattern>
<url-pattern>/keepAlive</url-pattern>
<url-pattern>/launch/*</url-pattern>
<url-pattern>/PublicQueue/*</url-pattern>
<url-pattern>/putContent/*</url-pattern>
<url-pattern>/QuickSearch/*</url-pattern>
<url-pattern>/signingServlet/*</url-pattern>
<url-pattern>/transport/*</url-pattern>
<url-pattern>/upload/*</url-pattern>
<url-pattern>/vwsimsoapervlet</url-pattern>
<url-pattern>/vwsoaprouter</url-pattern>
<url-pattern>/Workflows/*</url-pattern> Move the closing comment tag from
here to the location indicated at the beginning of this example.
</web-resource-collection>
    
```

- c. Locate the section <auth-constraint>, comment the wild-card (\*) role-name as noted in the file comments below.

```

<auth-constraint>
  <!-- <role-name>*</role-name> -->
  <!-- For WebSphere 6, use the role-name line below instead of the wildcard role
  above.-->

  <role-name>All Authenticated</role-name>

  <!-- For WebSphere 6, add this security-role element below the login-config
  element (below).
  <security-role>
    <description>All Authenticated</description>
    <role-name>All Authenticated</role-name>
  </security-role>
  -->
</auth-constraint>
    
```

- d. Locate the end of the </login-config> element, and add the All Authenticated users role-element after the closing tag.

```

<security-role>
  <description>All Authenticated</description>
  <role-name>All Authenticated</role-name>
</security-role>
    
```

- e. 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>
-->

```

- f. Search for the first instance of a <security-constraint> tag, and add the following <security-constraint> tag before that tag.

**CAUTION** Enter the information below as single lines without line breaks.

```

<security-constraint>
  <web-resource-collection>
    <web-resource-name>action</web-resource-name>
    <description>Define the non-secured resource</description>
    <url-pattern>/P8BPMWSBroker/*</url-pattern>
  </web-resource-collection>
</security-constraint>

```

3. Save your changes to **web.xml** and close the file.
4. If your site uses SSO, Continue on with [“\(SSO Only\) To edit web.xml for SSO \(optional\)” on page 359](#), otherwise continue with [“To configure Application Engine \(WebSphere 5.1.x\)” on page 361](#) or [“To configure Application Engine \(WebSphere 6.1\)” on page 368](#).

### (SSO Only) To edit web.xml for SSO (optional)

---

**NOTE** Perform this procedure only if your site uses SSO with a proxy server. You must use this procedure to modify web.xml to enable SSO.

1. Edit **web.xml**.

As needed, set the ssoProxyContextPath, ssoProxyHost, ssoProxyPort, and ssoProxySSLPort.

These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, you must set the corresponding sso<proxy host element> parameter for that element in the URL to the value for the SSO proxy host.

The default settings are (in **bold** below):

```

<init-param>
  <param-name>ssoProxyContextPath</param-name>
  <param-value></param-value>
</init-param>
<init-param>
  <param-name>ssoProxyHost</param-name>
  <param-value></param-value>
</init-param>
<init-param>
  <param-name>ssoProxyPort</param-name>
  <param-value></param-value>
</init-param>
<init-param>
  <param-name>ssoProxySSLPort</param-name>
  <param-value></param-value>
</init-param>

```

In general, the init parameters above must be configured as follows:

- **ssoProxyContextPath**: Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

For example, if the Workplace deploy host URL is  
`http://deploy_server:2809/Workplace` and the SSO proxy host URL is  
`http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

```
<param-name>ssoProxyContextPath</param-name>  
<param-value>/fn/Workplace</param-value>
```

- **ssoProxyHost**: Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

For example, if the host URL where Workplace is deployed is  
`http://deploy_server/Workplace` and the corresponding SSO proxy host URL is  
`http://sso_proxy_server/Workplace`, then use the following:

```
<param-name>ssoProxyHost</param-name>  
<param-value>sso_proxy_server</param-value>
```

- **ssoProxyPort**: Set the value to the http port on the SSO proxy host.

For example:

```
<param-name>ssoProxyPort</param-name>  
<param-value>80</param-value>
```

- **ssoProxySSLPort**: Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

For example:

```
<param-name>ssoProxySSLPort</param-name>  
<param-value>443</param-value>
```

2. Save your changes to **web.xml** and close the file.



## To configure Application Engine (WebSphere 5.1.x)

---

1. Open the WebSphere administrative console.
2. Set the JVM settings for JAAS login configuration, memory settings, and UTF-8 encoding.
  - a. Expand **Servers**.
  - b. Click **Application Servers**.
  - c. Click **<server name>**.
  - d. Under “Additional Properties,” select **Process Definition**, and then select **Java Virtual Machine**.
  - e. Set the JAAS login entry in the JVM generic argument field to one of the following (do not enter the linebreaks):

### New Install

```
-Djava.security.auth.login.config=<AE_install_path>\CE_API\config\jaas.conf.WebSphere
```

### Upgrade

```
-Djava.security.auth.login.config=<AE_install_path>\CE_API\config\jaas.conf.WebSphere
```

**CAUTION** (Windows only) On WebSphere/Windows the path cannot contain a space. You must use 8.3-notation for the install path information.

If **<AE\_install\_path>** is:

```
C:\Program Files\FileNet\AE
```

use:

```
C:\Progra~1\FileNet\AE
```

- f. Set the Initial and Maximum Heap Size.

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- g. Click **Apply** to save your changes.
- h. Click **Custom Properties**.
- i. Click **New**.
- j. In the **Name** field, enter  
`client.encoding.override`
- k. In the **Value** field, enter  
`UTF-8`
- l. Click **Apply**, click **Save**, and then click **Save Changes to the Master Configuration**.

3. Enable Application Integration Installation.
  - a. From the Administrative Console, expand **Environment**. Click **Virtual Hosts**.
  - b. Click the **default\_host** (or the host your application is deployed under).
  - c. Click **MIME Types**, then click **New**.
  - d. In the **MIME Type** field, enter **application/octet-stream**.
  - e. In the **Extension(s)** field, enter **exe**. Click **OK**.
  - f. Click **Apply**, **Save**, then **Save changes to the Master Configuration**.
4. Configure Light weight Third Party Authentication (LTPA).

**NOTE** If you are already using LTPA with your CE application server, you only need to export the existing keys per [Step vii](#) thru [Step ix](#) below.

  - a. On the Content Engine server, do the following:
    - i. Log in to the WebSphere administrative console.
    - ii. From the administrative console, Navigate to **Security > Authentication Mechanisms > LTPA**.
    - iii. Enter a password to create the *<LTPA password>*.

**NOTE** For password restrictions, see the WebSphere documentation. If you have already configured CE for LTPA, use the existing password in the AE configuration below.
    - iv. Enter a path for the Key File Name.  
For example, **C:\LTPA\<ltpa\_key\_name>**.
    - v. Click **Generate Keys**.
    - vi. Click **Save**, and then click **Save changes to the Master Configuration**.
    - vii. Navigate to **Security > Authentication Mechanisms > LTPA**.
    - viii. Click **Export Keys**.
    - ix. Copy the key file.  
Copy the file from the location you specified above to a directory on the Application Engine server. On Windows, for example, **C:\LTPA\<ltpa\_key\_name>**.
    - x. Navigate to **Security > Global Security**.
    - xi. From the **Active Authentication Mechanism** drop down list, select **LTPA (Light weight Third Party Authentication)**.
    - xii. From the **Active User Registry** drop down list, select **LDAP**.
    - xiii. Click **Apply**, and then click **Save changes to the Master Configuration**.
    - xiv. Stop and restart WebSphere.
  - b. On the Application Engine server, do the following:

- i. Log in to the WebSphere administrative console.
- ii. Navigate to **Security > Authentication Mechanisms > LTPA**.
- iii. Enter a value for the timeout that is larger than the default.  
 For example, if the timeout value is set to 2 hours, the LTPA key expires and end users will not be able to log in to Workplace after having been logged in for 2 hours.
- iv. Enter the *<LTPA password>* you created for Content Engine above.
- v. Enter the path for the key file that you copied to the Application Engine server. For example, **C:\LTPA\<ltpa\_key\_name>**.
- vi. Click **Import Keys**.
- vii. Navigate to **Security > User Registries > LDAP**.
- viii. Configure the LDAP provider to exactly match the settings from the Content Engine server.
  - Server User ID
  - Server User Password
  - Type
  - Host
  - Port
  - Base distinguished name (DN)
  - Bind distinguished name (DN)
  - Bind password
- ix. Click **Apply**.
- x. Click **Advanced LDAP Settings**.
- xi. Configure the settings here to exactly match the settings from the Content Engine server. For example, Group Filter and Group ID Map must match the Content Engine settings.
  - User filter
  - Group Filter
  - User ID map
  - Group member ID map
  - Certificate map mode
  - Certificate filter
- xii. Navigate to **Security > Global Security**.
- xiii. Select (select) **Enabled** flag.

xiv. Turn off (deselect) **Enforce Java 2 Security**.

**NOTE** IBM FileNet P8 platform utilizes LDAP-based security, and does not support Java 2 security. Enforcing Java 2 security will cause unexpected behavior.

xv. From the **Active Authentication Mechanism** drop down list, select **LTPA (Light weight Third Party Authentication)**.

xvi. From the **Active User Registry** drop down list, select **LDAP**.

xvii. Click **Apply**, and then click **Save changes to the Master Configuration**.

5. Stop and restart WebSphere.

6. Set permissions for the user running the application server.

If the local user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the following (default) folders:

**<WAS\_HOME>/profiles/default/installedApps/<node\_name>/app\_engine\_war.ear/  
app\_engine.war**

**<AE\_install\_path>**

7. Continue with “[To configure the server ports](#)” on page 371.

### **To configure Application Engine (WebSphere 6.0)**

---

1. Open the WebSphere administrative console.

2. Set JVM settings for JAAS login configuration, memory settings, and UTF-8 encoding.

a. Expand **Servers**.

b. Click **Application Servers**.

c. Click **<server name>**.

d. Under “Server Infrastructure”, select **Java & Process Management**.

e. Select **Process Definition**, and then select **Java Virtual Machine**.

f. Set the JAAS login entry in the JVM generic argument field to one of the following (do not enter the linebreaks):

```
-Djava.security.auth.login.config=<ae_install_path>\CE_API\config  
\jaas.conf.WebSphere
```

**CAUTION** (Windows only) On WebSphere/Windows the path cannot contain a space. You must use 8.3-notation for the install path information.

If **<AE\_install\_path>** is:

```
C:\Program Files\FileNet\AE
```

use:

```
C:\Progra~1\FileNet\AE
```

- g. Set the Initial and Maximum Heap Size.  
Refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).
  - h. Click **Apply**.
  - i. Click **Custom Properties**.
  - j. Click **New**.
  - k. Verify that the Name field is set to:  
`client.encoding.override`
  - l. Verify that the Value field is set to:  
`UTF-8`
  - m. Click **Apply**, click **Save**, and then click **Save Changes to the Master Configuration**.
3. Verify Application Integration.
    - a. From the Administrative Console, expand **Environment**. Click **Virtual Hosts**.
    - b. Click the **default\_host** (or the host your application is deployed under).
    - c. Click **MIME Types**.
    - d. Verify that MIME type is set to application/octet-stream or use the following steps to set it.
      - i. Click **New**.
      - ii. In the **MIME Type** field, enter **application/octet-stream**.
      - iii. In the **Extension(s)** field, enter **exe**. Click **OK**.
      - iv. Click **Apply**, click **Save**, and then **Save changes to the Master Configuration**.
  4. Configure Light weight Third Party Authentication (LTPA).
    - a. On the Content Engine server, do the following:  
**NOTE** If you are already using LTPA with your CE application server, you only need to export the existing keys per [Step viii](#) thru [Step viii](#) below.
      - i. Log in to the Administrative Console.
      - ii. Navigate to **Security > Global Security**.
      - iii. From the right side of the panel, select **Authentication > Authentication Mechanisms > LTPA**.
      - iv. Enter a password to create the **<LTPA password>**.  
**NOTE** For password restrictions, see the WebSphere documentation. If you have already configured CE for LTPA, use the existing password in the AE configuration below.

- v. Enter a path for the Key File Name. For example, **C:\LTPA\<ltpa\_key\_name>**.
  - vi. Click **Generate Keys**.
  - vii. Click **Save**, and then click **Save changes to the Master Configuration**.
  - viii. Navigate to **Security > Global Security**.
  - ix. From the right side of the panel, select **Authentication > Authentication Mechanisms > LTPA**.
  - x. Click **Export Keys**.
  - xi. Copy the key file.

Copy the file from the location you specified above to a directory on the Application Engine server. On Windows, for example, **C:\LTPA\<ltpa\_key\_name>**.
  - xii. Navigate to **Security > Global Security**.
  - xiii. From the **Active Authentication Mechanism** drop down list, select **LTPA (Light weight Third Party Authentication)**.
  - xiv. From the **Active User Registry** drop down list, select **LDAP**.
  - xv. Click **Apply**, and then click **Save changes to the Master Configuration**.
  - xvi. Stop and restart WebSphere.
- b. On the Application Engine server, do the following:
- i. Log in to the Administrative Console.
  - ii. Navigate to **Security > Global Security**.
  - iii. From the right side of the panel, select **Authentication > Authentication Mechanisms > LTPA**.
  - iv. Enter a value for the timeout that is larger than the default.

For example, if the timeout value is set to 2 hours, the LTPA key expires and end users will not be able to log in to Workplace after having been logged in for 2 hours.
  - v. Enter the **<LTPA password>** you created for Content Engine above.
  - vi. Enter the path for the key file that you copied to the Application Engine server. For example, **C:\LTPA\<ltpa\_key\_name>**.
  - vii. Click **Import Keys**.
  - viii. Navigate to **Security > Global Security**.
  - ix. From the right side of the panel, select **LDAP User Registry**.
  - x. Configure the LDAP provider to exactly match the corresponding settings on the Content Engine application server.
    - Server User ID
    - Server User Password

- Type
  - Host
  - Port
  - Base distinguished name (DN)
  - Bind distinguished name (DN)
  - Bind password
- xi. Click **Apply**.
- xii. Click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
- xiii. Configure the settings here to exactly match the corresponding settings from the Content Engine application server.
- User filter
  - Group Filter
  - User ID map
  - Group member ID map
  - Certificate map mode
  - Certificate filter
- xiv. Navigate to **Security > Global Security**.
- xv. Turn on (check) **Enabled** flag.
- xvi. Turn off (uncheck) **Enforce Java 2 Security**.
- NOTE** The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.
- xvii. From the **Active Authentication Mechanism** drop down list, select **LTPA (Light weight Third Party Authentication)**.
- xviii. From the **Active User Registry** drop down list, select **LDAP**.
- xix. Click **Apply**, and then click **Save changes to the Master Configuration**.
5. Set permissions for the user running the application server.
- The user that will be running the application server must have read/write permissions on the following (default) folders:
- <WAS\_HOME>/profiles/default/installedApps/<node\_name>/app\_engine\_war.ear/  
 app\_engine.war**
- <AE\_install\_path>**
6. Continue with [“To configure the server ports” on page 371](#).

## To configure Application Engine (WebSphere 6.1)

---

1. Open the WebSphere administrative console.
2. Set JVM settings for JAAS login configuration, memory settings, and UTF-8 encoding.
  - a. Expand **Servers**.
  - b. Click **Application Servers**.
  - c. Click **<server name>**.
  - d. Under “Server Infrastructure”, select **Java & Process Management**.
  - e. Select **Process Definition**, and then select **Java Virtual Machine**.
  - f. Set the JAAS login entry in the JVM generic argument field to one of the following (do not enter the linebreaks):

```
-Djava.security.auth.login.config=<ae_install_path>\CE_API\config  
\jaas.conf.WebSphere
```

**CAUTION** (Windows only) On WebSphere/Windows the path cannot contain a space. You must use 8.3-notation for the install path information.

If **<AE\_install\_path>** is:

```
C:\Program Files\FileNet\AE
```

use:

```
C:\Progra~1\FileNet\AE
```

- g. Set the Initial and Maximum Heap Size.

Refer to your application server vendor's recommendation for Initial and Maximum heap size values. For IBM specific recommendations, see the *IBM FileNet P8 Platform Performance Tuning Guide*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
    - h. Click **Apply**.
    - i. Click **Custom Properties**.
    - j. Click **New**.
    - k. Verify that the Name field is set to:

```
client.encoding.override
```
    - l. Verify that the Value field is set to:

```
UTF-8
```
    - m. Click **Apply**, click **Save**, and then click **Save Changes to the Master Configuration**.
  3. Verify Application Integration.
    - a. From the Administrative Console, expand **Environment**. Click **Virtual Hosts**.
    - b. Click the **default\_host** (or the host your application is deployed under).



- c. Click **MIME Types**.
  - d. Verify that MIME type is set to application/octet-stream or use the following steps to set it.
    - i. Click **New**.
    - ii. In the **MIME Type** field, enter **application/octet-stream**.
    - iii. In the **Extension(s)** field, enter **exe**. Click **OK**.
    - iv. Click **Apply**, click **Save**, and then **Save changes to the Master Configuration**.
4. Configure Light weight Third Party Authentication (LTPA).
- a. On the Content Engine server, do the following:

**NOTE** If you are already using LTPA with your CE application server, you only need to export the existing keys per [Step vi](#) thru [Step viii](#) below.

    - i. Log in to the Administrative Console.
    - ii. Navigate to **Security > Secure administration, applications, and infrastructure**.
    - iii. From the right side of the panel, select **Authentication Mechanisms and expiration**.
    - iv. In the box titled "Cross-cell single sign-on, enter a password to create the <LTPA password>.

**NOTE** For password restrictions, see the WebSphere documentation. If you have already configured CE for LTPA, use the existing password in the AE configuration below.

    - v. Enter a path for the Key File Name. For example, **C:\LTPA\<ltpa\_key\_name>**.
    - vi. Click **Export Keys**. Verify that the following message is displayed: The keys were successfully exported to the file <file name>.
    - vii. Click **OK**, and then click **Save changes to the Master Configuration**.
    - viii. Copy the key file.

Copy the file from the location you specified above to a directory on the Application Engine server. On Windows, for example, **C:\LTPA\<ltpa\_key\_name>**.
    - ix. Navigate to **Security > Global Security**.
    - x. From the **Active User Registry** drop down list, select **LDAP**.
    - xi. Click **Apply**, and then click **Save changes to the Master Configuration**.
    - xii. Stop and restart WebSphere.
  - b. On the Application Engine server, do the following:
    - i. Log in to the Administrative Console.
    - ii. Navigate to **Security > Secure administration, applications, and infrastructure**.
    - iii. From the right side of the panel, select **Authentication Mechanisms and expiration**.

- iv. Enter a value for the timeout that is larger than the default.  
 For example, if the timeout value is set to 2 hours, the LTPA key expires and end users will not be able to log in to Workplace after having been logged in for 2 hours.
- v. In the box titled cross-cell single sign-on, enter the *<LTPA password>* you created for Content Engine above. Confirm the password.
- vi. Enter the path for the key file that you copied to the Application Engine server. For example, **C:\LTPA\<ltpa\_key\_name>**.
- vii. Click **Import Keys**. Verify that the following message is displayed: The keys were successfully imported from the file *<file name>*.
- viii. Navigate to **Security > Secure administration, applications, and infrastructure**.
- ix. Turn on (check) **Enable Administrative Security flag**.
- x. Turn on (check) **Enable application security flag**.
- xi. Turn off (uncheck) **Java 2 Security**.  
**NOTE** The IBM FileNet P8 Platform utilizes LDAP-based security, and does not support Java 2 security. Enabling Java 2 security will cause unexpected behavior.
- xii. From the **Active Authentication Mechanism** drop down list, select **LTPA (Light weight Third Party Authentication)**.
- xiii. From the bottom of the panel, in the box titled "available realm definitions," select **Standalone LDAP registry** and click **Configure**.
- xiv. Configure the LDAP provider to exactly match the corresponding settings on the Content Engine application server.
  - Primary administrative user name  
 Select "Automatically generated server identity."
  - Type
  - Host
  - Port
  - Base distinguished name (DN)
  - Bind distinguished name (DN)
  - Bind password
- xv. Click **Apply**.
- xvi. Click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings**.
- xvii. Configure the settings here to exactly match the corresponding settings from the Content Engine application server.
  - User filter
  - Group Filter

- User ID map
- Group member ID map
- Certificate map mode
- Certificate filter

xviii. Save these settings

xix. Click **Test 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.

xx. Next to "Available realm definitions," make sure "Standalone LDAP registry" is still selected, and click **Set as current**.

xxi. Click **Save**.

5. Set permissions for the user running the application server.

The user that will be running the application server must have read/write permissions on the following (default) folders:

**<WAS\_HOME>\profiles\default\installedApps<node\_name>\app\_engine\_war.ear/  
 app\_engine.war**

**<AE\_install\_path>**

6. Continue with ["To configure the server ports" on page 371](#).

### **To configure the server ports**

---

This configuration is required for WebSphere 5.x and recommended for WebSphere 6.x.

1. Stop the WebSphere server.
2. Make a backup copy of **serverindex.xml** located in:

- WebSphere 5.x

**<WAS\_HOME>\profiles\default\config\cells<machine\_name>\Node01Cell\nodes\  
 <machine\_name>Node01\**

- WebSphere 6.x

**<WAS\_HOME>\profiles\default\config\cells<machine\_name>\Node01Cell\nodes\  
 <machine\_name>Node01\**

3. Edit **serverindex.xml**.

Locate the <specialEndpoints> section, and change the port numbers for the three SSL listener addresses to "0" as shown below:

```
<specialEndpoints xmi:id="NamedEndPoint_1155689929072"
endPointName="SAS_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929072" host="host_name" port="0"/>
</specialEndpoints>
<specialEndpoints xmi:id="NamedEndPoint_1155689929073"
endPointName="CSIV2_SSL_SERVERAUTH_LISTENER_ADDRESS">
  <endPoint xmi:id="EndPoint_1155689929073" host="host_name" port="0"/>
```

```
</specialEndpoints>  
<specialEndpoints xmi:id="NamedEndPoint_1155689929074"  
endPointName="CSIV2_SSL_MUTUALAUTH_LISTENER_ADDRESS">  
  <endPoint xmi:id="EndPoint_1155689929074" host="host_name" port="0"/>  
</specialEndpoints>
```

4. Save and close the **serverindex.xml** file.
5. Restart WebSphere.
6. Continue with the following task:
  - If you are performing a fresh install:  
Do [Task 26 "Install Application Engine Software Updates"](#) on page 383.
  - If you are performing an upgrade:  
Do ["Manually copy custom data."](#) on page 555.

## Task 25b: Configure Application Engine (WebLogic)

This topic covers the configuration of your Application Engine application (Workplace) on WebLogic. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

1. If you are using SSO, edit **web.xml**. See [“To edit web.xml for SSO \(optional\)”](#) on page 373.
2. Modify the application server startup script. See [“To modify the application server startup script”](#) on page 374.
3. Configure Application Engine. See [“To configure Application Engine \(WebLogic 8.1.x\)”](#) on page 376 or [“To configure Application Engine \(WebLogic 9.2\)”](#) on page 378.
4. Modify **config.xml** to support passing user credentials to clients such as Application Integration and WebDAV. See [“To enable passing user credentials to client applications”](#) on page 379.

### To edit web.xml for SSO (optional)

---

**NOTE** Perform this procedure only if your site uses SSO with a proxy server. You must modify web.xml to enable SSO.

1. Make a backup copy of **web.xml**.  
**<AE\_install\_path>/Workplace/WEB-INF/web.xml**
2. Edit **web.xml**.
  - a. Set the parameter `perimeterChallengeMode` to true, as in:

```
<init-param>
  <param-name>perimeterChallengeMode</param-name>
  <param-value>true</param-value>
</init-param>
```

- b. As needed, set the `ssoProxyContextPath`, `ssoProxyHost`, `ssoProxyPort`, and `ssoProxySSLPort`.

These parameter values are used to modify one or more elements of the native URL that Workplace sees on a request. Wherever the value of an SSO proxy host element in the URL request is different from the equivalent information for the host where Workplace is deployed, then you must set the corresponding `sso*` parameter for that element in the URL to the value for the SSO proxy host.

The default settings are (in **bold**):

```
<init-param>
  <param-name>ssoProxyContextPath</param-name>
  <param-value></param-value>
</init-param>
<init-param>
  <param-name>ssoProxyHost</param-name>
  <param-value></param-value>
</init-param>
<init-param>
  <param-name>ssoProxyPort</param-name>
  <param-value></param-value>
</init-param>
```

```
<init-param>
  <param-name>ssoProxySSLPort</param-name>
  <param-value></param-value>
</init-param>
```

In general, the init parameters above should be configured as follows:

- **ssoProxyContextPath:** Set the value to the context path of the SSO proxy host URL. This is the path portion of the URL that appears after the server name, and which represents top-level access to the Workplace application.

For example, if the Workplace deploy host URL is

`http://deploy_server:2809/Workplace` and the SSO proxy host URL is

`http://sso_proxy_server.domain.com/fn/Workplace`, then use the following:

```
<param-name>ssoProxyContextPath</param-name>
<param-value>/fn/Workplace</param-value>
```

- **ssoProxyHost:** Set the value to the SSO proxy host server name. Typically, this will be a full domain-qualified hostname.

For example, if the host URL where Workplace is deployed is

`http://deploy_server/Workplace` and the corresponding SSO proxy host URL is

`http://sso_proxy_server/Workplace`, then use the following:

```
<param-name>ssoProxyHost</param-name>
<param-value>sso_proxy_server</param-value>
```

- **ssoProxyPort:** Set the value to the http port on the SSO proxy host.

For example:

```
<param-name>ssoProxyPort</param-name>
<param-value>80</param-value>
```

- **ssoProxySSLPort:** Set the value to the https port on the SSO proxy host, if defined and/or used to access Workplace pages.

For example:

```
<param-name>ssoProxySSLPort</param-name>
<param-value>443</param-value>
```

3. Save your changes to **web.xml** and close the file.

### To modify the application server startup script

---

1. Stop the WebLogic application server if running.
2. Make a backup copy of the application server startup script.

Backup **startWebLogic.cmd** for Windows or **startWebLogic.sh** for UNIX.

**NOTE** If you are not using a WebLogic domain, backup **startWLS.cmd** for Windows or **startWLS.sh** for UNIX.

3. Edit the application server startup script.

- a. Configure JAAS login.

Add one of the following right after the classpath entry for WebLogic.

**CAUTION** The set jaas\_login entry should be entered as a single line without line breaks.

- **Windows**

```
@REM Jaas Login configuration setting
set jaas_login=%jaas_login% -Djava.security.auth.login.config=
"<AE_install_path>\CE_API\config\jaas.conf.WebLogic"
```

- **UNIX**

```
# Jaas Login configuration setting
jaas_login="{jaas_login} -Djava.security.auth.login.config=
"<AE_install_path>/CE_API/config/jaas.conf.WebLogic"
```

- **AIX**

```
# Jaas Login configuration setting
jaas_login="{jaas_login} -Djava.security.auth.login.config=
"<AE_install_path>/CE_API/config/jaas.conf.WebLogic"
jaas_login="{jaas_login} -
Dlogin.configuration.provider=com.ibm.security.auth.login.ConfigFile
```

b. Add %jaas\_login% section as indicated in the examples below in **bold**.

- **WebLogic 8 - Windows - after the CLASSPATH settings**

```
echo .
echo CLASSPATH=%CLASSPATH%
echo .
echo PATH=%PATH%
echo .
echo *****
echo * To start WebLogic Server, use a username and *
echo * password assigned to an admin-level user. For *
echo * server administration, use the WebLogic Server *
echo * console at http:\[hostname]:[port]\console *
echo *****

%JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %jaas_login% -
Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
Dweblogic.management.password=%WLS_PW% -
Dweblogic.ProductionModeEnabled=%PRODUCTION_MODE% -
Djava.security.policy="%WL_HOME%\server\lib\weblogic.policy" weblogic.Server
```

- **WebLogic 8 - UNIX - after the CLASSPATH settings**

```
# JAAS login config file
JAAS_LOGIN="{JAAS_LOGIN} -Djava.security.auth.login.config=/opt/FileNet/AE/
CE_API/config/jaas.conf.weblogic"

# Call WebLogic Server
echo "."
echo "CLASSPATH=${CLASSPATH}"
echo "."
echo "PATH=${PATH}"
echo "."
echo "*****"
echo "* To start WebLogic Server, use a username and *"
echo "* password assigned to an admin-level user. For *"
echo "* server administration, use the WebLogic Server *"
echo "* console at http://[hostname]:[port]/console *"
echo "*****"
{JAVA_HOME}/bin/java {JAVA_VM} {MEM_ARGS} {JAVA_OPTIONS} {JAAS_LOGIN} -
Dweblogic.Name={SERVER_NAME} -
```

```
Dweblogic.ProductionModeEnabled=${PRODUCTION_MODE} -
Djava.security.policy="${WL_HOME}/server/lib/weblogic.policy" weblogic.Server
```

- **WebLogic 9 - Windows** - in the WLS\_REDIRECT\_LOG settings

```
If "%WLS_REDIRECT_LOG"==" " (
    echo Starting WLS with line:
    echo %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %jaas_login% -
    -Dweblogic.Name=%SERVER_NAME% -
    Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
    %SERVER_CLASS%
    %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %jaas_login% -
    Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
    Dweblogic.management.password=%WLS_PW% -
    Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
    %SERVER_CLASS%
) else (
    echo Redirecting output from WLS window to %WLS_REDIRECT_LOG%
    %JAVA_HOME%\bin\java %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS% %jaas_login% -
    Dweblogic.Name=%SERVER_NAME% -Dweblogic.management.username=%WLS_USER% -
    Dweblogic.management.password=%WLS_PW% -
    Djava.security.policy=%WL_HOME%\server\lib\weblogic.policy %PROXY_SETTINGS%
    %SERVER_CLASS% >"%WLS_REDIRECT_LOG%" 2>&1
)
```

- **WebLogic 9 - UNIX** - in the WLS\_REDIRECT\_LOG settings

```
#{JAVA_HOME}/bin/java #{JAVA_VM} -version

if [ "${WLS_REDIRECT_LOG}" = " " ] ; then
    echo "Starting WLS with line:"
    echo "#{JAVA_HOME}/bin/java #{JAVA_VM} #{MEM_ARGS} #{JAVA_OPTIONS} -
    Dweblogic.Name=#{SERVER_NAME} -Djava.security.policy=#{WL_HOME}/server/lib/weblogic.policy
    #{PROXY_SETTINGS} #{SERVER_CLASS}"
    #{JAVA_HOME}/bin/java #{JAVA_VM} #{MEM_ARGS} #{JAVA_OPTIONS}
    #{JAAS_LOGIN} -Dweblogic.Name=#{SERVER_NAME} -Djava.security.policy=#{WL_HOME}/server/lib/weblogic.policy
    #{PROXY_SETTINGS} #{SERVER_CLASS}
else
    echo "Redirecting output from WLS window to ${WLS_REDIRECT_LOG}"
    #{JAVA_HOME}/bin/java #{JAVA_VM} #{MEM_ARGS} #{JAVA_OPTIONS}
    #{JAAS_LOGIN} -Dweblogic.Name=#{SERVER_NAME} -Djava.security.policy=#{WL_HOME}/server/lib/weblogic.policy
    #{PROXY_SETTINGS} #{SERVER_CLASS} >"${WLS_REDIRECT_LOG}" 2>&1
fi
```

4. Save and close the server startup script.

### To configure Application Engine (WebLogic 8.1.x)

1. (If you selected Container-Managed Authentication during the installation) Enable trust between WebLogic domains for the Content Engine domain and the Application Engine domain.

Do the following on both the Content Engine application server and the Application Engine application server.

- a. Log on to the WebLogic Administration Console.
- b. Click **<my\_domain>**.



- c. Click **View Domain-Wide Security Settings**.
- d. Click the **Advanced** tab.
- e. Uncheck (turn off) **Enable Generated Credential**.
- f. Enter a password for the domain in the **Credential** field. You must enter the same password for both Content Engine and Application Engine.
- g. Confirm the password by entering it in the **Confirm Credential** field.
- h. Click **Apply**.

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

2. (If you selected Container-Managed Authentication during the installation) Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
  - a. Refer to your Content Engine installation checklists and the WebLogic Administration Console settings for Compatibility Security > Realms for Authentication Provider, users, and groups on Content Engine.

Configure the LDAP provider to exactly match the settings from the Content Engine server.

    - Group Base DN:
    - User Name Attribute:
    - Port:
    - User Base DN:
    - Principal:
    - Credential:
    - Confirm Credential:
    - Host:
    - User From Name Filter:
    - Group From Name Filter:
  - b. Restart the application server.
3. Set permissions for the user running the application server.

**NOTE** On Windows, the following is only required for NTFS formatted partitions

If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed AE (**<AE\_install\_path>**), see [“Specify Installation Location” on page 346](#).

4. Continue with the following task:
  - If you are performing a fresh install:  
Do [Task 26 “Install Application Engine Software Updates” on page 383](#).
  - If you are performing an upgrade:  
Do [“Manually copy custom data.” on page 555](#).

### To configure Application Engine (WebLogic 9.2)

---

1. (If you selected Container-Managed Authentication during the installation) Enable trust between WebLogic domains for the Content Engine domain and the Application Engine domain.  
  
Do the following on both the Content Engine application server and the Application Engine application server.
  - a. Log on to the WebLogic Administration Console.
  - b. In the Change Center of the Administration Console, click **Lock & Edit**.
  - c. Click the **name** of the domain.
  - d. Click the **Security** tab.
  - e. Click **General**.
  - f. Click **Advanced**.
  - g. Enter a password for the domain in the **Credential** field. You must enter the same password for both the Content Engine domain and Application Engine domain.
  - h. Click **Save**.
  - i. Click **Activate Changes**.
  - j. Restart the server if needed.
  - k. Repeat this procedure in each domain for which you want to enable trust.
2. (If you selected Container-Managed Authentication during the installation) Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
  - a. Refer to your Content Engine installation checklists and the WebLogic Administration Console settings for Compatibility Security > Realms for Authentication Provider, users, and groups on Content Engine.  
  
Configure the LDAP provider to exactly match the settings from the Content Engine server.
    - Group Base DN:
    - User Name Attribute:
    - Port:
    - User Base DN:
    - Principal:

- Credential:
  - Confirm Credential:
  - Host:
  - User From Name Filter:
  - Group From Name Filter:
- b. Restart the application server.
3. Set permissions for the user running the application server.
- If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed AE (`<AE_install_path>`), see [“Specify Installation Location” on page 346](#).
4. Continue with [“To enable passing user credentials to client applications” on page 379](#).

### **To enable passing user credentials to client applications**

---

Perform this procedure to enable passing user credentials between Application Engine and its client applications such as WebDAV and Application Integration.

**CAUTION** If you do not make this change to `config.xml`, then end users will be prompted to enter their user name and password to complete any client operations, such as adding a document.

1. Stop the WebLogic server.
2. Make a backup copy of `config.xml` located in deployment directory.

For example:

WebLogic 8.x

`<BEA_home>/bea/user_projects/domains/<domain_name>/config.xml`

WebLogic 9.x

`<BEA_home>/bea/user_projects/domains/<domain_name>/config/config.xml`

3. Edit `config.xml`.

**CAUTION** The `enforce-valid-basic-auth-credentials` entry should be entered as a single line without line breaks.

- a. Locate the `<security-configuration>` section and add the following line to the end of the section, just before the `</security-configuration>` tag:

```
<enforce-valid-basic-auth-credentials>>false</enforce-valid-basic-auth-credentials>
```

- b. Save your changes to `config.xml` and close the file.
4. Restart WebLogic.

5. Continue with the following task:
  - If you are performing a fresh install:  
Do [Task 26 “Install Application Engine Software Updates”](#) on page 383.
  - If you are performing an upgrade:  
Do [“Manually copy custom data.”](#) on page 555.

## Task 25c: Configure Application Engine (JBoss)

This topic covers the configuration of your Application Engine application (Workplace) on JBoss. Perform the following high-level steps in the order listed, using the referenced detailed procedures for each step.

### To modify the application server startup script

---

1. Stop the JBOSS application server if running.
2. Make a backup copy of the application server startup script.

Backup **run.bat** (Windows) or **run.sh** (UNIX).

3. Configure JAAS login.
  - a. Open the application server startup script for editing.
    - (UNIX) **run.sh**
    - (Windows) **run.bat**

- b. Configure JAAS login.

Add one of the following right after the RESTART entry in the startup script.

**CAUTION** Enter the `jaas_login` entry (bold below) as a single line without line breaks.

- Windows

```
:RESTART

"%JAVA%" %JAVA_OPTS% "-Djava.security.auth.login.config=C:\Program
Files\FileNet\AE\CE_API\config\jaas.conf.JBoss" "-
Djava.endorsed.dirs=%JBOSS_ENDORSED_DIRS%" -classpath "%JBOSS_CLASSPATH%"
org.jboss.Main %*
```

- UNIX

```
:RESTART

"%JAVA%" %JAVA_OPTS% -Djava.security.auth.login.config="/opt/FileNet/AE/
CE_API/config/jaas.conf.WSI" "-Djava.endorsed.dirs=%JBOSS_ENDORSED_DIRS%" -
classpath "%JBOSS_CLASSPATH%" org.jboss.Main %*
```

4. Save and close the server startup script.
5. Configure LDAP settings on Application Engine to exactly match the Content Engine settings.
  - a. On the Application Engine server, open **login-config.xml**, located in `<JBoss_home>/server/default/conf`, for editing.
  - b. Set the `<application-policy name="FileNet">` entry identical to the corresponding entry in the **login-config.xml** file on the Content Engine server.
  - c. Restart the application server.

6. Set permissions for the user running the application server.

**NOTE** On Windows, the following is only required for NTFS formatted partitions

If the user that will be running the application server is different from the user that installed Application Engine, you must give the user read/write permissions on the folder where you installed Application Engine (*<AE\_install\_path>*), see [“Specify Installation Location” on page 346](#).

7. Continue with the following task:
  - If you are performing a fresh install:  
Do [Task 26 “Install Application Engine Software Updates” on page 383](#).
  - If you are performing an upgrade:  
Do [“Manually copy custom data.” on page 555](#).

## Task 26: Install Application Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Application Engine.

### To install the Application Engine software updates

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your support representative.
2. Open the readmes for any subsequent fix packs or interim fixes (typically optional) and perform the installation procedures provided.
3. Install the latest Content Engine and Process Engine client file updates. See [Task 27 on page 384](#).

## Task 27: Install the Latest Content Engine and Process Engine Client Files on Application Engine Servers

Install any Content Engine or Process Engine client file updates that are available.

### **To install the Content Engine and Process Engine client files**

---

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 service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. Content Engine 4.0.1 Service Pack or any later Fix Packs.
  - b. Process Engine 4.0.2 Service Pack or any later Fix Packs.
3. Deploy Application Engine. See the appropriate task according to your application server type:
  - [Task 28a on page 385](#) (WebSphere)
  - [Task 28b on page 391](#) (WebLogic)
  - [Task 28c on page 394](#) (JBoss)



## Task 28a: Deploy Application Engine (WebSphere)

This topic covers the deployment and start of your Application Engine application (Workplace) on WebSphere.

### To recreate the WAR or EAR file

---

Any time that you make changes to files in the **/Workplace** directory, such as changes to **web.xml** for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

**NOTE** Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  - a. Verify that all modified **/Workplace** directory files have been saved.
  - b. Recreate the **app\_engine.war** file by running **create\_app\_engine\_war.sh** (UNIX) or **create\_app\_engine\_war.exe** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**
- If you will be deploying from an EAR file.
  - a. Verify that a newly recreated **app\_engine.war** file exists.
  - b. Recreate the **app\_engine.ear** file by running **create\_app\_engine\_ear.sh** (UNIX) or **create\_app\_engine\_ear.exe** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**

### To deploy Application Engine (WebSphere 5.1.x)

---

1. Start the application server.
2. Log on to the WebSphere administrative console.
3. From the Administrative Console, expand **Applications**. Click **Install New Application**. The "Preparing for the application installation" dialog box opens.
4. Select file to deploy.
  - (If the Administrative Console is running *locally*) Select **Local Path** and enter or browse to the location of the **app\_engine.war** or **app\_engine.ear** file created by Setup (see below for the default path). Do not enter the machine name.
  - (If the Administrative Console is *remote*) Select **Server path** and enter the fully-qualified path-name to the **app\_engine.war** or **app\_engine.ear** file. Do not enter the machine name.  
**<AE\_install\_path>/deploy**
5. If you are deploying from a WAR file, enter `Workplace` as the context root, and click **Next** to proceed to deploying a new application.

**NOTE** The context root is the name of the application you log in to using the web interface, such as:

`http://<ApplicationEngineServerName>:<port#>/<Context Root>`.

6. From the “Generate Default Bindings” screen, leave the defaults, and click **Next**.
7. From the “Application Security Warning” screen, click **Continue**.
8. At “Install New Application”, Step 1, specify the application name.  
Enter `Workplace`, or the name you chose to call the application, click **Next**.
9. At “Install New Application”, Step 2, specify the **Virtual Host** for the virtual host that you are planning to use. Check **Workplace** and keep the default virtual host (`default_host`), click **Next**.
10. At “Install New Application”, Step 3, configure your application server, and then click **Next**.
11. At “Install New Application”, Step 4, verify your configuration and click **Finish**. Once the configuration is saved, click **Save to Master Configuration**.
12. Configure the Classloader settings.
  - a. From the Administrative Console, expand **Applications**. Click **Enterprise Applications**, and click your application (default `Workplace`).
  - b. From the **Configuration** tab, set Classloader Mode to **PARENT\_LAST**.
  - c. Verify that the **WAR Classloader Policy** is set to **Module**.  
**NOTE** Do this only for the specific web application. Do not change the similar settings for the entire application server.
13. Configure the Web Module Classloader setting.
  - a. From the Administrative Console, expand **Applications**. Click **Enterprise Applications**, and click your application (default `Workplace`). Under Related Items, click **Web Modules**. Click **app\_engine.war**.
  - b. From the **Configuration** tab, set Classloader Mode to **PARENT\_LAST**.  
**NOTE** Do this only for the specific web application. Do not change the similar settings for the entire application server.
  - c. Click **Apply**, click **Save**, and then click **Save changes to the Master Configuration**.
14. Stop and restart the application server.
15. If this is an upgrade, before you bring up the Workplace login, continue with [“Complete Post-Upgrade Application Engine Configuration” on page 556](#).

### **To deploy Application Engine (WebSphere 6.0)**

---

1. Start the application server.
2. Log on to the WebSphere administrative console.
3. From the WebSphere administrative console, expand **Applications**. Click **Install New Application**. The “Preparing for the application installation” dialog opens.

4. Select file to deploy.
  - (If the Administrative Console is running *locally*) Select **Local Path** and enter or browse to the location of the **app\_engine.war** or **app\_engine.ear** file created by Setup (see below for the default path). Do not enter the machine name.
  - (If the Administrative Console is *remote*) Select **Server path** and enter the fully-qualified path-name to the **app\_engine.war** or **app\_engine.ear** file. Do not enter the machine name.  
**<AE\_install\_path>/deploy**
5. If you are deploying a WAR file, enter the context root:  
Enter `workplace` and click **Next** to proceed to deploying a new application.  
**NOTE** The context root is the name of the application you log in to using the web interface, such as:  
`http://<ApplicationEngineServerName>:<port#>/<Context Root>`
6. On the “Preparing for the application installation” screen, leave the defaults, and click **Next**.
7. On the “Application Security Warning” screen, click **Continue**.
8. At “Install New Application”, Step 1, specify the application name.  
Enter `workplace`, or the name you chose to call the application, click **Next**.
9. At “Install New Application”, Step 2, Map modules to servers, specify the **WebServer** you are planning to use. Check **Workplace** and click **Next**.
10. At “Install New Application”, Step 3, Map virtual hosts for Web Modules , check **Workplace** and keep the default virtual host (`default_host`), click **Next**.
11. At “Install New Application”, Step 4, verify your configuration and click **Finish**. Once the configuration is saved, click **Save to Master Configuration**.
12. Configure the Classloader settings.
  - a. From the Administrative Console, expand **Applications**. Click **Enterprise Applications**, and click your application (default `Workplace`).
  - b. From the **Configuration** tab, set Classloader Mode to **PARENT\_LAST**.
  - c. Verify that the **WAR Classloader Policy** is set to **Module**.  
**NOTE** Do this only for the specific web application. Do not change the similar settings for the entire application server.
  - d. Click **Apply**.
13. Configure the Web Module Classloader setting.
  - a. From the Administrative Console, expand **Applications**.
  - b. Click **Enterprise Applications**, and click your application (default `Workplace`).
  - c. Under Related Items, click **Web Modules**. Click **app\_engine.war**.
  - d. From the **Configuration** tab, set Classloader Mode to **PARENT\_LAST**.

**NOTE** Do this only for the specific web application. There are similar settings for the entire application server. Do not change these.

- e. If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Map security roles to users/groups, and verify that the **Everyone?** column is checked for the "Everyone" role .
  - f. Click **Apply**, click **Save**, and then click **Save changes to the Master Configuration**.
14. Stop and restart the application server.
  15. If this is an upgrade, continue with the Upgrade Core Components [Task 13 "Complete Post-Upgrade Application Engine Configuration"](#) on page 556. You will not return to this task.
  16. Start the Enterprise Application.
    - a. From the Administrative Console, expand **Applications**.
    - b. Click **Enterprise Application**.
    - c. Check the box to the left of the **Workplace** application (or whatever you named it), and click **Start**.

#### To deploy Application Engine (WebSphere 6.1)

---

1. Start the application server.
2. Log on to the WebSphere administrative console.
3. From the WebSphere administrative console, expand **Applications**. Click **Install New Application**. The "Preparing for the application installation" dialog opens.
4. Select file to deploy.
  - a. Select Path
    - (If the Administrative Console is running *locally*) Select **Local Path** and enter or browse to the location of the **app\_engine.war** or **app\_engine.ear** file created by Setup (see below for the default path). Do not enter the machine name.
    - (If the Administrative Console is *remote*) Select **Server path** and enter the fully-qualified pathname to the **app\_engine.war** or **app\_engine.ear** file. Do not enter the machine name.
  - b. If you are deploying a WAR file, enter the context root:  
Enter `workplace` and click **Next** to proceed to deploying a new application.

**NOTE** The context root is the name of the application you log in to using the web interface, such as:

```
http://<ApplicationEngineServerName>:<port#>/<Context Root>.
```
5. Click **Next**.

6. At “Install New Application”, Step 1, specify the application name.  
 Enter `workplace`, or the name you chose to call the application, click **Next**.
7. At “Install New Application”, Step 2, Map modules to servers, specify the **WebServer** you are planning to use. Check **Workplace** and click **Apply**. Verify that the webserver you specify is displayed to the right of Workplace. Click **Next**.
8. At “Install New Application”, Step 3, Map virtual hosts for Web Modules , check **Workplace** and keep the default virtual host (`default_host`), click **Next**.
9. At “Install New Application”, Step 4, verify your configuration and click **Finish**. Once the configuration is saved, click **Save**.
10. Configure the Classloader settings.
  - a. From the Administrative Console, expand **Applications**. Click **Enterprise Applications**, and click your application (default Workplace).
  - b. From the **Configuration** tab, click **Class loading and update detection**.
  - c. Change Classloader order from "Classes loaded with parent class loader first" to "Classes loaded with application class loader first."  
**NOTE** Do this only for the specific web application. Do not change the similar settings for the entire application server.
  - d. Change the polling interval to a number appropriate for your environment.
  - e. Click **Apply**.
11. Configure the Web Module Classloader setting.
  - a. From the Administrative Console, expand **Applications**.
  - b. Click **Enterprise Applications**, and click your application (default Workplace).
  - c. Under Modules, click **Manage Modules**. Click **Workplace**, or your application name.
  - d. Change Classloader order from "Classes loaded with parent class loader first" to "Classes loaded with application class loader first."  
**NOTE** Do this only for the specific web application. There are similar settings for the entire application server. Do not change these.
  - e. If you are using container-managed authentication, navigate to Enterprise Applications > Workplace > Map security roles to users/groups, and verify that the **Everyone?** column is checked for the "Everyone" role .
  - f. Click **Apply**, click **Save**, and then click **Save changes to the Master Configuration**.
12. Stop and restart the application server.
13. If this is an upgrade, continue with the Upgrade Core Components [Task 13 “Complete Post-Upgrade Application Engine Configuration” on page 556](#). You will not return to this task.
14. Start the Enterprise Application.
  - a. From the Administrative Console, expand **Applications**.

- b. Click **Enterprise Application**.
- c. Check the box to the left of the **Workplace** application (or whatever you named it), and click **Start**.

## Task 28b: Deploy Application Engine (WebLogic)

This topic covers the deployment of your Application Engine application (Workplace) on WebLogic.

### To recreate the WAR or EAR file

---

Any time that you make changes to files in the **/Workplace** directory, such as changes to **web.xml** for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

**NOTE** Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  - a. Verify that all modified **/Workplace** directory files have been saved.
  - b. Recreate the **app\_engine.war** file by running **create\_app\_engine\_war.sh** (UNIX) or **create\_app\_engine\_war.bat** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**
- If you will be deploying from an EAR file.
  - a. Verify that a newly recreated **app\_engine.war** file exists.
  - b. Recreate the **app\_engine.ear** file by running **create\_app\_engine\_ear.sh** (UNIX) or **create\_app\_engine\_ear.bat** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**

### To deploy as “Workplace” or custom name using a WAR file

---

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users type to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the **app\_engine.war** file to reflect the name you want to use using the format **<Application Name>.war**.

Example:

The default **app\_engine.war** will generate the following context root:

```
http://<server_name>:<port#>/app_engine
```

Renaming the WAR file **Workplace.war** will generate the following context root:

```
http://<server_name>:<port#>/Workplace
```

**CAUTION** You must rename the WAR file every time you regenerate it. The **create\_app\_engine\_war.sh/.bat** script will by default create a file with the name **app\_engine.war**.

### To deploy Application Engine (WebLogic 8.1.x)

---

1. Start the application server.
2. (For WAR file or exploded directory deployment only) Create a new Web Application for WAR file or exploded directory deployment.
  - a. From the WebLogic Server Console, click **<mydomain>**, click **Deployments**, and then click **Web Application Module**.
  - b. Click **Deploy a new Application**.
  - c. From the right pane of the WebLogic Administration Console, do one of the following:
    - To deploy from an exploded folder, browse to and select the radio button for the Workplace folder in:  
**<AE\_install\_path>**
    - To deploy from a WAR file, browse to and select the radio button for the WAR file you want to deploy (Default: **app\_engine.war**) in:  
**<AE\_install\_path>/deploy**
  - d. Click **Target Module** to select **Workplace**.
  - e. Verify that the **Name** field has the value **Workplace**.
  - f. Click **Deploy**.

**NOTE** To verify that the deployment was successful, expand **Applications**. The web application Workplace will be listed.
3. (For EAR file deployment only) Create a new Web Application for EAR file deployment.
  - a. From the WebLogic Server Console, click **<mydomain>**, click **Deployments**, and then click **Application**.
  - b. Click **Deploy a new Web Application Module**.
  - c. From the right pane of the WebLogic Administration Console, browse to and select the radio button for the **app\_engine.ear** file in:  
**<AE\_install\_path>/deploy**
  - d. Click **Target Module** to select **Workplace**.
  - e. Verify that the **Name** field has the value **Workplace**.
  - f. Click **Deploy**.

**NOTE** To verify that the deployment was successful, expand **Web Applications**. The web application Workplace will be listed.
4. (Upgrades only) If this is an upgrade, continue with Upgrade Core Components [Task 13](#) "Complete Post-Upgrade Application Engine Configuration" on page 556. You will not return to this task.



5. Continue with [“To start the web application” on page 393](#).

### **To deploy Application Engine (WebLogic 9.2)**

---

1. Start the application server.
2. From the WebLogic Administration Console, navigate to the domain you created for the Application Engine in [“Configure an Application Server for Application Engine \(WebLogic\)” on page 124](#).
3. Click **Deployments**.
4. In the Change Center, click **Lock & Edit**.
5. From the right pane of the WebLogic Administration Console, do one of the following:
  - To deploy from an exploded folder, browse to and select the radio button for the Workplace folder in:  
**<AE\_install\_path>**
  - To deploy from a WAR or EAR file, browse to and select the radio button for the WAR file you want to deploy (Default: **app\_engine.war** or **app\_engine.ear**) in:  
**<AE\_install\_path>/deploy**
6. Click **Next**.
7. Accept the defaults for the deployment, except for the name for the deployment. Use “Workplace” instead of “appengine”.
8. Click **Finish**.
9. Click **Save**, and then click **Activate Changes**.

**NOTE** To verify that the deployment was successful, expand **Web Applications**. The web application Workplace will be listed.
10. If this is an upgrade, continue with the Upgrade Core Components [Task 13 “Complete Post-Upgrade Application Engine Configuration” on page 556](#). You will not return to this task.
11. Continue with [“To start the web application” on page 393](#).

### **To start the web application**

---

After Workplace is deployed, start the web application.

1. From the WebLogic Administration Console, navigate to **<your\_domain> > Deployments** and select the box next to Workplace.
2. Click **Start**, and then select **Servicing all requests**.
3. Click **Yes** to start the application.

## Task 28c: Deploy Application Engine (JBoss)

This topic covers the deployment and start of your Application Engine application (Workplace) on JBoss.

### To recreate the WAR or EAR file

---

Any time that you make changes to files in the **/Workplace** directory, such as changes to **web.xml** for container-managed authentication, SSO support, or any other edits, you must recreate the WAR or EAR file and redeploy your changes.

**NOTE** Before recreating the EAR file, you must also recreate the WAR file.

- If you will be deploying from a WAR file.
  - a. Verify that all modified **/Workplace** directory files have been saved.
  - b. Recreate the **app\_engine.war** file by running **create\_app\_engine\_war.sh** (UNIX) or **create\_app\_engine\_war.exe** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**
- If you will be deploying from an EAR file.
  - a. Verify that a newly recreated **app\_engine.war** file exists.
  - b. Recreate the **app\_engine.ear** file by running **create\_app\_engine\_ear.sh** (UNIX) or **create\_app\_engine\_ear.exe** (Windows) from the following location:  
**<AE\_install\_path>/deploy/**

### To deploy as “Workplace” or custom name using a WAR file

---

Perform this step only if you are using a WAR file for deployment, and you want to use “Workplace” or a custom name for the context root of the application. The context root is part of the URI that end users type to access Workplace. By default, when you deploy from a WAR file, the context root is the first part of the WAR filename.

Rename the **app\_engine.war** file to reflect the name you want to use using the format **<Application Name>.war**.

Example:

The default **app\_engine.war** will generate the following context root:

```
http://<server_name>:<port#>/app_engine
```

Renaming the WAR file **Workplace.war** will generate the following context root:

```
http://<server_name>:<port#>/Workplace
```

**CAUTION** You must rename the WAR file every time you regenerate it. The **create\_app\_engine\_war.sh/.exe** script will by default create a file with the name **app\_engine.war**.

## To deploy and start Application Engine

---

- To deploy from exploded directory.
  - a. On the JBoss server, copy the **Workplace** folder from:  
**<AE\_install\_path>**  
to:  
**<JBOSS\_home>/server/default/deploy/**
  - b. Append the extension .war to the Workplace folder:  
**<JBOSS\_home>/server/default/deploy/Workplace.war**
- To deploy from WAR file.

On the JBoss server, copy the **app\_engine.war** file from:

**<AE\_install\_path>/deploy**  
to:  
**<JBOSS\_home>/server/default/deploy/**
- To deploy from EAR file.

On the JBoss server, copy the **app\_engine.ear** file from:

**<AE\_install\_path>/deploy**  
to:  
**<JBOSS\_home>/server/default/deploy/**
- 4. Set permissions for the user running the application server.

If the user that will be running the application server is different from the user that installed AE, you must give the user read/write permissions on the following folders:

**NOTE** For Windows this is only required for NTFS formatted partitions:

**<JBOSS\_home>/server/default/deploy/app\_engine.war/.ear**  
**<AE\_install\_path>**
- 5. Start or restart the JBoss application server.
- 6. Verify that the application deployed successfully.

Verify that the **server.log** file located in **<JBOSS\_home>/server/default/log** lists deployment of the WAR or EAR file you used.
- 7. If this is an upgrade, continue with the Upgrade Core Components [Task 13 "Complete Post-Upgrade Application Engine Configuration"](#) on page 556. You will not return to this task.

# Configuration/Startup Tasks

## To configure the IBM FileNet P8 Platform components

---

1. Configure Content Engine for content-based retrieval (CBR). Do [Task 29 on page 397](#).
2. Set bootstrap preferences. Do [Task 30 on page 403](#).
3. Initialize the isolated region. Do [Task 31 on page 408](#).
4. Create a Process Engine connection point. Do [Task 32 on page 409](#).
5. Configure the Process Engine connection point for Application Engine. Do [Task 33 on page 410](#).
6. Set up Content Engine and client transport SSL security. Do [Task 34 on page 412](#).
7. Set up Application Engine SSL security. Do [Task 35 on page 416](#).
8. Perform additional configuration tasks. Do [Task 36 on page 420](#).
9. Familiarize yourself with IBM FileNet P8 system startup and shutdown procedures. See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup](#).

## Task 29: Configure Content Engine for Content-Based Retrieval

This task is for new installations only and is required if you have installed the Content Search Engine. It covers how to use Enterprise Manager to configure an index area and enable the Content-Based Retrieval (CBR) feature provided by the IBM FileNet P8 Content Search Engine. This search engine is based on Autonomy K2.

If you are upgrading and have 3.5.x content-search indexes that you want to upgrade to 4.0.0 index areas (now called K2 collections), do not complete this task. Instead, follow the upgrade procedure, [Task 3 Upgrade Content Engine Software on page 487](#).

### NOTES

- Before you complete the procedure in this topic, ensure the IBM FileNet P8 Content Search Engine has been installed and configured on at least one server (In effect, this means you have already installed and configured an Autonomy K2 Master Administration Server). For details, see [Task 8 Install and Configure Content Search Engine on page 147](#).
- Numerous K2 security accounts are referenced within this procedure. For more information on the accounts required, see [“To create Content Search Engine accounts” on page 60](#) for details on which accounts to designate and the permissions to assign.
- Various server names and related ports assigned during the installation and configuration of Autonomy K2 will be required during this procedure. If you do not have a record of the servers created and the ports that have been assigned, log into the Autonomy K2 dashboard to obtain the information necessary.
- If you unimport the style set, the original files will be deleted from your system. In this scenario, if you wish to re-import the style set, you will need to recover it from your installation disk. In order to avoid this situation, you can either enter a unique name for the Style Set Alias during the initial Content Search Engine (Autonomy K2) installation, or make a backup copy of the original style set. If you entered a unique name for the style set during installation, ensure you use that Style Set Alias name for this procedure.

### **To create a collections directory**

---

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

Security and communication between Autonomy K2, Content Engine, and the collections directory is handled through the user accounts and permissions provided to those accounts. For more information on the accounts required, see [“To create Content Search Engine accounts” on page 60](#) for details on which accounts to assign and the permissions to add. For detailed information on security, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Authorization > Security for integrated components and third-party products > Autonomy K2 Server > Security for Autonomy K2 Server](#).

1. Create a directory on the Verity server on which you will store collections (VerityIndexArea.RootDirectoryPath). This directory must be located on a disk that is local to the Verity server. Using a

remote mounted disk accessed via the network (NFS, PCNFS, or CIFS) will cause stability problems under load and corrupt collections. Set permissions to allow access to the K2 Operating System User.

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

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

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

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

### For upgrades of 3.5.x collections

Windows:

- a. Open the following K2 configuration file in a text editor:

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

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

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

- c. Modify the next available alias settings by adding an entry for each 3.5.x collections path, expressed as UNC, that you want to upgrade. For example, change alias7, mapping7, and dirmode7 to the following:

```
alias7=path2
mapping7=\\<server_host_name>\<Collections_Directory>
dirmode7=wr
```

UNIX:

- a. Open the following K2 configuration file in a text editor:

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

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

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

**NOTE** The Collections\_Directory\_Path must be a local path and not a mount point.

**For new installations where you are not upgrading existing collections**

Windows:

- a. Open the following K2 configuration file in a text editor:

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

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

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

UNIX:

- a. Open the following K2 configuration file in a text editor:

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

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

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

**NOTE** The Collections\_Directory\_Path must be a local path and not a mount point.

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

**To Configure Content Engine for CBR**

---

This procedure covers the minimum setup and configuration steps to get CBR configured and running with Autonomy K2. For more detail on Content-Based Retrieval and Content Engine, see the IBM FileNet P8 help topic Configure CBR found at [FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Configure CBR](#).

**NOTE** Where machine name variables are required, IP addresses will not validate. In these cases, you must enter the host name for the machine.

1. Launch Enterprise Manager and log in as the GCD Administrator.
2. Create a Verity Domain Configuration (VDC) for K2.
  - a. Right-click **Enterprise Manager [domain]** in the Enterprise Manager tree and select Properties.
  - b. Click the **Verity Domain Config.** tab.

- c. Enter the following K2 Master Administration Server access information:
    - Host Name - the name of the host of the K2 Master Administration Server.
    - Port - the K2 Master Administration Server port.
    - User Domain - the authentication domain in which your K2 services are installed.
    - User Group - K2 Security Group
    - Verity Username - K2 Security User
    - Password - the K2 Security User password.
  - d. Click Create Configuration to save your settings and create a Verity Domain Configuration object.
3. Assign a K2 Broker Server:
    - a. Click the **Verity Server** tab.
    - b. From the Brokers AVAILABLE pane, select the broker and click **Add** to move the server to the Brokers Selected pane.
    - c. Click **OK**.

**NOTE** You may assign multiple Broker Servers to a K2 Administration Server, primarily for failover. If one Broker Server goes down, then K2 can switch to another. In this configuration, you must ensure that all Search Servers required to access K2 Collections (index areas) are attached to each Broker Server. Be aware that a given Content Engine server will neither call multiple Broker Servers nor merge associated search results. See steps 3 and 4 of ["To configure Autonomy K2 for Content-Based Retrieval" on page 153](#) for instructions on how to create Broker Servers and assign Search Servers.

4. Enter a CBR Locale.

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

- a. Right-click the Object Store and select **Properties**.
  - b. Click the **Locale** tab.
  - c. Enter a valid Autonomy K2 locale for CBR.
5. Create an index area, as follows:
    - a. Right-click the name of the object store to which you want to add an index area (K2 collection) and select **New > Index Area**.
    - b. Enter the following information.
      - Display Name: *<Index\_Area\_Name>*
      - Descriptive Text: *<Description>*



- Site: Choose which site to associate with the index area.
  - Template Type: FileNet\_FileSystem\_PushAPI.
  - File system root directory for Verity Collections: Enter the full path to the collections directory (verity.rootdirectory.path) for this index area (for example, \\<myserver>\<Verity\_Collections>).
  - File system temporary directory for Verity Collections: Enter the full path to the temporary collections directory (verity.temporary.path) for this index area (for example, \\<myserver>\<Verity\_Collections\_temp>).
- c. Select <server\_name>\_search\_server in the Verity Search Servers selection window.
  - d. Highlight <server\_name>\_index\_server in the Verity Index Servers selection window.
  - e. Click **Create Index Area**.
  - f. Click **OK** in the confirmation window.
6. Enable CBR for class definitions by activating the CBR Enable flag of the class you want available for CBR, as follows:
    - a. Right-click the class you want to configure in the Enterprise Manager tree and click **Properties**.
    - b. Select **CBR Enabled** and click **OK**.
  7. Enable CBR for the class properties you want available for CBR, as follows:
    - a. Right-click the class you want to configure and click **Properties**.
    - b. Click the **Property Definitions** tab.
    - c. Click the string property you want to enable for CBR indexing and click **Edit**.
    - d. Check **CBR Enabled** and click **OK**.

### To enable additional K2 Index Servers and Search Servers

---

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

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

5. Enable the new Index Servers as follows:
  - a. Click **Edit Index Servers**.
  - b. In the Index Servers Available pane, Highlight any servers you want to enable for this index area and click **Add** to add the server to the Index Servers Selected list.
  - c. Click **OK** to save the settings and enable the new servers.

## Task 30: Set Application Engine Bootstrap Preferences

Bootstrap preferences are a category of site preferences. The first time you sign into Workplace after Application Engine installation, the Bootstrap Preferences page opens.

### ***Bootstrap Preferences***

The following six bootstrap preference groups are available the first time you sign in:

- Security Info (required for SSL only)
- User Token Settings
- Preference Settings (required)
- Banner Image
- Application Integration
- Administrator Access Role

For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site Preferences > Bootstrap preferences](#).

### ***Enhanced Timezone Detection***

In addition to these settings you can also set the *useEnhancedTimeZoneDetection* parameter to accurately detect a client browser's time zone. This setting cannot be modified through the Site Preferences page. To enable this feature you must manually modify the **bootstrap.properties** file. For more information, see the IBM FileNet P8 help topic [FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > bootstrap.properties](#).

#### **NOTES**

- By successfully signing in to Workplace and saving the bootstrap preferences, you are verifying the Application Engine's basic functionality such as user authentication as well as communication and storing of data in Content Engine.
- In addition to the preferences covered in this topic, more preferences can be set for the Workplace application using Workplace Site Preferences. For more information, see the the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences](#).
- After the initial bootstrap configuration, users with the Application Engine Administrators role can change any of these preferences by signing into Workplace and navigating to Admin > Site Preferences > Bootstrap.
- When you access the bootstrap preference page via the Site Preferences application, an additional preference, **Guest info** (to allow guest sign ins), is also available.
- In a web farm/clustered environment, all Application Engines share site preferences by using the same **bootstrap.properties** file. For more information, see the *IBM FileNet P8 Platform High Availability Technical Notice*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).

- Bootstrap properties are stored in a separate file from other site preference. The default location for this file, **bootstrap.properties**, is:
  - WebSphere  
**<WebSphere\_home>/installed Apps/Workplace.ear/app\_engine.war/WEB-INF**
  - WebLogic  
**<AE\_install\_path>/Workplace/WEB-INF**
- (New installations only) To allow users to create workflows subscriptions, you *must* configure the PWDesigner access role. For more information, see [“\(New installations only\) To enable user access to the Workflow Subscription Wizard” on page 407](#).

### To set the bootstrap properties on first login

---

1. Sign in to Workplace:
  - a. On any computer, open a browser and type:  
`http://<ApplicationEngineServerName>:<port#>/Workplace`
  - b. Enter a user name and password, and then click **Sign in**. The Bootstrap Preferences page opens.

**NOTE** The user who initially logs in and sets the bootstrap preferences is automatically added to the Application Engine Administrators role. For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences](#).
2. Enter security info (required for SSL only).
  - a. Enter the SSL Host and Port information for the SSL server.
  - b. Enter the Java Server HTTP port.

Use the Security info preference to redirect sign-ins through a Secure Socket Layer (SSL) server and to identify a remote Java server. This encrypts the user IDs and passwords when they travel over the network. See [“Set Up Application Engine SSL Security” on page 416](#) for instructions on setting up SSL security for one or more Application Engines.

**CAUTION** Once you’ve configured SSL, the Site Preferences application also runs under SSL to protect the guest account’s user ID and password. This means that when you run Site Preferences on an unsecured server that redirects sign-ins to an SSL server, you will be editing the Bootstrap preferences of the SSL server (stored in the bootstrap.properties file). This does not affect the General, Object Store, and Shortcut preferences, which are retrieved from the preferences file saved in the object store.
3. Configure user token settings.

User Tokens are used by IBM FileNet P8 applications to launch into each other without the need for an additional login.

  - a. Select whether or not to create user tokens for your Application Engine (Default: **Yes**).

- b. Select whether or not the application will pick up generated tokens from other applications (Default: **Yes**).
  - c. Specify a Token timeout interval (1 - 15 minutes).
4. (Required) Specify preference settings.

Preference settings specify the name of the site preference file, its storage location (object store), and the documentation server URL (if installed). The site preferences file is checked into the object store the first time you log on to Workplace. The site preferences are stored as XML data in this file, **<Site Preferences for Preferences name>.xml**. Use Enterprise Manager to access this file, and navigate to **Object Stores > Object Store location > Root Folder > Preferences**.

**NOTE** The bootstrap preferences are saved in the **bootstrap.properties** file, and not in the site preferences file.

- a. Select an object store from the **Object store location** choice list. The preferences file will be saved in this object store. Workplace users must have access to this object store.
- b. Enter a preference file name in the **Preferences name** field.
- c. Enter the documentation server URL in the **Documentation server** field.

The format of the URL is:

```
http://<DocServerName>:<port#>/ecm_help/
```

where *<DocServerName>* is the name of your Java application server where the documentation is installed,  
*<port#>* is the port number,  
and *<ecm\_help>* is the root directory of the documentation web site.

**NOTE** If no documentation URL is specified, the Workplace Help link will default to `http://localhost`.

- d. Enter the ISRA Interface Servlet URL.

For more information, see ["Enable Application Engine to Use ISRA" on page 451](#).

5. Set Banner Image.

The banner image is the graphic that appears at the upper left -hand side of the Workplace application pages. If you have a banner image that you would like to use in place of the default, follow this procedure.

- a. Copy the new graphic file to the location of your choice on Application Engine in the **/FileNet/AE/Workplace** folder.
- b. In the Path field, type the path (relative to the **/Workplace** folder) to the new banner graphic file.
- c. In the Image Width field, type the width of the image (in pixels).
- d. In the Image Height field, type the height of the image (in pixels).

6. Configure Application Integration.

Select **No** (default), if you do not want users to be prompted to add an email to an object store each time the email is sent.

Select **Yes**, if you want users prompted to add an email to an object store when the email is sent.

This preference setting only affects Outlook integration.

7. Add Application Engine Administrators.

Add the users and groups that will perform Application Engine administration tasks to the *Application Engine Administrators* role.

**NOTES**

- The user who initially signs in and sets the bootstrap preferences is automatically added to the Application Engine Administrators role. For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences](#).
- To modify the access roles after the initial bootstrap configuration, users with the Application Engine Administrators role can use the access roles section of the Workplace Site Preferences. Launch Workplace and navigate to Admin > Site Preferences > Access Roles.

8. Click **Apply** to save your bootstrap settings.

**To add a single index for Application Name on the site preferences object store**

---

To properly index access roles and improve login performance on Application Engine, you must create an index for Application Name on the object store that contains the Workplace site preferences. Do this after you have successfully configured the bootstrap preferences.

1. On Content Engine, launch the Enterprise Manager.
2. In the left pane, expand the **Object Stores** folder.
3. Expand the object store node where your preferences are stored. See [“\(Required\) Specify preference settings.” on page 405](#) above.
4. Expand **Other Classes** and then **Custom Object**.
5. Right click **Access Role** and select **Properties**.
6. Select the **Property Definitions** tab.
7. Select **Application Name** and click **Edit**.
8. On the General tab of the Application Name Properties, if the **Indexed** field shows 'not indexed', click **Set/Remove**.
9. Select **Set** and then **Single Indexed**.
10. Click **OK** to set the index.
11. Click **OK** to apply the change and close the Application Name Properties window.
12. Click **OK** to close the Access Role Class Properties window.

**NOTE** If you are performing an upgrade, continue with [Step 4 on page 557](#).

### **(New installations only) To enable user access to the Workflow Subscription Wizard**

---

To allow users to create workflow subscriptions, you must configure the PWDesigner access role using the Workplace Site Preferences, and give the users appropriate access rights to the workflow subscriptions classes. You can perform these steps in any order, and you must perform both steps any time you need to add or remove users.

1. Assign users as members of the PWDesigner access role. See the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences](#).
2. Run the security script wizard, and load the **workplacescript.xml** file to add accounts to the Workflow Designer role.

For more information about how to use the Security Script wizard to assign the Workflow Designer role to user or group accounts, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content Engine Wizard Help > Security Script](#).

For more information about the **workplacescript.xml** file and how roles are defined in the Enterprise Manager, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Managing Security > Security Script Wizard](#).

## Task 31: Create a Process Engine Isolated Region

Process Engine communicates to its database using a connection point. Each connection point is associated with an isolated region. In this task you will create an isolated region. In [Task 32 on page 409](#) you will define a connection point to this isolated region.

### To create a Process Engine isolated region

---

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn 4.0 on the desktop, or by navigating to Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn 4.0. Log on as a GCD administrator.
2. Connect to the FileNet P8 domain you created in ["Install and Deploy Content Engine" on page 160](#).
3. Right-click **PE Region ids > New PE Region ids**.
4. Click **Next** on the Specify a Site screen to select a site named *initial site*.
5. Enter the DNS name for the Process Engine server.
6. Enter the region ID.
7. Modify the communication port if needed.
8. Click **Next** when done.
9. Enter the password for the isolated region. This password must match that entered in the Process Engine Task Manager for the isolated region in ["Configure Process Task Manager" on page 340](#).
10. Click **OK** on the Confirmation Window.
11. Click **Finish** to finish create a new region for Process Engine.



## Task 32: Create a Process Engine Connection Point

A connection point identifies a specific isolated region of the workflow database, and gives it a name that workflow-related applications use to access the region. Follow these procedures to create a connection point.

### To create a Process Engine connection point

---

1. Start Enterprise Manager by double-clicking the FileNet Enterprise Manager SnapIn 4.0 on the desktop, or by navigating to Start > All Programs > FileNet P8 Platform > Enterprise Manager SnapIn 4.0. Log on as a GCD administrator.
2. Connect to the FileNet P8 domain you created in [“Install and Deploy Content Engine” on page 160](#).
3. Right-click PE Connection Points > New PE Connection Points.
4. Enter a Process Engine Connection Point name and click **Next**.
5. Choose the region which is created in [“Create a Process Engine Isolated Region” on page 408](#), and click **Next**.
6. Click **Finish** to finish creating the Connection Point.
7. Click **OK**.

## Task 33: Configure the Process Engine Connection Point for Application Engine

Before users can access tasks and work items from Workplace, you must configure the Connection Point on the Application Engine. Make sure that you have already completed these Tasks:

- “Create a Process Engine Isolated Region” on page 408
- “Create a Process Engine Connection Point” on page 409

### To configure the connection point

---

1. Log onto Application Engine:
  - a. On any computer, open a browser and navigate to:  
`http://<ApplicationEngineServerName>:<port#>/Workplace`
  - b. Sign in using the same account that you used to set the bootstrap preferences.
2. Click **Admin**.
3. Click **Site Preferences**.
4. Under General Settings > Tasks, select a **Process Engine Connction Point** from the drop down list.
5. Click **Apply**.
6. Click **Exit**.
7. Initialize the isolated region.
  - a. Click **Admin**.
  - b. Click **Process Configuration Console**.

**NOTE** If your computer does not have the appropriate Java Runtime Environment (JRE) installed, you will be prompted to download the JRE at this point; follow the prompts to complete the download. During the installation process, click the **Browser** tab and enter the following settings:

    - De-select (clear) the Internet Explorer option.
    - If you will be using Netscape 6.0, select the Netscape option.
  - c. Right-click the **icon** or **name** of the isolated region you want to initialize, and select **Connect** from the context menu.
  - d. Click **Action**.
  - e. Click **Initialize Isolated Region**.
  - f. Click **Yes** at the prompt asking if you want to continue.
  - g. Close the Process Configuration Console.

8. In Workplace, click **Tasks** to confirm that Application Engine is communicating with Process Engine.
9. Sign out of Workplace.

## Task 34: Set Up Content Engine and Client Transport SSL Security

Configuring SSL enables secure communications between the Content Engine and the directory service, as well as between Content Engine clients and the Content Engine server. In addition, setting up Content Engine SSL provides secure authentication for Process Engine.

**CAUTION** IBM strongly recommends enabling SSL for the Content Engine and Process Engine web services. Authentication over these two web services is usually performed by providing username and password credentials. If these web services are not configured to run over an SSL connection, clear text passwords will be sent across the network. (However, this is not true when Kerberos-based authentication is used. In the P8 4.0.0 release, Kerberos authentication is available only for the Content Engine web service.) The option not to use SSL over these two web services is provided primarily for development systems or other non-production systems where the security provided by SSL may not be required.

For access to the Content Engine through the EJB transport (IIOP or T3 protocol), an SSL connection is necessary to provide privacy for data sent across the network, but user passwords would not be compromised if SSL were not used. While it is preferable to use SSL with the EJB transport (IIOP or T3 protocol), it is not a requirement.

### NOTES:

- The Content Engine web service is used:
  - By all clients of the Content Engine 4.0 .NET API
  - By all clients of the Content Engine 4.0 COM Compatibility API (CCL)
  - By the Enterprise Manager tool
  - By the Content Engine 3.5.2 to 4.0.0 Upgrade tool
  - By the Process Engine, when making calls to the Content Engine to retrieve user and group information
  - By the Component Manager, running on the Application Engine, which is an integral component for BPM Process Orchestration framework
  - By customer and 3rd party tools written against the CE 3.5 web service API, including Altien Document Manager and the Sharepoint integration done by Vorsite.
- Certain Java applications (written against the Content Engine 3.5.x Java API or the Content Engine 4.0.0 Java API) may use the Content Engine web service transport, but typically they would use EJB transport (IIOP or T3 protocol).
- The IBM FileNet Application Engine server will use only the EJB transport to communicate with the Content Engine in the P8 4.0.0 release.
- The Process Engine web service is used by customer and third-party applications to write runtime applications (typically step processor applications) against the Process Engine. The Process Engine Java API does not make use of the Process Engine web service.

## To enable SSL for Content Engine

---

**NOTE** In the steps below, a server certificate certificate will be added to the Directory Services server (for authentication). In addition, the CA certificate will be added in two different locations on the Content Engine server (the JDK path location is for authorization). Follow the steps closely to ensure that the proper certificate is added to each of the three locations.

1. Obtain and install a server certificate and a CA certificate on the directory service. These certificates are available from third-party certificate authorities, such as VeriSign, or you can generate your own certificates if you have the necessary certificate management software installed.
2. Enable SSL on the directory service and set the SSL port number. The default SSL port number is 636; however, if you have more than one directory service that is using SSL on the server, you may need to use a non-default port number. See your directory server documentation for instructions.
3. On the Content Engine server, add the CA certificate to the application server keystore, if it does not already contain it.
4. On the Content Engine server, add the CA certificate to the JDK (Java) keystore, if it does not already contain it. You can use the default key store, in [Step a](#), or create a custom location, in [Step b](#).
  - a. To use the JDK default java key store, do the following:
    - i. Determine the java version your application server uses and the JAVA\_HOME location.
    - ii. Use the keytool to import the CA certificate to the Java keystore at **%JAVA\_HOME%\jre\lib\security\cacerts**.
    - iii. To improve security, change the default password.
  - b. To use your own key store (rather than the JDK default key store), do the following:
    - i. Add the following system parameters to the Java command line in your application server's startup script:

```
-Djavax.net.ssl.trustStore=<path_to_your_keystore_file>  
-Djavax.net.ssl.trustStorePassword=<password_of_your_keystore>
```
    - ii. Use the Java keytool to import the CA certificate to your own keystore.
5. Use Enterprise Manager to enable SSL for Content Engine and set the port number to match the SSL port on the directory server, as described in ["To enable SSL between Enterprise Manager and the directory service" on page 414](#).
6. Obtain another server and CA certificate for the Content Engine.
7. Create a custom identity keystore on the Content Engine server, and add the server certificate to the custom keystore.

8. Using the application server administration tool, enable SSL and point to the custom identity keystore. Directions vary by application server type; see your application server documentation for detailed procedures.

- WebLogic

Set up a custom identity keystore. In the left pane of the WebLogic Administration Console, navigate to *DomainName* > Servers > *ServerName*. In the right pane, select Keystores and SSL and specify the keystore information.

- WebSphere

Configure an SSL repertoire. In the left pane of the WebSphere administrative console, navigate to Security > SSL. In the right pane, select your JSSE repertoire and specify key and trust file names and passwords.

**NOTE** (WebLogic only) The name in your certificate must match the host name specified in your WebLogic application server. If the name in the certificate is fully qualified (for example, Host1.filenet.com), the same fully qualified host name must appear in the Host field (WebLogic > Authentication Provider > Active Directory tab > Host field).

9. Configure clients to use a particular URL for connecting to Content Engine based on the application server type and the client transport (protocol) type. The following table provides details:

Protocol	SSL	Port	App Server	Sample URL
HTTP	no	7001	WebLogic	<http://mycorp.com:7001/wsi/FNCEWS40DIME/>
HTTPS	yes	7002	WebLogic	<https://mycorp.com:7002/wsi/FNCEWS40DIME/>
T3 (IIOP)	no	7001	WebLogic	t3://mycorp.com:7001/FileNet/Engine
T3S (IIOP)	yes	7002	WebLogic	t3s://mycorp.com:7002/FileNet/Engine
HTTP	no	9080	WebSphere	<http://mycorp.com:9080/wsi/FNCEWS40DIME/>
HTTPS	yes	9403	WebSphere	<https://mycorp.com:9403/wsi/FNCEWS40DIME/>
IIOP	no	2809	WebSphere	iiop://mycorp.com:2809/FileNetEngine
IIOP	yes	9443	WebSphere	iiop://mycorp.com:9443/FileNetEngine

**NOTE** The port values in the table above are default values. If you change the port that your application server listens on, you must also change the port number used by the Content Engine client.

**To enable SSL between Enterprise Manager and the directory service**

1. Launch Enterprise Manager and log on as a GCD administrator.
2. In the tree view, right-click the root node and choose Properties.

3. In the Enterprise Manager Properties dialog box, click the Directory Config. tab, select a directory service, and click **Modify**.
4. In the General tab of the Modify Directory Configuration dialog box, set the Is SSL Enabled parameter to True and modify the port number appropriately.
5. Click **OK** in each open dialog box.

## Task 35: Set Up Application Engine SSL Security

This topic describes how to configure an Application Engine to direct sign-ins through a Secure Socket Layer (SSL) https connection. It assumes that Application Engine(s) have already been installed.

IBM FileNet Application Engine supports the following methods of configuring an SSL environment:

- Full SSL support - A single Application Engine server, where all of the software is running under SSL.
- One server SSL redirect - One Application Engine server set up to redirect logon attempts on the non-SSL port to the SSL port.
- Two server SSL redirect - Two Application Engine servers, where one is SSL-enabled, and the other redirects users to the SSL-enabled Application Engine server to log on.

### **To set up full SSL support on a single Application Server**

---

1. Enable SSL on the application server that runs Application Engine (see your SSL documentation).
2. Test the SSL connection by signing into Workplace using one of the following URLs:

`https://<Application_Engine_server_name>:<SSL_port>/Workplace`

The entire sign-in process will be handled by the SSL-enabled host.

For more information about SSL port numbers, see ["IBM FileNet P8 Port Numbers" on page 643](#).

### **To set up SSL redirect on a single Application Engine server**

---

1. Enable SSL on the application server that runs Application Engine (see your SSL documentation).
2. Sign in to Workplace:

- a. On any computer, open a browser and type the following URL address:

`http://<Application_Engine_server_name>:<port#>/Workplace`

- b. Sign in as a user with Application Engine Administrator access role privileges. For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences](#).

3. Set bootstrap preferences:

- a. Navigate to Admin Site Preferences > Bootstrap.
- b. Set the Security info Site Preference SSL Host:Port to identify the alias host name and port number.

Use the IP address of the Application Engine server for the SSL Host entry.

For more information, see ["Enter security info \(required for SSL only\)." on page 404](#).



- c. Click **Apply** to save your bootstrap settings.
4. Update the base URL:
  - a. Navigate to Admin > Site Preferences > Refresh.
  - b. Enter the Workplace Base URL value in the provided field. The URL must contain a valid host name, and not contain "localhost" or an IP number. For example, `http://myserver:7001/Workplace`  
 For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Refresh preferences](#).
  - c. Click **Refresh** to update the base URL.
  - d. Click **Exit** to close Site Preferences.
5. Sign out of Workplace, and close your browser.
6. Test the SSL connection by signing into Workplace using the following URL:

`http://<Application_Engine_server_name>:<non-SSL port>/Workplace`

**NOTE** You will be redirected to the SSL-enabled port for sign in, then back to the non-SSL enabled port after sign-in is complete. Before sign-in, you should receive a warning that you are accessing pages over a secure connection (unless you turned this dialog box off), and then Workplace will open.

### **To set up SSL redirect on two Application Engine servers**

---

1. Install Application Engine on both computers so that both Application Engines use the same **bootstrap.properties** file and site preferences file (the Setup program will prompt you for a shared location).

During setup of the first Application Engine, create a share on the folder where the **bootstrap.properties** file is installed (the **WEB-INF** folder). Then during setup of the second Application Engine, specify the shared location from the first installation. The **bootstrap.properties** file must already exist when specifying a shared location. See "Setup WebLogic clusters" or "Setup WebSphere clones" in the *IBM FileNet P8 Platform High Availability Technical Notice* for specific instructions. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).

**CAUTION** The system clocks on the two Application Engine servers must be synchronized to within the Token time-out interval. For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Bootstrap Preferences > User token settings](#).

2. Copy the **UTCryptokeyFile.properties** file.

For SSL redirect to work, each Application Engine must use the same User Token cryptographic key file.

After installing the second Application Engine, copy the **UTCryptoKeyFile.properties** file from the first Application Engine server to the same location on the second Application Engine

server. See [“Make a note of the user token crypto key path.” on page 353](#) for information on the default location for the UTCryptoKeyFile.properties file.

**NOTE** IBM recommends copying the file over a secure link.

3. Enable SSL on the application server that you are using for the SSL-enabled Application Engine (see your SSL documentation).
4. Sign in to Workplace on the non-SSL enabled Application Engine.
  - a. On any computer, open a browser and type:  

```
http://<ApplicationEngineServerName>:<port#>/Workplace
```
  - b. Sign in as a user with Application Engine Administrator access role privileges. For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Access Roles preferences](#).
5. Set bootstrap preferences:
  - a. Navigate to Admin > Site Preferences > Bootstrap.
  - b. Set the Security info Site Preference SSL Host:Port to identify the alias host name and port number.  
  
For more information, see [“Enter security info \(required for SSL only\).” on page 404](#).
  - c. Click **Apply** to save your bootstrap settings.
6. Update the base URL:
  - a. Navigate to Admin > Site Preferences > Refresh.
  - b. Enter the Workplace Base URL value in the provided field. The URL must contain a valid host name, and not contain localhost or an IP number. For example, `http://myserver:7001/Workplace`  
  
For more information, see the IBM FileNet P8 help topic [User Help > Actions, preferences, and tools > Site preferences > Refresh preferences](#).
  - c. Click **Refresh** to update the base URL.
  - d. Click **Exit** to close Site Preferences.
7. Sign out of Workplace, and close your browser.
8. Test the SSL connection by signing into Workplace using the following URL:

```
http://<Application_Engine_server_name>:<non-SSL port#>/Workplace
```

**NOTE** You will be redirected to the SSL-enabled server for sign in, then back to the non-SSL enabled server after sign-in is complete. Before sign-in, you should receive a warning that you are accessing pages over a secure connection (unless you turned this dialog box off), and then Workplace will open.

### ***Additional procedure for WebSphere 5.1***

If you are using Workplace Application Integration to connect to port 80 on WebSphere, you need to enable URL rewriting for SSL sign-in to work properly.

#### **To enable URL rewriting**

---

1. Open the WebSphere administrative console and select **Applications > Enterprise Applications** from the left pane. Then select the Application Engine application (the default is **app\_engine.war**).
2. From Additional Properties, select **Web Container**.
3. Click **Session Management** and click **Edit Properties**.
4. Select **Overwrite** for Overwrite Session Management.
5. Select **Enable URL rewriting** for Session tracking mechanism.
6. Click **OK**, then **Save**, and **Save** again to save changes to the Master configuration.
7. Stop and restart the application server.

### ***Using Java Applets in an SSL Environment***

If you are using a Java applet in an SSL environment, you may experience an SSLHandshakeException because the appropriate certificate does not exist on your computer. Follow the instructions in the the IBM FileNet P8 help topic [User Help > Using Workplace > Basics > Use Java applets](#) to resolve this issue.

## Task 36: Perform Additional Configuration Tasks

Once you have completed the Installation Tasks, your core IBM FileNet P8 system will be up and running. Below is a list of additional configuration tasks you should complete (or at least review) to prepare the system for general use. Except where noted, the links go to the IBM FileNet P8 Help, and start from the Contents panel in:

<Documentation URL, in the form [http://webserver:port#/ecm\\_help/\\_start\\_here.htm](http://webserver:port#/ecm_help/_start_here.htm)

- Configure Content Federation Services for Image Services Guidelines. Refer to the *IBM FileNet P8 Content Federation Services for Image Services Guidelines*. To download this guide from the IBM support page, see “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on page 21.
- Configure Application Engine to set the file types you want to open in a browser window rather than using the Image Viewer. Refer to [FileNet P8 Administration > Application Engine Administration > Key configuration files and logs > content\\_redir.properties file](#).
- Set site preferences for the Workplace application. Refer to [User Help > Actions, preferences, and tools > Site preferences](#).
- Design searches and/or search templates for Workplace users. Refer to [User Help > Actions, preferences, and tools > Tools > Search Designer > About Search Designer](#).
- Design publishing templates for Workplace users. Refer to [User Help > Actions, preferences, and tools > Tools > Publishing Designer > About Publishing Designer](#).
- Configure security for publishing. Refer to [User Help > Actions, preferences, and tools > Tools > Publishing Designer > Security > Specify publication document security](#).
- Configure automatic workflow launch. Refer to [FileNet P8 Administration > Content Engine Administration > Events and subscriptions > Concepts: workflow subscriptions](#).
- Create and configure the object stores that will contain business objects, folders, documents, workflow definitions, searches, and other objects. Refer to [FileNet P8 Administration > Content Engine Administration > Object stores > How to... > Create object store](#).
- Define document classes and folders and set security for each class. Refer to [FileNet P8 Administration > Content Engine Administration > Classes > Concepts](#).
- Review and, if necessary, edit the security of the network shared folders containing any file stores created for the object store. Refer to [FileNet P8 Administration > Content Engine Administration > Content storage > File storage areas](#).
- Configure Process Engine for automatic startup. Refer to [FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Process Service > Start and stop Process Service > Configure the Process Service for automatic startup \(Windows\)](#).
- Configure email notification. Refer to [FileNet P8 Administration > Process Engine Administration > Workflow administration tasks > Coordinating workflow design > Email notification](#).
- Set Process Engine runtime options. Refer to [User Help > Integrating workflow > Process Configuration Console > VWServices > View or modify VWService properties > Set runtime options](#).

- Set the default date/time mask for the Process Service. Refer to [Process Engine > Process Task Manager > Process Service > Configuring Process Service > General properties](#).
- Create content cache area. Refer to [FileNet P8 Administration > Content Engine Administration > Content storage > Content cache areas > How to... > Create content cache](#).
- Create additional authentication realms. Refer to [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > How to > Configure for multiple realms](#).
- Define additional isolated regions. Refer to [User Help > Integrating workflow > Process Configuration Console > Isolated regions](#).
- For each isolated region:
  - Define workflows. Refer to [User Help > Integrating workflow > Process Designer](#).
  - Configure event logging options. Refer to [User Help > Integrating workflow > Process Configuration Console > Isolated regions > View or modify isolated region properties > Configure event logging options](#).
  - Configure step processors. Refer to [User Help > Integrating workflow > Process Configuration Console > Isolated regions > View or modify isolated region properties > Configure custom step processors](#).
  - Define and configure work queues. Refer to [User Help > Integrating workflow > Process Configuration Console > Queues > Configuring work queues](#).
  - Define and configure component queues. Refer to [User Help > Integrating workflow > Process Configuration Console > Queues > Configuring component queues](#).
  - Define and configure workflow rosters. Refer to [User Help > Integrating workflow > Process Configuration Console > Queues > Rosters](#).

# Optional Installation Tasks

## To install optional IBM FileNet P8 components

---

**NOTE** You can install the additional or optional IBM FileNet P8 components listed below in any order.

- Install and configure IBM FileNet Publishing components. Do [Task 37 on page 423](#).
- Enable Process Engine Component Integrator. Do [Task 38 on page 424](#).
- Install Enterprise Manager on a dedicated computer. Do [Task 39 on page 427](#).
- Set up remote Content Engine file storage areas. Do [Task 40 on page 428](#).
- Install Workplace Application Integration. Do [Task 41 on page 429](#).
- Install File Tracker. Do [Task 42 on page 433](#).
- Deploy Multiple Content Engine Instances. Do [Task 43 on page 436](#).
- Deploy Content Engine to Other Application Servers. Do [Task 44 on page 437](#).
- Deploy Multiple Application Engine Instances. Do [Task 45 on page 443](#).
- Install additional Content Search Engine administration servers. Do [Task 46 on page 447](#).
- Enable Application Engine to use ISRA. Do [Task 47 on page 451](#).
- Set up IBM FileNet System Manager. Do [Task 48 on page 456](#).
- Install interim fixes for optional components. Do [Task 49 on page 457](#).

## Task 37: Install and Configure IBM FileNet Publishing

Install the IBM FileNet Rendition Engine to establish publishing capabilities. For instructions, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide* at [FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

## Task 38: Enable the Process Engine Component Integrator

Via the Component Integrator functionality included in the IBM FileNet P8 Platform, a step in a workflow can access properties of documents, folders, and other objects in an object store. Using this functionality requires configuration on both Application Engine and Process Engine, as described in this task.

Post-installation, you will also have to define workflows that incorporate Content Engine (CE) operations in order to use the out-of-the-box Component Integrator functionality. For further details on defining such workflows, see the IBM FileNet P8 help topic [Steps > Component Steps > General Properties > Using Content Engine \(CE\) operations in a workflow](#).

**NOTE** Also post-installation, users can extend the out-of-the-box Component Integrator functionality so that a workflow step can interact with an external entity such as a Java object or JMS messaging system. For further information, see the IBM FileNet P8 help topic [Developing Process Applications > Developing Work Performers / Component Integrator Operations > Developing Component Integrator-Based Workflow Applications](#).

### To specify user name and password for the Java adaptor (on Application Engine)

---

1. Sign in to Workplace.

If you defined the *Process Engine Configuration Group* on the Security tab of Process Task Manager (when completing [“Configure Process Task Manager” on page 340](#)), you must log on as a member of either that group or the *Process Engine Administrators Group*, which was also defined on the Security tab, in order to complete the following steps.

2. In Workplace, click **Admin** and then click **Process Configuration Console**.

**NOTE** If your computer does not have the appropriate Java Runtime Environment (JRE) installed, you will be prompted to download the JRE at this point; follow the prompts to complete the download. During the installation process, click the **Browser** tab and enter the following settings:

- Clear the Internet Explorer option.
- If you will be using Netscape 6.0, select the Netscape option.

For further information about the JRE download, click **Help** in Process Configuration Console, click **Process Reference** on the help page toolbar, and see the IBM FileNet P8 help topic [Concepts > Java Runtime Environment \(JRE\)](#).

3. Select the Isolated Region icon that corresponds to the isolated region you initialized in [“Create a Process Engine Isolated Region” on page 408](#).
4. Right-click the **CE\_Operations** component queue and select **Properties**.
5. On the Adaptor tab of the displayed dialog box, enter a user name and password that will be used for identification and permissions for both Process Engine (PE) and potentially any external systems that will be accessed. By default, the user name and password are set to *Administrator* and *<no password>*, respectively. If you choose to use another user name and password, they must already exist in the directory service.



For additional information about the fields on the Adaptor tab, click the **Help** button. To use the out-of-the-box functionality, it is necessary to modify only the user name and password fields.

6. Click **OK**.

### **To configure and start the Component Manager (on Application Engine)**

---

Execute [Step 1](#) below if Application Engine is configured to use maximum strength symmetric encryption. The JRE used to run the Process Task Manager that contains the Component Manager (which uses JDK 1.4.x) must be updated with Unlimited Strength Jurisdiction Policy Files. Otherwise proceed to [Step 2](#) below.

1. Install unlimited strength .jar files.

**NOTE** Perform this step only if you are using JDK 1.4 or higher and have selected the **Create unlimited strength keys** option in the *Application Engine User Security* and/or *User Token Security* steps of the Application Engine installer. Failure to perform the step will cause EncryptionException messages or other errors indicating that a Java Security API provider for Blowfish is not available. The EncryptionException is caused by the wrong versions of (or absence of) required .jar files that provide unlimited strength security policy files in a Sun JDK 1.4 or higher environment.

For more information, see the IBM FileNet P8 help topic [FileNet P8 Administration > Application Engine Administration > Application Engine Security](#).

- a. Obtain the JDK version specific unlimited strength .jar files.
  - Sun JDK - Obtain the Sun unlimited strength policy files from the Sun product web site (<http://java.sun.com/j2se/>).
  - IBM JDK - Obtain the IBM unlimited jurisdiction policy files from the IBM web site (<http://www.ibm.com/developerworks/java/jdk/security>).

**CAUTION** Make sure you install .jar files specific to the JDK version you are using.

- b. Install the files into the JRE's `/jre/lib/security` folder by replacing files with the same names.
- c. Restart the application server.

2. Start Process Task Manager on Application Engine.

Launch the Process Task Manager from `<AE_install_path>/FileNet/AE/Router`.

- UNIX: Run `routercmd.sh`
- Windows: Run `routercmd.bat`

**NOTE** If the port number assigned to Component Manager conflicts with the port number required by another application or service running on the Application Engine server, then Process Task Manager will not start up as expected. See ["IBM FileNet P8 Port Numbers" on page 643](#) for details on how to resolve this condition.

3. Select **Component Manager** in the left pane (also referred to as the feature pane).
4. Right-click and select **New** to define a new connection point. You will be prompted to enter the Content URI, Service Username, and Service Password to authenticate to the CE.

5. Enter or modify the component properties as appropriate. For details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Application Engine > Component Manager > Configure the Component Manager -> General](#).

**NOTE** In an environment configured for single sign-on (SSO), do not use the SSO server name in the URL, even if Process Task Manager displays it by default.

6. Click **Start** on the toolbar.

---

**To specify connection between Process Engine and Component Manager (on Process Engine)**

1. On Process Engine, start the Process Task Manager.
  - a. (Windows) Select **Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager**.
  - b. (UNIX) Enter **vwtaskman** from the command prompt.
2. Select **Process Engine** in the left pane (also referred to as the feature pane).
3. In the Component Manager connection section, select **Server Connections** tab.
4. In the Host field, enter the host name of the Application Engine server where Component Manager is running.
5. In the Event Port field, enter the port that the Component Manager listens to for incoming events. The default is 32773. The port number you enter must match the number you entered in [Step 3](#) of “[To configure and start the Component Manager \(on Application Engine\)](#)” on page 425.

## Task 39: Install an Additional Instance of Enterprise Manager

Do this task only if you want to install an instance of Enterprise Manager in addition to the one you installed in [“Install Enterprise Manager” on page 183](#).

**CAUTION** Do not install Enterprise Manager 4.0.x on any machine running the 3.5.x version, at least until the Content Engine 4.0.x upgrade is complete. Otherwise, you will no longer be able to run Enterprise Manager 3.5.x against any remaining 3.5.x object stores.

**NOTE** You can install Enterprise Manager only on a Windows machine, and only using the Windows version of the Content Engine installation media.

### To install an additional instance of Enterprise Manager,

---

1. If you have not already done so, install Microsoft .NET Framework 2.0 and Web Services Enhancements (WSE) 3.0. Enterprise Manager on the Windows machine where you are going to install Enterprise Manager.
2. Do a silent or interactive installation of Enterprise Manager as shown in [“Install Enterprise Manager” on page 183](#).
3. Go to [“Install Content Engine Software Updates” on page 177](#) to install service packs, fix packs and/or interim fixes required for Content Engine software.

## Task 40: Create Additional File Storage Areas

Do this task to create additional file storage areas for existing object stores. To create additional fixed storage areas, navigate instead to the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content storage > Fixed storage areas](#).

Do the following procedures for each file storage area you want to create.

### To create a file storage area

---

1. Prepare a location for the file storage area, as shown in “[Prepare Storage Areas for Object Stores](#)” on page 252, and then continue at [Step 2](#).
2. Start Enterprise Manager.
3. Select a FileNet P8 domain and log on as an administrator of the object store in which you will create a file storage area.
4. Right-click the Storage Areas node and then choose New Storage Area.
5. When the Create a Storage Area wizard opens, click **Next** and complete the wizard screens as shown in the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content Engine Wizard Help > Create a Storage Area](#).

Before storing content in the file storage area, do the following procedure to verify that it was properly created.

### To verify the file storage area

---

1. Log on to the machine where Content Engine Server is installed.
2. List the contents of *<fsa1>*, the directory you created or designated on the file server in one of the following procedures:
  - (UNIX) “[To configure a UNIX-based file server](#)” on page 255
  - (Windows) “[To configure a Windows-based file server for a Windows client using CIFS](#)” on page 255
  - (Windows) “[To configure a Windows-based file server for a UNIX client using NFS](#)” on page 256
3. Verify that *<fsa1>* contains an XML file, named **fn\_stakefile**, and two subdirectories, **content** and **inbound**.
4. Verify that *<fsa1>* has the ownership and access permissions you specified.

## Task 41: Install Application Integration

Install Application Integration if you want to integrate IBM FileNet Workplace with your Microsoft Office applications and Outlook. Complete the following procedure on each machine that will use Workplace Application Integration.

**NOTE** You cannot collocate Workplace Application Integration with clients running IDM Desktop Application Integration.

Verify that the client machine meets the platform requirements documented in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### To install the Application Integration software interactively

---

1. Log onto the client machine using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download Application Integration for Microsoft Office**, and then do one of the following:
  - Click **Open** to run the program from its current location.
  - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the ApplicationIntegration.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the installer.

The Welcome Wizard dialog box for Application Integration appears. Another Welcome dialog box appears.

5. Click **Next**.
6. Read the license agreement, and then select **I accept the terms to the license agreement**, and then click **Next**. If you do not accept the license agreement, you cannot continue with the install.
7. Do the following:
  - Select the applications you want to integrate, and then click **Next**.  
**NOTE** The Application Integration Toolkit Components option is required to use Application Integration.
  - Under **Install to**, the default installation path is displayed. Click **Change** to specify a different location on the Change Current Destination Folder dialog box, and then click **OK**. Click **Next**.  
**NOTE** You may see two default installation paths - one for Microsoft Office and Outlook, and another for the Toolkit Components. The Toolkit Components path only appears when the system on which you are installing Application Integration has the Toolkit Components currently installed. You cannot modify the Toolkit Components installation path.
8. Enter the server name, port number and application name that defines the Workplace address. The *server name* is the name of the web server running Workplace, *port number* is the web

server's assigned port, *application* is the directory where you installed the IBM FileNet Workplace application files.

Check **Server uses secure connection (SSL)** if you are running full SSL to encrypt all communication with Workplace.

**NOTE** You can also leave these fields blank and enter the information when you log on to Workplace Application Integration.

9. Click **Next**.
10. Click **Install**.
11. After the install is complete, click **Finish** to complete the setup process.

### **To install the Application Integration software silently**

---

1. Log onto the client machine using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download Application Integration for Microsoft Office**, and then click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the ApplicationIntegration.exe file locally, and then click **Save**.
5. Open a DOS command window and change the current directory to the one where **ApplicationIntegration.exe** resides.
6. Type the following at the command line:

```
ApplicationIntegration.exe /s/v"/qn <additional msi arguments included in string>
LICENSE_ACCEPTED=true"
```

Use the /s switch to launch the execution silently and include the /qn switch in the msi string to make msi run silently.

Refer to the following optional command line values you can also use. Append the values within the string containing the msi arguments.

For example, ApplicationIntegration.exe /s/v"/qn /L\*v C:\temp\AppIntSetup.txt  
 LICENSE\_ACCEPTED=true"

Command Line Values	Installs
ADDLOCAL=ALL	All Features
ADDLOCAL=ALL REMOVE=OutlookIntegrationFeature	Office Only
ADDLOCAL=ALL REMOVE=OfficeIntegrationFeature	Outlook Only
ADDLOCAL=ALL REMOVE=OutlookIntegrationFeature, OfficeIntegrationFeature	Core Only

Command Line Values	Settings
HOST=<host name>	Enter the name of the web server running Workplace.
PORT=<port number>	Enter the web server's assigned port number.
APPLICATION=<application name>	Enter the directory in which you installed the Workplace application files.
SERVER_CONNECTION=1	Set Application Integration to use an https connection
SERVER_CONNECTION=0	Set Application Integration to use http connection. This is the default if this parameter is not passed.
/L*v C:\temp\AppIntSetup.txt	Verbose installation log and specify log location.

**To verify your Workplace Application Integration installation**

1. Start Microsoft Word.
2. From the **File** menu, click **FileNet P8**, click **Open Document**, and then click **Select Item**. The Logon dialog box opens.
3. Log on using any valid domain account. The available object stores in your environment are displayed.

**NOTE** If you didn't enter the Workplace Address information in [Step 8](#) above, enter the server name, port number and application name that defines the Workplace address. The *server*

*name* is the name of the web server running Workplace, *port number* is the web server's assigned port, *application* is the directory where you installed the IBM FileNet Workplace application files. Check **Server uses secure connection (SSL)** if you use a full SSL to encrypt all communication with Workplace. Do not select this option if you use a SSL redirect during login.

4. Close all dialog boxes and close Microsoft Word.

### **To uninstall or modify Workplace Application Integration**

---

1. From the **Start** menu, click **Settings**, and then click **Control Panel**.
2. Click **Add/Remove Programs**, and then click **FileNet Workplace Application Integration 4.0**.
3. Do one of the following:
  - Click **Remove**, and then click **Yes** to confirm you want to uninstall Workplace Application Integration.
  - Click **Change** to access maintenance tasks, and then click **Next**. You can modify, repair, or remove Application Integration using the maintenance tasks.

Do one of the following:

- Select **Modify** to add or remove integration with Microsoft applications from your previous install. For example, if you have both Microsoft Office and Outlook installed, you can remove one of the applications using this option. The Custom Setup dialog box appears, where you highlight the option you want to add or remove. Click **Next**, and then click **Install**. Click **Finish** to complete the process.
- Select **Repair** to re-install Workplace Application Integration to repair installation errors, and then click **Next**. Click **Install** to start the repair process. Click **Finish** to complete the process.
- Select **Remove** to remove Workplace Application Integration from your system, and then click **Next**. Click **Remove**. Once the application is removed from your system, click **Finish** to complete the process.

### **To silently uninstall Workplace Application Integration**

---

1. Open a command prompt.
2. Enter the following command to uninstall Workplace Application Integration:

```
msiexec.exe /X{35907B7D-02E2-490C-8F3B-54C4E3729D90} /qn
```



## Task 42: Install File Tracker

Install File Tracker if you want to use the Workplace file tracking feature without installing Application Integration. Complete the following procedure on each machine that will use Workplace File Tracker.

### NOTES

- If you have already installed or upgraded to Application Integration 3.5.1-002 or higher, then the File Tracker feature has already been installed. Do not perform this procedure if you already installed a version of Application Integration that includes File Tracker, including Application Integration 4.0.
- If you have already installed Application Integration 3.5.1-001 or earlier, upgrade it to 4.0 plus the latest fix pack, and this will include the File Tracker installation. For details about the upgrade process, see [“Upgrade Application Integration and File Tracker” on page 560](#).

Verify that the client machine meets the platform requirements documented in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### To install the File Tracker software interactively

---

1. Log onto the client machine using an account that has Administrator privileges.
  2. Log into Workplace.
  3. Click **Author**, and then click **General Tools**.
  4. Scroll down and click **Download File Tracker** and do one of the following:
    - Click **Open** to run the program from its current location.
    - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the FileTracker.exe file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the installer.
- The Welcome Wizard dialog box for File Tracker appears. Another Welcome dialog box appears.
5. Click **Next**.
  6. Read the license agreement, and then select **I accept the terms to the license agreement**, and then click **Next**. If you do not accept the license agreement, you cannot continue with the install.
  7. Do one of the following:
    - Click **Change** if you want to install File Tracker to a different location. Specify the location to which you want to install File Tracker, and then click **OK**. Click **Next**.
    - Click **Next** to accept the default location.
  8. Click **Install**.
  9. After the install is complete, click **Finish** to complete the setup process.

### To install the File Tracker software silently

---

1. Log onto the client machine using an account that has Administrator privileges and log into Workplace.
2. Click **Author**, and then click **General Tools**.
3. Scroll down and click **Download File Tracker** and click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the FileTracker.exe file locally, and then click **Save**.
4. Open a DOS command window and change the current directory to the one where **FileTracker.exe** resides.
5. Type the following at the command line:

```
FileTracker.exe /s /v"/qn <additional msi arguments included in string>
LICENSE_ACCEPTED=true"
```

**NOTE** Use the /s switch to launch the execution silently and include the /qn switch in the msi string to make msi run silently. In addition, be aware of spaces specified between switches in the above example. Using the correct spacing ensures a successful silent install.

Refer to the following optional command line values you can also use. Append the values within the string containing the msi arguments.

For example:

```
FileTracker.exe /s /v"/qn /L*v C:\temp\FileTrackerSetup.txt LICENSE_ACCEPTED=true"
```

Command Line Values	Settings
/L*v C:\temp\FileTrackerSetup.txt	<p>Verbose installation log and specific log location.</p> <p><b>NOTE</b> If you intend to specify a log location, create the directory before running the silent install. If the directory is not created ahead of time, the install will fail.</p>

### To uninstall Workplace File Tracker

---

1. From the **Start** menu, click **Settings**, and then click **Control Panel**.
2. Click **Add/Remove Programs**, and then click **FileNet Workplace File Tracker**.
3. Click **Remove**, and then click **Yes** to confirm you want to uninstall Workplace File Tracker.

### **To silently uninstall Workplace File Tracker**

---

1. Open a command prompt.
2. Enter the following command to uninstall Workplace File Tracker:  
`msiexec.exe /X{4291FBBC-C585-43ED-9416-5F22D8C6FEE9} /qn`

## Task 43: Deploy Multiple Content Engine Instances

You can deploy multiple instances of Content Engine on a single application server. Moreover, the instances do not need to all have the same database type or directory server.

To prepare for deploying a subsequent instance of Content Engine on an application server, you will do only the applicable tasks in [“Prerequisite Tasks” on page 53](#) and [“Installation Tasks” on page 128](#).

1. As needed, configure the directory server for your subsequent Content Engine instance.
2. Do the database task for your subsequent Content Engine instance.

For example, if the initial instance uses DB2 for its database, and a subsequent instance will use MS SQL Server, do the prerequisite MS SQL Server task. Even if your Content Engine instances will all use the same database type, you must create at least a GCD database/tablespace for each instance.

3. For *each* Content Engine instance, set up a dedicated configuration on your particular application server, as follows:
  - Create and configure a WebSphere profile, as shown in [“Configure an Application Server for Content Engine \(WebSphere\)” on page 106](#).
  - Create and configure a WebLogic domain, as shown in [“Configure an Application Server for Content Engine \(WebLogic\)” on page 110](#).
  - Set up and configure a JBoss server, as shown in [“Configure an Application Server for Content Engine \(JBoss\)” on page 118](#).
4. [“Install and Deploy Content Engine” on page 160](#).
5. Do all subsequent Content Engine-related tasks in [“Installation Tasks” on page 128](#).

## Task 44: Deploy Content Engine to Other Application Servers

After doing the first installation and deployment of Content Engine on an application server, you can deploy Content Engine to additional application servers of the same type (for example, if the first deployment of Content Engine is on a WebSphere server, each additional deployment of Content Engine must be on a WebSphere server).

If you deployed Content Engine to a server that ...	Continue at...
Does <i>not</i> function as an administrative server for your application server environment,	<a href="#">“Deploy Content Engine Into a Non-Managed Server Environment” on page 437</a> to deploy Content Engine to additional machines.
Does function as an administrative server for your application server environment	<a href="#">“Deploy Content Engine Into a Managed Server Environment” on page 439</a> to redeploy Content Engine to managed server instances that will run Content Engine.

### Deploy Content Engine Into a Non-Managed Server Environment

Use one of the following procedures to deploy Content Engine from one (*source*) application server to another (*destination*) application server

#### To deploy Content Engine to another WebSphere server

1. Create a profile for Content Engine on the destination WebSphere server (see [“Configure an Application Server for Content Engine \(WebSphere\)” on page 106](#)).
2. Copy the file **Engine-ws.ear** (the Content Engine EAR file) from the source WebSphere server to the destination WebSphere server, for example (on a UNIX machine) **/opt/FileNet/ContentEngine/lib/Engine-ws.ear**.
3. Log on to the WebSphere administrative console on the destination WebSphere server.
4. Deploy Content Engine, as follows:
  - a. Navigate to Applications > Install New Application.
  - b. Click **Local file system** and browse to the file **Engine-ws.ear file** and click **Next**.
  - c. Accept defaults for all options except the following: provide options to perform the EJB deploy
  - d. Accept defaults for most options except the following:
  - e. Provide options to perform the EJB Deploy: Change Database Type to the desired database (e.g., ORACLE\_v9).
  - f. At the last page, click **Finish** and the application server will start the actual deployment.

- g. Update the master configuration by selecting Save on the toolbar then click the **Save** button.
5. Assign FileNet configuration property values, as follows:
  - a. Navigate to Applications > Enterprise Applications > FileNetEngine.
  - b. Change the following configuration settings to the values shown and then click **Apply**:
    - Classloader Mode: PARENT\_LAST
    - WAR Classloader Policy: Application
  - c. Click **Apply**.
6. Assign values to the connection factory properties, as follows:
  - a. Navigate to Applications > FileNetEngine > Connector Modules > engine.rar > Resource Adaptor.
  - b. Click **J2C Connection Factories**.
  - c. Click **New**.
  - d. Set the connection factory values, as follows:
    - Name: FileNetConnectionFactory
    - JNDI: FileNet/Local/ConnectionFactory
    - Authentication Preference: None
    - Component-managed Authentication Alias: *<node name>/<J2C connection alias for GCD DB user>*
  - e. Click **Apply**.
7. Assign values to the connection pool properties, as follows:
  - a. Click **Connection Pool**.
  - b. Set the connection pool values, as follows:
    - Connection Timeout: 1000
    - Max Connections: 100
    - Min Connections: 10
  - c. Click **Apply**.
8. Update the master configuration by selecting Save on the toolbar then click the **Save** button
9. Navigate to **Applications > Enterprise Applications** and verify that FileNet Engine is listed.
10. Click **Start** to verify that Content Engine Server starts.
11. (Optional) To access Centera fixed content devices from this server, continue at ["Install Centera Shared Libraries" on page 178](#).

### To deploy Content Engine to another WebLogic server

---

1. Create a domain for Content Engine on the destination WebLogic server (see [“Configure an Application Server for Content Engine \(WebLogic\)”](#) on page 110).
2. Locate the file **Engine-wl.ear** on the source WebLogic server.
3. Copy the file **Engine-wl.ear** to a directory in the domain you created in the destination WebLogic server in [Step 1](#).
4. Stop and start the destination WebLogic server.
5. (Optional) To access Centera fixed content devices from this server, continue at [“Install Centera Shared Libraries”](#) on page 178.

### Deploy Content Engine Into a Managed Server Environment

Use one of the following procedures to redeploy Content Engine from an administrative server to a managed server instance.

#### To redeploy Content Engine to a managed WebSphere server

---

1. Start the WebSphere administrative console (if it is not already running).
2. Undeploy the Content Engine using WebSphere administrative console on the deployment manager node, as follows:
  - a. Navigate to Applications > Enterprise Applications, select the FileNet Engine application, and click **Uninstall**.
  - b. Save your changes and synchronize the configuration with all nodes in the cell.
  - c. Restart the deployment manager instance, all node agents, and application server instances in the cell which will host the Content Engine application. Wait until all instances have restarted before continuing to next step.
3. (WebSphere 5.1.1 only) Deploy the Content Engine application to the managed nodes in the application server configuration, as follows:
  - a. Navigate to Applications > Install New Application , select the file **Engine-ws.ear**, or type the full path to the file, and then click **OK**.

By default, **Engine-ws.ear** has the following path:

    - (UNIX) **<Content Engine Install Path>/FileNet/ContentEngine/lib**
    - (Windows) **<Content Engine Install Path>\FileNet\Content Engine\lib**
  - b. Accept the defaults in the *Preparing for the application installation/You can choose to generate default bindings and mappings* screen.
  - c. Accept the defaults in the screen *Install New Application\Step 1:Provide options to perform the installation*.

- d. On the screen *Install New Application\Step 2:Provide options to perform the ejb deploy*, select the database type used in the configuration from the *Deploy EJBs Option - Database Type* drop-down box, and then click **Next**.

Select the check box next to each of the modules listed, and then click **Apply**. The server field for each module should now appear with the values selected in the Clusters and Servers selection box. Click **Next**.

- e. Accept the defaults in the screen *Install New Application\Step 3:Provide JNDI Names for Beans*.
  - f. Accept the defaults in the screen *Install New Application\Step 4:Map EJB references to beans* screen.
  - g. Accept the defaults for the screen *Install New Application\Step 5:Map virtual hosts for web modules*.
  - h. On the screen *Install New Application\Step 6:Map modules to servers*, select the server names that will provide access to Content Engine from the Clusters and Servers selection box.
  - i. On the screen *Install New Application\Step 7:Summary*, review the summary and click **Finish**.
4. (WebSphere 6.0.2 only) Deploy the Content Engine application to the managed nodes in the application server configuration, as follows:
    - a. Navigate to Applications > Install New Application, select the file **Engine-ws.ear**, or type the full path to the file, and then click **OK**.  
 By default, **Engine-ws.ear** has the following path:
      - (UNIX) **<Content Engine Install Path>/FileNet/ContentEngine/lib**
      - (Windows) **<Content Engine Install Path>\FileNet\Content Engine\lib**
    - b. Accept the defaults in the screen *Preparing for the application installation/Choose to generate default bindings and mappings*.
    - c. Accept the defaults in the screen *Install New Application\Step 1:Select installation options*.
    - d. On the screen *Install New Application\Step 2:Map modules to servers*, select the server name, and any web servers that will provide access to Content Engine from the Clusters and Servers selection box.  
 Select the check box next to each of the modules listed, and then click **Apply**. The server field for each module should now show up with the values selected in the Clusters and Servers selection box. Click **Next**.
    - e. On the screen *Install New Application\Step 3:Provide options to perform the ejb deploy*, select the database type used in the configuration from the *Deploy EJBs Option - Database Type* drop-down box, then click **Next**.
    - f. Accept the defaults in the screen *Install New Application\Step 4:Provide JNDI Names for Beans*.



- g. Accept the defaults in the screen *Install New Application\Step 5:Map JCA resource references to resources*.
  - h. Accept the defaults in the screen *Install New Application\Step 6:Map EJB references to beans*.
  - i. Accept the defaults in the screen *Install New Application\Step 7:Map virtual hosts for web modules*.
  - j. On the screen *Install New Application\Step 8:Summary*, review the summary and click **Finish**.
5. Perform post-deployment tasks, as follows:
- a. After the application has been loaded and configured in WebSphere, save the changes and synchronize all nodes.
  - b. Continue at [Step 5 on page 438](#).

**To redeploy Content Engine to a managed WebLogic 8.1.x server**

---

1. If they aren't already running, start all instances of WebLogic Server in the WebLogic domain, including the administrative server.
2. From the WebLogic Server Administration Console, go to the Target & Deploy tab for each of the connection pools and data sources created for the GCD database.
3. Select the check box next to the WebLogic instances that will be running Content Engine, clear the check box next to the administration server name, and click **Apply**. This will redeploy Content Engine to the managed servers (WebLogic instances) you selected.
4. Select Engine-wl from the WebLogic Administration Console.
5. Select the Targets tab for the application, and select the check boxes next to the WebLogic instances that will be running Content Engine. Clear the check box for the administration server name and click **Apply**.
6. Restart all administration server instances and managed nodes.

**To redeploy Content Engine to a managed WebLogic 9.2.x server**

---

1. If they aren't already running, start all instances of WebLogic Server in the WebLogic domain, including the administrative server.
2. From the WebLogic Server Administration Console, go to the Target tab for data sources created for the GCD database and click **Lock & Edit**
3. Select the check box next to the WebLogic instances that will be running Content Engine and clear the check box next to the administration server name. This will redeploy content Engine to the managed servers (WebLogic instances) you selected.
4. Select Engine-wl from the WebLogic Administration Console.
5. Select the Targets tab for the application and select the check boxes next to the WebLogic instances that will be running Content Engine. Clear the check box for the administration server name.

6. Click **Activate Changes**.
7. Restart all administration server instances and managed nodes.

## Task 45: Deploy Multiple Application Engine Instances

This topic covers deployment of multiple instances of Workplace on a single application server. Each deployment of Workplace must use the same Content Engine, Process Engine, and connection point. Each deployment of Workplace may use different Site Preference settings and may provide access to different object stores.

### NOTES

- The following procedure assumes that you have already installed Application Engine and performed the following configuration tasks according to your application server type:
  - “Configure Application Engine (WebSphere)” on page 356
  - “Configure Application Engine (WebLogic)” on page 373
  - “Configure Application Engine (JBoss)” on page 381
- When deploying multiple instances of Workplace, make copies of all the Workplace configuration and working files. Each instance of Workplace will use separate configuration, deploy, download, upload, and Workplace files. Leave the default installed files unmodified.
- For more information on how to deploy and manage multiple identical applications, see your application server documentation.

### To deploy a second instance of the Workplace application

---

1. Make a copy of the **/FileNet/Config/AE** directory, including all of its contents, for each instance you plan to deploy. For example, if you are deploying two instances, you would create
  - <install\_path>/FileNet/Config/AE1**
  - <install\_path>/FileNet/Config/AE2**
2. Make copies of the upload and download directories in the **<install\_path>/FileNet/AE** directory. For example, you would create:
  - <install\_path>/FileNet/AE/download1**
  - <install\_path>/FileNet/AE/upload1**
  - <install\_path>/FileNet/AE/download2**
  - <install\_path>/FileNet/AE/upload2**
3. Make a copy of the deploy directory and all of its contents for each Workplace instance. For example, you would create:
  - <install\_path>/FileNet/AE/deploy1**
  - <install\_path>/FileNet/AE/deploy2**

4. Make a copy of the Workplace directory and all of its contents for each Workplace instance.

For example, you would create:

**<install\_path>/FileNet/AE/Workplace1**

**<install\_path>/FileNet/AE/Workplace2**

5. Navigate to each custom copied Workplace web.xml instance and update the path for the configuration directory, upload directory, and download directory locations.

For example, in the **<install\_path>/FileNet/AE/Workplace1/WEB-INF/web.xml**, you would make the following changes (in **bold**):

```
<context-param>
  <param-name>configurationDirectory</param-name>
  <param-value>/opt/FileNet/Config/AE1</param-value>
</context-param>

<context-param>
  <param-name>uploadDir</param-name>
  <param-value>/opt/FileNet/AE/Upload1</param-value>
</context-param>
<context-param>
  <param-name>downloadDir</param-name>
  <param-value>/opt/FileNet/AE/Download1</param-value>
</context-param>
```

### To deploy each additional Workplace instance as an EAR file

**NOTE** Perform the following steps for each custom Workplace instance you plan to deploy.

1. Modify the **application.xml** file located in the copied deploy directories:
  - a. Open each instance of the application.xml file, for example, **<install\_path>/FileNet/AE/deploy1/META-INF/application.xml**.
  - b. Change the <display-name> and the <context-root> elements to your custom name, for example, Workplace1 (shown in **bold**, below).

```
<display-name>Workplace1</display-name>
<description>FileNet Application Engine</description>

<module>
  <web>
    <web-uri>app_engine.war</web-uri>
    <context-root>Workplace1</context-root>
  </web>
</module>
```

2. In the **create\_app\_engine\_war** file, change the *install\_home* path and the *deploy* directory to match your custom names.

For example, you would make the following changes, shown in **bold**.

```
install_home="/opt/FileNet/AE/Workplace1"
"${install_home}/../_AEjvm/bin/jar" -cf "${install_home}/../deploy1/app_engine.war"*
```

3. In the **create\_app\_engine\_ear** file, set the install home, deploy directory, and EAR file to match your custom names.

For example, you would make the following changes, shown in **bold**.

```
install_home="/opt/FileNet/AE/Workplace1"
cd "${install_home}/../deploy1"
"${install_home}/../AEjvm/bin/jar" -cvf "${install_home}/../deploy1/
app_engine1.ear" META-INF *.war
```

4. Delete the existing **app\_engine.war** and **app\_engine.ear** files.
5. Create your custom WAR and EAR files by running the **create\_app\_engine\_war** and then the **create\_app\_engine\_ear** files.
6. Deploy the EAR file for each custom Workplace instance according to the procedures provided for your application server type:
  - [“Deploy Application Engine \(WebSphere\)” on page 385](#)
  - [“Deploy Application Engine \(WebLogic\)” on page 391](#)
  - [“Deploy Application Engine \(JBoss\)” on page 394](#)

### To apply Application Engine software updates to custom Workplace instances

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your support representative.
2. Open the readmes for any subsequent fix packs or interim fixes (typically optional) and perform the installation procedures provided.
3. After the updates have been installed, you will need to check if the configuration files have been updated. If they have been updated, a **backup-401-<patch>** directory will have been created in the **<install\_path>/FileNet/AE** directory. The master configuration files are located in the **/FileNet/Config/AE** directory.
4. If the configuration files have been updated by the Application Engine software updates, you must run the **repatch.sh** script to update the custom configuration files for each custom Workplace instance.

**NOTE** As an initial step, prior to running **repatch.sh**, you should back up the existing **.old** files (don't delete them) from the "master" patch install. These files are located in the **backup-401-<patch>** directory.

5. Copy the **repatch.sh** file into the **backup-401-<patch>** directory for the version you installed.
6. Execute the **repatch.sh** from a command prompt:

```
./repatch.sh <install_path>/FileNet Workplace1
```

7. Rerun the script for each custom instance of Workplace.

**NOTE** The ".old" files for each custom deployment will be backed up automatically by the script.

8. Copy the **recopy.sh** file to **<install\_path>/FileNet/AE**

9. Execute the **recopy.sh** from a command prompt:  
**./recopy.sh Workplace1**
10. Rerun the script for each custom instance of Workplace.
11. Navigate to the custom **/FileNet/AE/deploy** directories and delete the existing **app\_engine.ear** and **app\_engine.war** files
12. Create the WAR and EAR files which now will contain updated files by running the **create\_app\_engine.war.sh** file and then the **create\_app\_engine.ear.sh** file.
13. Repeat [Step 11](#) and [Step 12](#) for each custom Workplace instance.
14. Deploy the EAR file for each custom Workplace instance according to the procedures provided for your application server type:
  - [“Deploy Application Engine \(WebSphere\)” on page 385](#)
  - [“Deploy Application Engine \(WebLogic\)” on page 391](#)
  - [“Deploy Application Engine \(JBoss\)” on page 394](#)

## Task 46: Install Additional Content Search Engine Servers

This task describes how to install and configure additional Autonomy K2 Administrative Servers for IBM FileNet P8 Content Search Engine, an optional component based on the Autonomy K2 product. This task is required for both new installations and upgrades. In effect, you will set up Autonomy K2 Administration Servers required for any multi-server K2 configuration.

**NOTE** To install the Content Search Engine software silently, you can create your own script to run the command-line steps presented in the procedures in this topic.

### To install an Autonomy K2 Administration Server on Windows

---

1. Create the required accounts and groups and related permissions as specified in [“To create Content Search Engine accounts”](#) on page 60.

2. Access the host machine and log on to the directory service as K2 Operating System User .

**NOTE** Ensure K2 Operating System User is an operating system administrator on the host machine.

3. Insert the IBM FileNet Autonomy K2 installation CD and extract the contents of **K2-win.zip** to the following location:

**C:\Program Files\filenet\contentengine\verity**

4. Set the Java\_Home environment variable as follows:

- a. Open the System control panel.
- b. Click the **Advanced** tab.
- c. Click **Environment Variables**.
- d. Click **New**.
- e. Set the variable information as follows:

Variable name: **Java\_Home**

Variable value: `<Java1.5.0xx_JDK_install_path>`

5. Open **C:\Program Files\filenet\contentengine\verity\config.vcnf** in a text editor and make the following modifications in the file:

- Replace the instance of `<myMode>` with Local.
- Replace all instances of `<myLocalHostName>` with the name of the machine on which you are installing.
- Replace all instances of `<myMasterHostName>` with the machine name that the K2 Master Administration Server is installed on.
- Replace all instances of `<installDir>` with **C:\Program Files\filenet\contentengine\verity**.
- Replace the instance of `<JavaHome>` with the path to the installed Java 1.5.0\_xx JDK.

6. Open a command line and change directory to **C:\Program Files\filenet\contentengine\verity**.
7. Enter the following at the command line and substitute the applicable value for each variable:

```
k2\_nti40\bin\vconfig -cfg "c:\program
files\filenet\contentengine\verity\config.vcnf" -dir "c:\program
files\filenet\contentengine\verity" -username "<K2 Security User>" -pwd "<password>"
-domain "<domain>" -verbose -log log.txt
```

The Autonomy K2 Administration Server service will be installed and running at the completion of the vconfig command.

8. Set K2 Administration Server service to run as K2 Operating System User, as follows:
  - a. Access Component Services.
  - b. Stop the Verity K2 6.1.1 Administration Server service.
  - c. Change the logon settings and set the service to Log On as K2 Operating System User.
  - d. Start the Verity K2 6.1.1 Administration Server service.
9. See [“Configure Content Engine for Content-Based Retrieval” on page 397](#) to configure this Autonomy K2 installation.

**To install an Autonomy K2 Administration Server on UNIX**

**NOTE** For HP-UX installations, manually configure the kernel with following parameters:

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

1. Create the required accounts and groups and related permissions as specified in [“To create Content Search Engine accounts” on page 60](#).



2. Insert the IBM FileNet Autonomy K2 installation CD and extract the contents of **K2-*<platform>.tar.gz*** to **/opt/verity** using the following commands:

- a. `gzip -d <platform>.tar.gz`
- b. `tar -xvf <platform>.tar`

**NOTE** If you install to a directory other than **/opt/verity**, you must create a soft link using the following command:

```
ln -s <InstallPath> /opt/verity
```

3. Change directory to **/opt/verity**.
4. Edit **/opt/verity/config.vcnf** as follows:
  - Replace the instance of *<myMode>* with Master.
  - Replace all instances of *<myLocalHostName>* with the name of the machine on which you are installing.
  - Replace all instances of *<myMasterHostName>* with the name of the machine on which you are installing.
  - Replace the instance of *<JavaHome>* with the path to the installed Java 1.5.0\_xx JDK.
5. Set the Java\_home variable:

```
JAVA_HOME=/usr/java/jdk1.5.0_09_xx
export JAVA_HOME
```

**NOTE** Enter the Java\_home variable in the .profile file to set this variable each time the user logs in.

6. Append the following environment variables according to your platform:

(HP-UX)

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

(AIX)

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

(Solaris)

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

```
export LD_LIBRARY_PATH
```

(Linux)

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

7. Change directory to **/opt/verity/** and enter:

```
k2/<platform>/bin/vconfig -cfg "/opt/verity/config.vcnf" -dir "/opt/verity" -
verbose -log log.txt
```

Substitute one of the following for *<platform>*:

- `_ssol26` (Solaris 8.0, 9.0 or 10.0)
- `_hpux` (HP-UX 11i with -AA compiler flag)
- `_rs6k43` (AIX 5.2 and 5.3)
- `_ilnx21` (Red Hat Advanced Server 3.0 and 4.0, SUSE 8 and 9)

The Autonomy K2 Administration Server service and Tomcat will be installed and running at the completion of the `vconfig` command.

---

**To configure Autonomy K2 Administration Servers from the K2 Dashboard**

If you are adding Content Search Engine to an already functioning and updated IBM FileNet P8 Platform system, skip to [Task 18 Install Content Search Engine Software Updates](#) to apply any available service packs, fix packs and/or interim fixes and then complete [Task 29 Configure Content Engine for Content-Based Retrieval](#). Otherwise, move on to [Task 9 "Install and Deploy Content Engine" on page 160](#).

## Task 47: Enable Application Engine to Use ISRA

Image Services Resource Adapter (ISRA) is a J2EE connector to the FileNet Image Services (IS) libraries. Using ISRA, Workplace users can view IS documents and their associated annotations in the FileNet P8 Image Viewer and, if they have the appropriate permissions, update the annotations.

To enable Workplace users to view documents using ISRA, the following steps must be completed:

- Install the Application Engine.
- Install FileNet ISRA.

For information on installing, configuring and deploying FileNet ISRA, refer to the Image Services Resource Adapter documentation on the FileNet ISRA Installation CD.

**TIP** Use the Sample Application shipped with FileNet ISRA to confirm that the ISRA installation was successful.

**WARNING** In an ISRA upgrade situation, take care to use the same library name (JNDI connection factory name) that has been previously set in the ISRA install. Changing this variable can cause conflicts when accessing documents.

- Install the Application Engine ISRA Servlet.
  - Deploy the servlet on the same application server as FileNet ISRA.
  - You need not the servlet with the Application Engine.
  - Install and deploy ISRA before installing and deploying the ISRA Servlet.
- Configure Workplace Site Preferences.

### ISRA SSL support

The following table details supported SSL configurations for ISRA.

SSL Configuration	SSL Support
ISRA Servlet and AE Collocated. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet and AE Collocated. AE Configured for SSL logon redirect to a local host.	Supported
ISRA Servlet and AE Collocated. AE, and ISRA Servlet running under SSL.	Not Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a local host.	Supported

SSL Configuration	SSL Support
ISRA Servlet remote from AE. AE running under SSL, ISRA Servlet not running under SSL.	Supported
ISRA Servlet remote from AE. AE, and ISRA Servlet running under SSL.	Not Supported

### To install and deploy the Application Engine ISRA Servlet

The FileNet P8 Application Engine installation software contains the ISRA servlet installation programs for the supported P8 AE operating systems.

1. Log on to the application server.

UNIX - logon as a user with write access to the **/bin** directory and read, write, execute access to the directory where you plan to install ISRA Servlet.

Windows - logon as a member of the local Administrators group or as a user with equivalent permissions.

2. Stop the application server if it is running.
3. Access the ISRA installation package, and start the Application Engine ISRA Servlet Setup wizard:

UNIX - Execute **<Platform>filenet\_ae\_israservlet\_setup.bin**.

Windows - Execute **WINfilenet\_ae\_israservlet\_setup.exe**

4. Complete the Setup screens as follows:

In this screen...	Perform this action...
License Agreement	Review and accept the license agreement for FileNet P8 software.
Directory Name	<p>For the Directory Name field, enter or browse to the location where you want to install the ISRA Servlet, or accept the default location:</p> <ul style="list-style-type: none"> <li>• UNIX - <b>/opt</b></li> <li>• Windows - <b>C:\Program Files\</b></li> </ul> <p><b>UNIX</b></p> <pre>&lt;AE_israservlet_install_path&gt;/FileNet/ ApplicationEngineISRAServlet&gt;</pre> <p><b>Windows</b></p> <pre>&lt;AE_israservlet_install_path&gt;\FileNet\ApplicationEngineI SRAServlet&gt;</pre>

In this screen...	Perform this action...
Create war File	Check this box if you use a WebSphere application server.  If the box is checked, the installation program will create the ae_isra.war file and a script that can also generate the ae_isra.war file.
User Token Security	Make a note of the user token crypto key path.
Ready to Install	Verify your selections, and click <b>Next</b> to install the ISRA Servlet.
Completing the Setup	Click <b>Finish</b> to complete the ISRA Servlet installation Wizard.

5. Check the file **filenet\_ApplicationEngineISRAServlet\_install\_log.txt**, located in the **<AE\_israservlet\_install\_path>\FileNet** directory, to see if any errors occurred during the installation.
6. Install unlimited strength jar files.  
  
Perform this step only if you selected the **Create unlimited strength key** option in the Application Engine User Token Security step of the Application Engine installation and the Application Engine ISRA Servlet is deployed on a different application server from the Application Engine. Failure to perform the step causes an EncryptionException when you login to the IS Server.
7. (WebSphere 5.x only) To allow users to save annotations through ISRA, copy the **<AE\_israservlet\_install\_path>\FileNet\jar\ISRA.jar** file to the **<AE\_install\_path>\FileNet\AE\Workplace\WEB-INF\lib** directory.
8. Start the application server.
9. Deploying the Application Engine ISRA Servlet is similar to deploying Workplace.
  - WebSphere:  
  
Deploy **<AE\_israservlet\_install\_path>\FileNet\ApplicationEngineISRAServlet\ae\_isra.war** in the same way you deployed the **app\_engine.war** file for Workplace.
  - WebLogic:  
  
Deploy **<AE\_israservlet\_install\_path>\FileNet\ApplicationEngineISRAServlet** in the same way you deployed **<AE\_install\_path>\FileNet\AE\Workplace**.
10. Stop and restart the application server.
11. Verify the Application Engine ISRA Servlet installation.  
  
A diagnostic tool is available to verify that the ISRA Servlet is installed and deployed correctly.
  - a. Launch your browser.
  - b. Enter the URL for the Application Engine ISRA Servlet, for example,  
  
`http://<ApplicationEngineISRAServlet_servername>:<port>/ApplicationEngineISRAServlet/ISRA`

**NOTE** ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

If the ISRA Servlet is installed and deployed correctly, a Congratulations message displays. For example:

```
Congratulations! ISRA Interface Servlet is configured at this URL.  
WcmApiConfigFile = D:\ISRAInterface\jsp\WEB-INF\WcmApiConfig.properties  
WcmApiConfig file exists  
  
CryptoKeyFile/UserToken = C:\Program  
Files\FileNet\Authentication\UTCryptoKeyFile.properties  
CryptoKeyFile/UserToken exists  
  
FileNet ISRA classes are in the classpath  
com.filenet.is.ra.cci.FN_IS_CciConnectionSpec
```

## To configure the Workplace site preferences

The Application Engine setup installs a pre-configured external service called Image Service, which includes the parameterized values necessary to access FileNet IS libraries from Workplace. Enable the service by setting the *Image Service* value in Site Preferences to **Show** (the default is Hide), as described in the following procedure.

### To enable the pre-configured Image Service

---

1. Sign in to Workplace as a user having the Application Engine Administrators access role.
2. Launch Site Preferences as follows:
  - a. Select **Admin**.
  - b. Select **Site Preferences**.
3. Select **External Services** from the left options list.
4. Select **Modify** for the Image Service (under External Reference Services).

The External Reference Service Settings site preference page displays.
5. Under General Information, locate **Show on Select File page** and change the value to **Show**.
6. Click **Accept**.
7. Click **Apply**.

## To set the ISRA Interface Servlet URL

1. Select **Bootstrap**.

2. Under Preferences Settings, set the value of ISRA Interface Servlet URL. For example:

`http://<servername>:<port>/ApplicationEngineISRAServlet/ISRA`

**NOTE** ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

3. Click **Apply**.
4. Click **Exit** to exit Site Preferences.

### ***To log on to Image Services via LDAP***

To log on to the Image Services library using your LDAP account, configure ISRA and Image Services for LDAP authentication. If the LDAP account with which you accessed Workplace is not valid for the Image Services library, or if LDAP authentication is not configured, you will be prompted to log on to the Image Services library.

For information on configuring LDAP authentication for ISRA, refer to the ISRA Installation and Deployment Guide. For information on configuring LDAP authentication for Image Services, refer to the Image Services System Tools Reference Manual.

### ***To access IS library documents***

For information about accessing IS library documents, see [User Help > Actions, preferences and tools > Actions > Documents > Add a document \(Workplace\)](#).

## Task 48: Install and Configure IBM FileNet System Manager

Content Engine, Application Engine, and Process Engine install, by default, the necessary software required for the System Manager performance component. To use System Manager, enable associated components and install IBM FileNet Dashboard to perform related configuration procedures to enable System Manager. Installing Dashboard is not necessary if you currently have IBM FileNet System Monitor installed.

Refer to the following IBM FileNet P8 help topic [FileNet P8 Documentation > FileNet P8 Administration > Enterprise-wide Administration > System Manager](#) for instructions on how to enable the associated System Manager components.

Refer to the documentation provided with IBM FileNet Dashboard for instructions on how to use Dashboard.



## Task 49: Install Software Updates for Optional Components

Install any service packs, fix packs and/or interim fixes required for the optional components. To determine whether such additional software updates are needed, contact your service representative.

## *Upgrade Planning and Procedures*

This upgrade section contains the following major topics:

- [“Plan the Upgrade” on page 459](#)
- [“Upgrade Core Components” on page 477](#)
- [“Upgrade Add-On Components” on page 558](#)

## Plan the Upgrade

This section includes the following topics:

- [“Upgrade Overview” on page 460.](#)
- [“Upgrade Planning Considerations” on page 461.](#)
- [“Upgrade Checklists” on page 470.](#)

## Upgrade Overview

The upgrades described in the guide assume that you will:

- Retain your basic platform configuration from the previous release (except as noted above for Content Engine).
- Make no changes to existing user or group definitions during the upgrade (although you must specify new administrative and program accounts, as noted in [“Specify IBM FileNet P8 Accounts” on page 60](#)).
- Apply only the necessary supported software updates, service packs, fix packs and patches, as noted in this guide and in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

Before upgrading to IBM FileNet P8 Platform 4.0.x you must verify that the following conditions have been met:

- An existing IBM FileNet P8 3.5.x system is installed and configured that meets the minimum requirements. You cannot upgrade directly to IBM FileNet P8 Platform 4.0.0 from any version prior to the levels shown in table in [“General Requirements for all IBM FileNet P8 Systems” on page 461](#).
- **WARNING** Do NOT install this Application Engine 4.0.1 Service Pack until all of your IBM FileNet P8 functional expansions support it. For information, see the latest Service Pack edition of the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

**CAUTION** You must make these operating system changes at the time of 4.0.x upgrade and before you make the system available to general users. To migrate Content Engine and Content Search Engine to UNIX at a later date, you must contact your IBM FileNet representative and arrange a possible IBM FileNet Professional Services engagement. Likewise, to migrate other components or third-party software from one platform to another (for example, migrating from an SQL Server database engine to Oracle), you must contact your IBM FileNet representative about a IBM FileNet Professional Services engagement.

**NOTE** This guide generally does *not* address migrations to different operating systems, databases, or directory services. However, because IBM FileNet P8 Platform 4.0.0 introduces UNIX support for Content Engine server and the new Autonomy K2-based Content Search Engine software, you can install these components on UNIX rather than Windows as part of the upgrade procedure. You can then upgrade your 3.5.x metadata, file stores, and index stores, and thereafter physically move the resulting file storage areas and index areas from their Windows locations to UNIX locations.

## Upgrade Planning Considerations

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

### General Requirements for all IBM FileNet P8 Systems

- **Gather auxiliary documentation.** Retrieve the following from the [IBM Information Management support page on www.ibm.com](#), or from your IBM FileNet P8 Help installation:

**NOTE** For general instructions on how to navigate to this and other IBM FileNet product documentation on the IBM web site, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

- *IBM FileNet P8 Hardware and Software Requirements.* This document provides details for all IBM FileNet P8 system components, as well as the minimum supported levels of third-party software components. The information throughout the *IBM FileNet P8 Platform Installation and Upgrade Guide* assumes you have met all applicable requirements listed in that document.
- The IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > FileNet P8 Security > Users and groups](#). This help topic provides a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain a IBM FileNet P8 system.
- *IBM FileNet P8 Release Notes.* This document provides details on new features, known issues, and resolved problems. Please review the “What's New” topic for recent product changes, particularly to Security features.
- *IBM FileNet P8 Platform Installing Non-English Environments Technical Notice.* This document will help you set up the product if your environment is not English-language based.
- *IBM FileNet P8 Platform High Availability Technical Notice.* This document provides details on how to set up your IBM FileNet P8 system using clusters, farms, and other high-availability software and hardware.
- The IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup](#). This help topic describes how to shut down and restart IBM FileNet P8 Platform components and some expansion products. Manual, command line, and some sample batch file procedures are provided.
- *IBM FileNet P8 Platform Planning and Deployment Guide.* This document provides information on deploying a IBM FileNet P8 system from a staging environment into a full production environment.
- *IBM FileNet P8 Platform Troubleshooting Guide.* This document provides troubleshooting information on all aspects of the product.
- *IBM FileNet P8 Platform 3.5.x Version Tools Technical Notice.* This document provides instructions on running component-specific tools to determine whether you are at the IBM FileNet P8 software release, fix pack, and interim fix levels required to upgrade.

- *IBM FileNet P8 Platform 4.0.x Version Tools Technical Notice*. This document provides instructions on running component-specific tools to determine whether you are at the IBM FileNet P8 4.0.x software release, fix pack, and interim fix levels required to upgrade in your environment (e.g., if you running expansion products that require a particular level of IBM FileNet P8 Platform 4.0.x software).
- *IBM FileNet P8 Platform Performance Tuning Guide*. This document provides performance tuning information on all aspects of the product.
- *IBM FileNet Rendition Engine Installation and Upgrade Guide*. This document provides details and information on installing and upgrading Rendition Engine components.
- **Apply the required minimum level of IBM FileNet P8 Service Packs, Fix Packs, or Interim Fixes to the currently installed software before you upgrade.** IBM FileNet Service Packs, Fix Packs and Test Fixes often include feature updates that are required to ensure a successful upgrade. Therefore, before you begin your upgrade to IBM FileNet P8 Platform 4.0.0, you must have applied the minimum level of Service Pack, Fix Pack, or Test Fix to your installed 3.5.x components, as noted in the table below.

**NOTE** Over the life of IBM FileNet P8 3.5.x, Service Pack, Fix Pack, and Interim Fix requirements could change. To determine which additional Service Packs are available, go to the [IBM Information Management support page on www.ibm.com](http://www.ibm.com) .

Component	Fix Pack	Software Build
Content Engine	CE-3.5.2-002	kl195.025
Process Engine	PE-3.5.2-002	pe185.013
Application Engine	AE-3.5.1-003	per185.027.a01
IBM FileNet P8 eForms	P8eF-3.5.1	raptor220.004
Process Analyzer	PA-3.5.1-002	pa185.012
Process Simulator	PS-3.5.2-002	ps185.008

For more details, see the “Fix Pack and Patch Dependencies” topic in the *IBM FileNet P8 Platform Release Notes*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

- **IBM recommends that you plan and test the upgrade on a designated test system first.** Verify the upgrade is successful by running functionality and stress tests. After successful verification, perform the production upgrade.

## Expansion Product Considerations

### eForms

If you have the IBM FileNet P8 eForms expansion product installed you must uninstall it before upgrading Application Engine.

## Records Manager

- **You must update the RM Datamodel on FPOS object stores before you attempt to upgrade Content Engine Data on these.** The Upgrader Tool will fail if run on an FPOS object store unless it has been updated to the 4.0.0 datamodel. This limitation only applies to FPOS object stores, you can update ROS object stores without first updating the datamodel. For more information, see [“Upgrade Content Engine Data” on page 498](#) and the *IBM FileNet Records Manager Installation and Upgrade Guide* task “Prepare to Upgrade Object Stores.”

### NOTES

- To complete your platform upgrade and confirm the upgrade you must update at least one object store. If your setup contains only RM object stores (FPOS and ROS) you must either update at least one of these or create a non-RM object store prior to upgrading and use this to confirm the upgrade.
  - If your site preferences reside on an RM object store you must update this object store to complete the platform upgrade. Alternatively, you can move the site preferences to another object store before the upgrade.
- **Do not install Enterprise Manager 4.0.x on any machine running the 3.5.x version until the RM 4.0.0 upgrade is complete.** This includes updating all RM object stores to version 4.0.0. Installing EM 4.0.x will cause unexpected behavior with the RM Data Upgrade Tool.

## Third-party Software Considerations

The 3.5.x and 4.0.0 versions of IBM FileNet P8 support common levels of third-party software. Exceptions are listed in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### Fixed Content Devices

- If you set up EMC Centera fixed content devices in version 3.5.x of IBM FileNet P8, make sure they are using the version of the CentraStar™ operating system supported in version 4.0.0 of IBM FileNet P8, as specified in the *IBM FileNet Hardware and Software Requirements*.
- If you set up NetApp SnapLock fixed content devices in version 3.5.x of IBM FileNet P8, make sure they are using the version of the Data ONTAP™ operating system supported in version 4.0.0 of IBM FileNet P8, as specified in the *IBM FileNet Hardware and Software Requirements*.

**NOTE** To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

## Operating System Considerations

### General

- **Upgrade to Windows 2003 before you begin the upgrade process to IBM FileNet P8 4.0.0.** See the *IBM FileNet P8 Hardware and Software Requirements* for details on any

required Windows 2003 Service Packs and patches. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

**NOTE** IBM FileNet field personnel indicate that most customers prefer to do a fresh installation of Windows 2003 rather than upgrading the Windows 2000 software. This will require you to install Content Engine 4.0.0 on its own Windows 2003 server (separate from the 3.5.x Content Engine server).

- **Synchronize the time and date on all servers.** System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

The Process Engine database server (the machine that hosts the database used by Process Engine) is considered the master time keeper; the UTC time of that machine is considered the correct time. The server hosting the Process Engine API and the server hosting Content Engine must have the UTC time set to match the UTC time on the Process Engine database server, plus or minus 15 minutes.

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

## Content Engine

### Windows 2003

- **Ensure proper upgrade to Windows 2003 on your existing Content Engine 3.5.2 servers prior to IBM FileNet P8 upgrade.** If you do want to upgrade to Windows 2003 on your Content Engine 3.5.2 servers and intend to run them for any length of time before you upgrade to Content Engine 4.0.0, be sure to see the “Configure Content Engine and SQL Servers for Windows 2003” in the *IBM FileNet P8 Platform Installation and Upgrade Guide v 3.5.x* for important details and procedures. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- **Ensure that you have the necessary Windows clients for Enterprise Manager 4.0.0.** Although Content Engine server now runs on UNIX as well as Windows, the Enterprise Manager administrative client still runs on Windows only. For upgrade purposes, IBM recommends you install Enterprise Manager 4.0.0 on a separate client from the 3.5.x version.

### UNIX

- **If you intend to move to UNIX for Content Engine 4.0.0, consider the following:**
  - You *must* install Content Engine server 4.0.0 on UNIX as an initial step in the upgrade process. To migrate after you have completed the upgrade and begun using your system, you



will have to contact your IBM FileNet representative and arrange a Professional Services engagement.

- You must temporarily leave your existing 3.5.x file stores (and associated index stores) on Windows to upgrade them using the CE 3.5.2 to 4.0 Upgrader Tool. This means that if you install Content Engine server on UNIX, you must have an NFS gateway in place (for example, Windows R2 Gateway or Samba) to enable communication between the new UNIX Content Engine server and Windows file storage areas and index areas.

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

## Process Engine

- **Reconcile the Process Engine user security information.**

The Process Engine duplicates certain parts of the user security information in its own database. Over time, the information in the directory service might be changed or updated. When this happens, the information in the Process Engine's environment records, whether cached or permanent, can end up containing old, invalid information about the Process Engine users and groups.

**WARNING** During the PE upgrade it is critical that this user information is correct and up to date. Before upgrading you must reconcile the cached and permanent user data environment records on the Process Engine with the possibly more-up-to-date data in an LDAP-based directory service.

For more information, see the IBM FileNet P8 help topic [FileNet P8 Administration > Process Engine Administration > Administrative tools > vwtool > Commands > environment](#).

- If you are using the Process Analyzer expansion product, several steps must be taken on the Process Analyzer, on the Process Engine database, and on Process Engine, before upgrading Process Engine. See ["Upgrade Process Engine \(UNIX\)" on page 517](#) and ["Upgrade Process Engine \(Windows\)" on page 527](#) for details.

## UNIX

- **Ensure minimum /tmp size.** The **/tmp** directory must have 510 MB free for use by Process Engine Setup.
- **Allocate a minimum of 500 MB of additional disk space to the /fnsw disk volume.** This space is required for upgrades and is in addition to the minimum space requirements called for in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see ["Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs" on page 21](#).
- For the SEC databases, expand the following raw partitions to the indicated sizes:
  - fn\_sec\_r10 - 64MB
  - fn\_sec\_db0 - 64MB

## Security Considerations

- **If you intend to migrate to a different directory service, do so prior to the 4.0.0 upgrade.** If you are expecting to migrate to a different directory service for use with IBM FileNet P8 Platform 4.0.0, you must do so *before* you begin the upgrade from 3.5.x. See the *IBM FileNet P8 Platform 3.5.x Directory Service Migration Guide* for complete information. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#). Migrating to a different directory service provider is not supported once you are running the 4.0.0 release but is likely to be provided in subsequent releases or Service Packs. Contact your IBM FileNet representative if you intend to switch to IBM Tivoli Directory Server, which is supported for IBM FileNet P8 4.0.0, but not for 3.5.x.
- **Be sure to understand updated user and group account requirements for release 4.0.0.** The 4.0.0 release requires you to designate or create several new accounts (for example, the K2 Operating System User for Content Search Engine). For details, see [“Specify IBM FileNet P8 Accounts” on page 60](#).
- **Determine whether you intend to use Secure Socket Layers (SSL).** IBM recommends you use SSL to strengthen security. Configuring SSL for 4.0.0 is noticeably different than for 3.5.x. For example, Content Engine 4.0.0 now relies on its own application server to direct calls to your chosen authentication provider. For configuration details, see [“Set Up Content Engine and Client Transport SSL Security” on page 412](#) and [“Set Up Application Engine SSL Security” on page 416](#).
- **Gather 3.5.x authentication information and plan 4.0.x authentication.**

In the 3.5.x release, Content Engine authentication and authorization information are both specified by configuring one or more authentication providers, using Active Directory, Novell eDirectory, or Sun Java System Directory Server.

- For Novell and Sun authentication providers, the root of the configured 3.5.x Directory Service contains one or many children, or naming contexts, each of which are used automatically by 3.5.x Content Engine as 3.5.x FileNet P8 authentication realms.
- For Windows Active Directory, the 3.5.x Content Engine uses the Windows domain (also known as *deployment domain*) it resides in and automatically gets all trusted domains, siblings, parents, children within the domain's forest. By default, the deployment domain is the "Default Realm."

3.5.x authentication realms are viewable (but not editable) in Enterprise Manager's root domain property sheet > Authentication Provider tab > DefaultRealm property. The names of all realms for the current FileNet P8 domain appear as values in the drop-down list for this property.

**NOTE** If you are using the 3.5.x Content Engine attributes DefaultRealm and RestrictToDefaultRealm to modify 3.5.x authentication behavior, then you must carefully map out which 3.5.x authentication realms are active and essential to your upgraded 4.0.x Content Engine authentication scheme. (The DefaultRealm and RestrictToDefaultRealm attributes are not supported by Content Engine 4.0.x, because of the introduction of the application server which uses Java Authentication and Authorization Service (JAAS) for authentication.)

By contrast, in Content Engine 4.0.x, configuring authentication and authorization are two separate steps, both of which must be completed on the Content Engine 4.0.x server

environment before upgrading object stores from 3.5.x. Authentication is configured through the application server's admin console. Authorization is configured by creating one or more Directory Configuration objects in the Content Engine 4.x for the FileNet P8 domain. (Note that 3.5.x authentication information is not carried forward by the CE 3.5.2 to 4.0 Upgrader Tool.) Therefore, the following are prerequisites to upgrading a 3.5.x domain:

- Install and configure one of the supported J2EE application servers, using the application server's install tools.
- Install a Content Engine 4.0.x server as an application into the J2EE application server, using the Content Engine Setup program.
- Configure the authentication environment, through the J2EE application server's console. (If you want the Content Engine Setup program to create the initial application server's authentication, as opposed to creating it yourself before installing the Content Engine, make sure you select to install the component Application Server Authentication Provider on the associated Content Engine Setup wizard screen.)
- Create a Content Engine 4.0.x domain, using the Enterprise Manager's Add Domain Configuration wizard.
- Create one or more Directory Configurations, using the Enterprise Manager's Create a Directory Configuration wizard. One Directory Configuration is required for each distinct 3.5.x authentication realm.

There are several cases that might exist for this last bullet. In all cases, Directory Configurations must be created that preserve access to all users that have used the object stores being upgraded. Access to all users' SIDs and group information is necessary for authorization to occur.

- In the typical case, you will create one 4.0.x Directory Configuration object for each realm used in 3.5.x, specifying the same LDAP host and port information that was used in Content Engine 3.5.x.
- However, in some cases customers will choose to configure the Content Engine 4.0.x server to obtain directory configuration information from a different LDAP server than was used in Content Engine 3.5.x (for instance by using a replica of the 3.5.x directory service which contains the same user or group information), or to expand the users of the Content Engine 4.0.x domain being upgraded by adding additional realms. If this is the case, then the set of 4.0.x Directory Configurations that must be created for the new 4.0.x domain could differ from the set of realms that were configured in 3.5.x. Just make sure that the 4.0.x Directory Configurations contain the same user and group principals, with the exact same SIDs, as those used in Content Engine 3.5.x. Remember that users in the newly added realms will have no authorization to access any Content Engine 4.0.x objects until they are explicitly granted authorization, typically by using Enterprise Manager in ways described in the 4.0.x *Help for Content Engine Administration*.

## Network Considerations

- **Ensure availability of the required port numbers.** For a composite list of port numbers required for IBM FileNet P8 4.0.0, see the ["IBM FileNet P8 Port Numbers" on page 643](#).

## Database Considerations

### General

- **Update to the appropriate database patches before you upgrade IBM FileNet P8 components.** For minimum patch requirements, see *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs”](#) on page 21.
- **Update database versions after you upgrade IBM FileNet P8 components.** If you will upgrade the version of the database from DB2 version 8 to 9, or from SQL Server 2000 to 2005, complete the upgrade of all IBM FileNet P8 components and verify a fully functional system before upgrading the database software.
- **Plan for upgrade implications regarding databases.** If you are upgrading to version 4.0.x, you will continue to use existing object store databases. However, you must create a new database for the 4.0.x global configuration data (GCD), which the Content Engine Upgrader Tool migrates from its 3.5.x file-based format (**sysinit.dat**).
- **Decide when to upgrade the database.** IBM FileNet P8 requires no re-configuration of the application server when the database server is undergoing an upgrade as long as connection information in the data sources does not change and no data is lost when performing the database upgrade procedures. Be sure to update to the JDBC driver associated with the new database version as specified in *IBM FileNet P8 Hardware and Software Requirements*.
- **Consider the following for Business Process Manager components (PE, PA, PS):**
  - The Process Engine database upgrade automatically renames the existing VWLogxxx\_yyy tables to VWLogxxx\_yyy\_Archive and creates the new VWLogxxx\_yyy table.
  - The Process Engine Log SeqNumber starts from 21000.
  - Process Analyzer gets its data from the Process Engine database. The 3.5.x and 4.0.0 database schemas have changed both on Process Engine and Process Analyzer. The 3.5.x data from Process Engine must be transmitted to Process Analyzer 3.5.x *before* you upgrade either of these components to 4.0.0. For details, see the 4.0.x *IBM FileNet Process Analyzer Installation and Upgrade Guide*.

### Oracle

- **Ensure you have applied appropriate Oracle patches to Oracle clients as well as servers.** Be sure that clients remote from Oracle database engines have patches that are comparable to the database server. Oracle clients include any machines remote from the Oracle database server where Content Engine, Process Engine, and Enterprise Manager software is installed. You can download all the required Oracle database server patches from [OracleMetaLink](#) and install them.
- **Determine when to execute SQL scripts.** A number of SQL scripts must be executed. These scripts can be executed manually, before starting Process Engine Setup, or from Process Engine Setup. See [“Process Engine SQL Scripts”](#) on page 647 for information on execution modes and associated security requirements as well as details about the scripts.

## Microsoft SQL Server

- **Determine when to execute SQL scripts.** A number of SQL scripts must be executed. These scripts can be executed manually, before starting Process Engine Setup, or from Process Engine Setup. See [“Process Engine SQL Scripts” on page 647](#) for information on execution modes and associated security requirements as well as details about the scripts.

## *Application Server Considerations*

- **Determine whether the number of object stores to be upgraded will require you to run a 64-bit Java Virtual Machine (JVM) on your Content Engine application servers.** If the 3.5.x FileNet P8 domain you are upgrading contains more than 50 object stores, IBM recommends you install Content Engine on application servers running a 64-bit, rather than 32-bit, JVM. Otherwise, you might experience significant Content Engine performance issues.

## Upgrade Checklists

The following topics provide checklists of items that must be completed for a successful upgrade of the P8 Platform. The lists are divided by user or role. Note that your organization might have different administrator roles, and that some of the responsibilities of listed roles might vary.

You can print these lists and gather and record the relevant information about your existing installation before you begin the upgrade. You can then record new information as you perform the upgrade tasks. Where applicable, the list items indicate which role will need the information you provide. IBM recommends coordinating communication among the various roles to facilitate an easier installation process.

### IT Administrator

This role administers hardware and operating systems for your environment, as well as helping make decisions about machine usage and configurations such as clustering and farming. The information confirmed and collected in this set of checklists should be provided to the Upgrade Administrator for each component.

#### Before the Upgrade

- Make sure the site has restorable backups of system and data for all IBM FileNet P8 components. In case of an emergency, you might need to back out of the upgrade. For details on what to back up for IBM FileNet P8 components, see the IBM FileNet P8 help topic [Enterprise-wide Administration > Backup and Restore](#).
- In some cases it might be necessary for IBM FileNet Engineering and Support personnel to access the system for support. To facilitate such access, we request that you implement a reliable remote access method.
- If you have configured clusters or farms, see *IBM FileNet P8 Platform High Availability Technical Notice*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- If you have more than 50 object stores in a single 3.5.x FileNet P8 domain, see the Notes section in [“Upgrade Content Engine Software” on page 487](#) for a recommended upgrade approach.

#### Content Engine

- The upgrade process for Content Engine includes a new installation of the 4.0.0 version. Refer to the items for Content Engine in the [“Installation Checklists” on page 42](#).
- Record the following information about the K2 Master Administration Server.

**K2 Master Administration Server host name:**

\_\_\_\_\_

**K2 Master Administration Server port:**

\_\_\_\_\_

**StyleSetAlias path (FileNet\_FileSystem\_PushAPI):**

\_\_\_\_\_

**K2 Server names:**

\_\_\_\_\_

**K2 Index server names:**

\_\_\_\_\_

**K2 Broker server names:**

\_\_\_\_\_

- For each object stores you plan to upgrade, record the following:

**Local JNDI name:** \_\_\_\_\_  
**Global JNDI name:** \_\_\_\_\_

- Specify the Content Engine Operating System user:

\_\_\_\_\_

## Process Engine

- (HP-UX only) Collect kernel parameter information before you upgrade the Process Engine software.

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

[http://www.hp.com/products1/unix/java/infolibrary/prog\\_guide/java1/configuration.htm](http://www.hp.com/products1/unix/java/infolibrary/prog_guide/java1/configuration.htm)

- (All UNIX platforms) Expand the raw partitions for the SEC databases:

**fn\_sec\_r10** \_\_\_\_\_  
**fn\_sec\_db0** \_\_\_\_\_

- (All UNIX platforms) Save the following files for the fns (or alias) and root users:

**.Xdefaults** \_\_\_\_\_  
**.Xresources** \_\_\_\_\_  
**.dbxinit** \_\_\_\_\_  
**.dtprofile** \_\_\_\_\_  
**.env** \_\_\_\_\_  
**.login** \_\_\_\_\_  
**.mwmrc** \_\_\_\_\_  
**.xinitrc** \_\_\_\_\_  
**.profile** \_\_\_\_\_  
**.cshrc** \_\_\_\_\_

## Application Engine

- Provide the following administrative credentials for installation on the Application Engine server:

**Windows (local administrator):** \_\_\_\_\_  
**UNIX (root)** \_\_\_\_\_

**NOTE** The root user must have read/write/execute access to both the **/bin** directory and the Content Engine installation directory.

- (Optional) Provide a local machine administrative account to run Application Engine. This user must have read and write privileges:

**Windows:** \_\_\_\_\_  
**UNIX:** \_\_\_\_\_

- (WebLogic) Determine the recommended MaxPermSize value for MEM\_ARGS:

**MaxPermSize:** \_\_\_\_\_

## Security Administrator

This role administers authentication, users and groups, passwords, encryption, and general network access considerations for the upgrade and eventual use of the IBM FileNet P8 Platform software. The information confirmed and collected in the following checklist items should be provided to the Upgrade Administrator for each component, as well as to the FileNet P8 administrators (for example, the GCD administrator).

### General

- Note the user name and password for a FileNet P8 GCD administrator:

**GCD Administrator:** \_\_\_\_\_  
**Password:** \_\_\_\_\_

### Content Engine

- The upgrade process for Content Engine includes a new installation of the 4.0.0 version. Refer to the items for Content Engine in the [“Installation Checklists” on page 42](#).
- Ensure that you have the necessary local and Content Engine administrative user accounts and passwords to run upgrade-related programs. You will need these various levels of administrative operating system access to run Content Engine Setup and the CE 3.5.2 to 4.0 Upgrader Tool against the Content Engine servers and Windows Enterprise Manager administrative clients. For details on accounts and privileges, see [“Upgrade Content Engine Software” on page 487](#) and [“Specify IBM FileNet P8 Accounts” on page 60](#).
- Supply the following K2 user and group information for the Content Search Engine:

**K2 Security User Group:** \_\_\_\_\_  
**K2 Security User:** \_\_\_\_\_  
**K2 Security User password:** \_\_\_\_\_  
**K2 Operating System User:** \_\_\_\_\_  
**K2 Operating System User password:** \_\_\_\_\_  
**(Windows) Domain name on which K2 services run:** \_\_\_\_\_

- (Windows) Verify directory permissions on file storage areas for the following users:

**User performing the upgrade:** \_\_\_\_\_  
**User running the application server:** \_\_\_\_\_



## Application Engine

- Specify the users and groups to add to the Application Engine Administrators access role:

Users or groups

- Specify the users and groups who will be allowed to create subscriptions to add to the PWDesigner access role.

**NOTE** These are not specifically required to complete an install, but required to create subscriptions later in Workplace, and can also be added later:

Users or groups

- For SSO, specify the following:

SSO proxy host URL: \_\_\_\_\_  
 SSO proxy host server name: \_\_\_\_\_  
 HTTP port on the SSO proxy host: \_\_\_\_\_  
 HTTPS port on the SSO proxy host: \_\_\_\_\_

- For SSL, specify the following:

SSL Host name and port number: \_\_\_\_\_  
 Java Server HTTP port: \_\_\_\_\_

## Database Administrator

This role oversees database creation and administration. The decisions, task confirmations, and information gathered in the following checklist items should be provided to the Upgrade Administrator for each component.

## General

- Note the following general database information:

**Database instance name:** \_\_\_\_\_  
**NLS character set:** \_\_\_\_\_

## Content Engine

- The upgrade process for Content Engine includes a new installation of the 4.0.0 version. Refer to the items for Content Engine in the [“Installation Checklists” on page 42](#).
- Note the following database information for each existing object store:

**Database user name:** \_\_\_\_\_  
**Database user password:** \_\_\_\_\_  
**Database type:** \_\_\_\_\_  
**Database class:** \_\_\_\_\_  
**Database URL:** \_\_\_\_\_

## Process Engine

- Determine when to execute Oracle SQL scripts.

A number of SQL scripts must be executed. These scripts can be executed manually, before starting Process Engine Setup, or executed from Process Engine Setup. See [“Process Engine SQL Scripts” on page 647](#) for information on execution modes and associated security requirements as well as details about the scripts.

## Upgrade Administrator

This role will perform upgrade tasks for your IBM FileNet P8 Platform software. This role might also perform initial configuration, setup, and startup tasks that would subsequently fall under the role of the FileNet P8 Administrator.

## Before the Upgrade

- Be sure to read the “What’s New” and “Known Issues” topics in the *IBM FileNet P8 Platform Release Notes*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#). New features might be interesting for your implementation of IBM FileNet P8 Platform, and known issues might impact the upgrade procedures.
- Before you begin the upgrade, make sure no users or processes are accessing the IBM FileNet P8 system.
- Before you begin the upgrade, disable tasks such as:
  - Scheduled automated backups
  - Cron jobs
  - Virus scanning

## Environment Considerations

- ❑ Determine the existing IBM FileNet P8 Platform documentation URL. If you have an existing application server in place with a IBM FileNet P8 Documentation web site established, you can determine its URL by checking the Documentation server site preference in Workplace.

**Documentation Server URL:** \_\_\_\_\_

**NOTE** You can now install and deploy the IBM FileNet P8 Documentation on the Content Engine application server, as well as on the Application Engine or any other supported application server.

## Content Engine

- ❑ Review the list of users and permissions required to perform the Content Engine Upgrade. See the overview of [“Upgrade Content Engine Software” on page 487](#).
- ❑ The upgrade process for Content Engine includes a new installation of the 4.0.0 version. Refer to the items for Content Engine in the [“Installation Checklists” on page 42](#).
- ❑ Note the GCD (sysinit file) location on the master GCD server. The GCD master is typically your initially installed Content Engine server.
- ❑ If you are going to upgrade content-search indexes (Verity collections), provide the UNC path to the collections directory for each collection you plan to upgrade:

**UNC path to collections directory:** \_\_\_\_\_

- ❑ After the install of the 4.0.0 Content Engine software is complete, provide the URL for the component manager on Content Engine:

**CE Component Manager URL:** \_\_\_\_\_

- ❑ After the install of the 4.0.0 Content Engine software is complete, provide the download and upload URLs for Content Engine (server name and port number for your Content Engine server for downloading and uploading document content):

**CE download URL:** \_\_\_\_\_

**CE upload URL:** \_\_\_\_\_

- ❑ After the install of the 4.0.0 Content Engine software is complete, provide the Content Engine client software URL (server name and port number for your Content Engine server on which the Application Engine Web Connectivity (Java) module resides):

**CE client software URL:** \_\_\_\_\_

## Process Engine

- ❑ Note the User and Group Base specified in the LDAP Connection tab in the Process Engine Task Manager. This information can be useful for troubleshooting during the upgrade process.

**User Base:** \_\_\_\_\_

**Group Base:** \_\_\_\_\_

## Application Engine

- Gather the J2EE application server Admin user name and password or have the administrator available during the upgrade.

**User name:** \_\_\_\_\_

**Password:** \_\_\_\_\_

- Determine which components and installation option you used to install Application Engine 3.5.x.

When you upgrade Application Engine, you must install the same components as are currently installed. If you installed Application Engine 3.5.x using a Typical install, perform a Typical install; if you used a Custom install, perform a Custom install using the same selection of components.

In particular, do NOT run a Typical upgrade if you initially ran a Custom installation (for example, to install the Java API). If you do this, your Custom software component choices will not be updated as expected.

To find out which Application Engine 3.5.x components are currently installed, look for the following folders:

- Process and Content Java APIs (only)
  - <AE\_install\_path>/FileNet/lib2**
  - <AE\_install\_path>/FileNet/Router**
- Workplace web applications and Process and Content Java APIs
  - <AE\_install\_path>/FileNet/lib2**
  - <AE\_install\_path>/FileNet/Router**
  - <AE\_install\_path>/FileNet/Workplace**
- Workplace Toolkit Source Code (for custom application development)
  - <AE\_install\_path>/FileNet/Workplace/source**

- Record the current Component Manager configuration data.

Existing Component Manager configuration data is not retained during the Application Engine upgrade. If you did not modify the default Application Engine 3.5.x configuration for the Component Manager, it is not necessary to record and re-enter the Component Manager configuration data; it will be set to the Application Engine 4.0.x default values during the upgrade.

To view the current configuration data, start the Process Task Manager on Application Engine. Click **Component Manager** and note the values displayed on the **General** and **Required Libraries** tabs.

- Verify that at least one object store exists and is available for upgrading. See [“Create Object Stores” on page 258](#).

# Upgrade Core Components

## To upgrade the core IBM FileNet P8 Platform components

---

**CAUTION** You must upgrade the core components as described in the topics listed below before upgrading the optional add-on components as described in [“Upgrade Add-On Components” on page 558](#).

1. Review IBM FileNet P8 Platform requirements and other planning considerations. See [“General Requirements for all IBM FileNet P8 Systems” on page 461](#).
2. Upgrade the IBM FileNet P8 Platform documentation. Do [Task 1 on page 478](#).
3. Ensure that you are running the minimum required release of IBM FileNet P8 Platform software and service packs for doing an upgrade. Do [Task 2 on page 486](#).
4. Upgrade Content Engine. Do [Task 3 on page 487](#).
5. Install Content Engine software updates. Do [Task 4 on page 497](#).
6. Upgrade Content Engine data. Do [Task 5 on page 498](#).
7. Install Content Search Engine software updates. Do [Task 6 on page 512](#).
8. Complete the post-upgrade Content Engine configuration. Do [Task 7 on page 513](#).
9. Upgrade Process Engine. Depending on your server platform, do one of the following:
  - [Task 8a on page 517](#) (UNIX).
  - [Task 8b on page 527](#) (Windows 2000).
10. Install Process Engine software updates. Do [Task 9 on page 537](#).
11. Install the latest Content Engine client files on Process Engine servers. Do [Task 10 on page 538](#).
12. Complete the post-upgrade Process Engine configuration. Do [Task 11 on page 539](#).
13. Upgrade Application Engine. Do [Task 12 on page 551](#).
14. Complete the post upgrade Application Engine configuration. Do [Task 13 on page 556](#).

# Task 1: Upgrade IBM FileNet P8 Documentation

## Overview

There are two methods available for updating to an IBM FileNet P8 4.0.x Documentation installation:

- [“Refresh IBM FileNet P8 4.0.0 Documentation Without Uninstalling” on page 479](#)  
Use this option if you have IBM FileNet P8 3.5.x documentation installed without any expansion product help.
- [“Update IBM FileNet P8 4.0.0 Documentation by Uninstalling and Reinstalling” on page 481](#)  
Use this option if you have IBM FileNet P8 3.5.x documentation installed with expansion product help included (for example: Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets).

## NOTES

- With the IBM FileNet P8 4.0.0 release, the documentation can be located on the Content Engine server. If you wish to install the 4.0.0 documentation on the Content Engine server, or install on a new application server, you will not need to perform an upgrade. Instead, use the install instructions in the following topics, depending on your application server type:
  - [“Install IBM FileNet P8 Platform Documentation \(WebSphere\)” on page 130](#)
  - [“Install IBM FileNet P8 Platform Documentation \(WebLogic\)” on page 135](#)
  - [“Install IBM FileNet P8 Platform Documentation \(JBoss\)” on page 141](#)
- The refresh procedure assumes the following:
  - You wish to replace the 3.5.x documentation at the existing location with updated files.
  - You are familiar with your application server’s procedures for reinstalling or redeploying web site applications.
- Before installing any IBM FileNet P8 documentation, review the *IBM FileNet P8 Hardware and Software Requirements* for the required software versions, Service Packs, and Hot Fix Packs for third-party software. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).
- Although some versions of BEA WebLogic support deployment of WAR-file-based web applications, you cannot deploy **ecm\_help.war** on this application server platform. You must instead use the fully expanded IBM FileNet P8 documentation directory (**ecm\_help**) structure. Otherwise, the searches within the IBM FileNet P8 documentation will not work and you will receive null pointer errors.
- In environments where Windows NTFS is used, there is a 256 character limit on file and folder names (folder depth). When trying to delete an existing IBM FileNet P8 documentation web site, you may encounter access denied errors. See Microsoft Knowledge Base article <http://support.microsoft.com/?kbid=320081> for more information.

- If your IBM FileNet P8 documentation server is a Windows server, verify that Microsoft Windows Internet Information Services (IIS Admin Service, Simple Mail Transport Protocol (SMTP), and World Wide Web Publishing Service) are stopped and set to *manual*.
- Before installing the latest IBM FileNet P8 documentation, be sure to back up your existing IBM FileNet P8 documentation web site according to your site and application server procedures. This precaution will allow you to restore the IBM FileNet P8 documentation web site quickly if for any reason you have to back out of or delay your IBM FileNet P8 software upgrade.
- If you install any IBM FileNet P8 expansion product(s) as part of your upgrade (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets), be aware that:
  - You must copy the associated help for all expansion products onto the IBM FileNet P8 documentation server.
  - You must update the index for the help Search feature as documented below. This action ensures that searches return all expected results. If you add help for other expansion products later, you must re-run the procedure below for updating the help Search index.

**IMPORTANT** You must also update the help Search index if you update the installed IBM FileNet-based help that you have customized or translated as part of your own application development. For further help-customization details, see the *IBM FileNet P8 4.0.0 Localizing and Customizing FileNet P8 Help Technical Notice*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

- Any time you update the documentation search index, a backup of the files in the existing **Index\core** directory will be automatically copied to the **Index\IndexOld** subdirectory. You can reapply these backed-up files to the **core** subdirectory (after first removing the new files created there) if you need to return to your previous indexed state.
- Depending on your operating system (Windows or UNIX) and application server version (WebSphere or WebLogic), some screens may be slightly different than those documented in the procedures listed below.

## Refresh IBM FileNet P8 4.0.0 Documentation Without Uninstalling

You can refresh an existing IBM FileNet 3.5.x (or earlier 4.0.0) P8 documentation installation by simply copying the newer documentation files over the existing files, and then reindexing for Search.

Use this option if you have IBM FileNet P8 3.5.x documentation installed without any expansion product help.

### NOTES

- If you have FileNet P8 3.5.x documentation installed with expansion product help included, use the procedure [“Update IBM FileNet P8 4.0.0 Documentation by Uninstalling and Reinstalling” on page 481](#).
- This refresh procedure requires that you copy the expanded **ecm\_help** directory from the IBM FileNet P8 Documentation package over the existing documentation. This stipulation applies

to all application servers, even those that required a WAR file for initial deployment of the existing 3.5.x documentation, namely WebSphere.

### To refresh the IBM FileNet P8 documentation without first uninstalling

---

1. Stop the application server (or IBM FileNet P8 documentation site) on which the existing documentation is deployed so that no processes can access the documentation.
2. This step removes the Search index, which you must recreate later for the newer set of documentation files. On the P8 documentation application server, locate the deployed FileNet P8 documentation directory, and back up (or move to a safe location) the files located in:

UNIX: `<deployment_path>/ecm_help/search/index/core`

Windows: `<deployment_path>\ecm_help\search\index\core`

3. Access the IBM FileNet P8 Platform Documentation package.
4. Copy the expanded **ecm\_help** directory from the package over the deployed **ecm\_help** directory.
  - a. (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 documentation server.

```
cp -r <mount_location> <target_destination>
```

**WARNING** Care should be taken when copying folders in UNIX. Dragging-and-dropping of folders replaces any existing folder(s) of the same name. Also note that your switch (`-r`) requirements may be different from the example shown. Contact your system administrator if you have questions about proper syntax.

- b. (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
  - c. (Content Federation Services for Image Services) Copy the **cfs\_guide.pdf** file from the **Documentation** directory on the software package into the **ecm\_help/cfs\_help** directory on the IBM FileNet P8 documentation server.

**NOTE** Repeat this step for each of your expansion products. You can copy more than one expansion product documentation set to the documentation application server before continuing (for example, Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets) so you end up with one **ecm\_help** directory containing multiple sets of expansion product files added to it.

5. Continue as follows:
  - If you have added expansion product documentation, you will need to update the search index. Go on to the procedure in the following topic, [“Update Help Search Index” on page 484](#).
  - If you have no further documentation to install, and you did not install any expansion products, then go on to the procedure in the topic [“Deploy and Verify IBM FileNet P8 Documentation Web Site” on page 485](#).



## Update IBM FileNet P8 4.0.0 Documentation by Uninstalling and Reinstalling

Use this option if you have IBM FileNet P8 3.5.x documentation installed with expansion product help included (for example: Process Analyzer, Process Simulator, IBM FileNet P8 eForms, Content Federation Services for Image Services, or IBM FileNet P8 Portlets).

Use the following procedures to completely remove an existing IBM FileNet P8 documentation site from the application server before updating the site.

### To update the IBM FileNet P8 documentation on WebLogic application servers

---

#### To update a Weblogic 8.x installation

1. Verify the Weblogic application server is running, then start the WebLogic Administration Console.
2. From the WebLogic Administration Console, stop the existing documentation site.
3. Delete the existing IBM FileNet P8 documentation site.
4. Delete any IBM FileNet P8 documentation files, leaving the **ecm\_help** directory in place.
5. Access the IBM FileNet P8 Platform Documentation package.
6. Copy the **ecm\_help** folder structure from the package to the location of the original IBM FileNet P8 documentation as follows:
  - (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 documentation server.

```
cp -r <mount_location> <target_destination>
```

**CAUTION** Take care when copying folders in UNIX. Dragging-and-dropping folders replaces any existing folder(s) of the same name. Also note that your switch (`-r`) requirements can be different from the example shown. Contact your system administrator if you have questions about proper syntax.
  - (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
7. From the WebLogic Administration Console, click *<mydomain>* > Deployments > Web Application Modules.
8. Select the existing Web Application site and click the Deploy tab.
9. Click **Start**.
10. Continue as follows:
  - If you have added expansion product documentation, you will need to update the search index. Go on to the procedure in the following topic, [“Update Help Search Index” on page 484](#).
  - If you have no further documentation to install, then go on to the procedure in the topic [“Deploy and Verify IBM FileNet P8 Documentation Web Site” on page 485](#).

### To update a Weblogic 9.x installation

1. Verify the Weblogic 9.x application server is running, then start the WebLogic Administration Console.
2. From the WebLogic Administration Console, click **Deployments** and then select the IBM FileNet P8 documentation web site.
3. Click the **Lock and Edit** button.
4. Select the IBM FileNet P8 Documentation and stop the existing documentation site. If prompted, click **Yes** to stop the site.
5. From the install location, delete any IBM FileNet P8 documentation files, leaving the **ecm\_help** directory in place.
6. Access the IBM FileNet P8 Platform Documentation package.
7. Copy the **ecm\_help** folder structure from the package to the location of the original IBM FileNet P8 documentation as follows:
  - (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 documentation server.

```
cp -r <mount_location> <target_destination>
```

**CAUTION** Take care when copying folders in UNIX. Dragging-and-dropping folders replaces any existing folder(s) of the same name. Also note that your switch (`-r`) requirements can be different from the example shown. Contact your system administrator if you have questions about proper syntax.
  - (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
8. From the WebLogic Administration Console, click **Deployments**.
9. Select the existing Web Application site and Click the **Deploy** Tab.
10. Click **Start** and then click **Servicing all requests**.
11. Continue as follows:
  - If you have added expansion product documentation, you will need to update the search index. Go on to the procedure in the following topic, ["Update Help Search Index" on page 484](#).
  - If you have no further documentation to install, then go on to the procedure in the topic ["Deploy and Verify IBM FileNet P8 Documentation Web Site" on page 485](#)

### To update the IBM FileNet P8 documentation on WebSphere application servers

---

1. Verify the WebSphere 5.x application server is running, then start the WebSphere administrative console.
2. Expand Applications > Enterprise Applications.
3. Select the IBM FileNet P8 Documentation site.

4. Stop the existing documentation site.
5. From the initial install location, delete the IBM FileNet P8 documentation installed files, leaving the **ecm\_help** directory in place.
6. Access the IBM FileNet P8 Platform Documentation package.
7. Copy the **ecm\_help** folder structure from the package to the location of the original IBM FileNet P8 documentation as follows:
  - (UNIX) When copying expansion product documentation, use a `cp` copy command from a terminal to copy the **ecm\_help** directory structure from the associated *Documentation* package over the **ecm\_help** directory installed on the existing IBM FileNet P8 documentation server.  

```
cp -r <mount_location> <target_destination>
```

**CAUTION** Take care when copying folders in UNIX. Dragging-and-dropping folders replaces any existing folder(s) of the same name. Also note that your switch (`-r`) requirements can be different from the example shown. Contact your system administrator if you have questions about proper syntax.
  - (Windows) Use a copy command from a command prompt or drag-and-drop the files to the destination.
8. Expand Applications > Enterprise Applications.
9. Select the IBM FileNet P8 Documentation site.
10. Start the existing documentation site.
11. Continue as follows:
  - If you have added expansion products you will need to update the search index add on documentation for expansion products, go on to the procedure in the following topic, [“Update Help Search Index” on page 484](#).
  - If you have no further documentation to install, then go on to the procedure in the topic [“Deploy and Verify IBM FileNet P8 Documentation Web Site” on page 485](#).

#### **To update the IBM FileNet P8 documentation on JBoss application servers**

---

1. Shut down the application server.
2. Remove the entire existing IBM FileNet P8 Platform documentation application directory.
3. Remove the temporary working directory for the IBM FileNet P8 Platform documentation from the **<JBoss\_home>\work\MainEngine\localhost** directory.
4. Access the IBM FileNet P8 Platform Documentation package.
5. Copy the **ecm\_help** directory from the package to the existing location on the application server from which you just removed the old version of the directory.
6. Continue as follows:
  - If you have added expansion products you will need to update the search index. Go on to the procedure in the following topic, [“Update Help Search Index” on page 484](#).

- If you have no further documentation to install, then go on to the procedure in the topic [“Deploy and Verify IBM FileNet P8 Documentation Web Site” on page 485](#).

## Update Help Search Index

Perform this procedure to update the Search index after you have installed all the IBM FileNet P8 Platform and expansion product documentation on a supported application server.

### To update the help Search index

---

**NOTE** Perform this procedure only if you refreshed the core documentation, or you have installed expansion product (or customized application) help onto your IBM FileNet P8 documentation application server. Otherwise, skip to [“Deploy and Verify IBM FileNet P8 Documentation Web Site” on page 485](#).

1. Make sure that the server on which the documentation is deployed is stopped, and that no processes are accessing the documentation.
2. If you have previously deployed IBM FileNet P8 documentation as a web application, undeploy that web application using the procedure for your application server.
3. Make sure you have copied the help for all your various expansion products to a designated application server location containing the IBM FileNet P8 help. Otherwise, you will have to repeat this procedure if you add new help later.
4. Open a command prompt or terminal on the application server.
5. From the command line, navigate to the **search** subdirectory under your **ecm\_help** root directory.
6. Using a text editor, open the search-indexing script file that is appropriate to your application server operating system:

(UNIX) **indexFiles.sh**

(Windows) **indexFiles.bat**

7. If necessary, set the JAVA\_HOME variable in the script file with the path to your JRE installation. The default examples are:

(UNIX) **JAVA\_HOME="/usr/java/j2sdk1.4.1\_02"**

(Windows) **SET JAVA\_HOME=c:\j2sdk1.4.2**

**NOTE** Your JDK reference may be different; use your installed version location if different than the example.

8. Save your changes and close the text editor.
9. If you intend to run the search indexer on a UNIX application server, ensure that you add execute permissions (chmod 755) to the **indexFiles.sh** file.

10. Run the updated search-indexing script file.

By default, the script backs up the existing index files to **indexOld**, and then re-indexes all the help files starting from the root **ecm\_help** directory.

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

11. Go on to the procedure in the next topic "[Deploy and Verify IBM FileNet P8 Documentation Web Site](#)" on page 485.

## ***Deploy and Verify IBM FileNet P8 Documentation Web Site***

Perform this procedure after you have installed (and, if necessary, reindexed) the IBM FileNet P8 documentation on a supported application server.

### **To deploy and verify the IBM FileNet P8 documentation web site**

---

1. Deploy or install the copied IBM FileNet P8 documentation as a web-site application. Use the appropriate instructions provided with your application server.
2. Verify that the application server and the new IBM FileNet P8 documentation web site are running, as follows:
  - a. From your web browser, access the following URL. The documentation's Help Directory should open.

***http://<docserver>:<port#>/<contextRoot>/***

where:

*docserver* is the name of the Java web server.

*port#* is the port number.

*contextRoot* is the value of the Map to URL field that you specified when you deployed the IBM FileNet P8 documentation application. If you specified **/ecm\_help**, then the *contextRoot* is **ecm\_help**.

**NOTE** You can use multi-part root folders (e.g., **/docs/ecm\_help**) if your application server supports them.

- b. Click the **Search** link on the Help Directory toolbar. The documentation Search page should open.
- c. Select one of the Search query result links. The associated help page should open.

**NOTE** When it is time to configure the online help location for the various IBM FileNet P8 components, either while running Setup programs or later via site preferences settings, use the URL in the example in [Step a](#) above.

## Task 2: Verify Current Release Level

Before you begin your IBM FileNet P8 Platform upgrade, you must first verify that your current system is running the software release and component fix pack levels supported for upgrade.

See the section “[Plan the Upgrade](#)” on page 459 for a list of upgrade paths, planning considerations, and pre-upgrade information and tasks that you must take into account, even before performing the task presented below.

### To verify your IBM FileNet P8 system is ready for upgrade

---

1. Run the IBM FileNet P8 Platform 3.5.x Version Tools. See the *IBM FileNet P8 Platform 3.5.x Version Tools Technical Notice*. These tools will identify what release and fix pack level of IBM FileNet P8 software you are currently running. To download this guide from the IBM support page, see “[Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs](#)” on page 21.
2. Ensure that all your users have closed their sessions and no longer have access to the system.

**NOTE** You may choose to shut down all your IBM FileNet P8 Platform components as part of this step. However, be aware that some components (e.g., Content Engine) must be running to begin the actual upgrade process, so the associated setup programs will restart the software as needed.

## Task 3: Upgrade Content Engine Software

Upgrading Content Engine software from version 3.5.x to 4.0.0 involves the following major steps (shown later in this task topic) to be done in the order shown:

1. Prepare Content Engine for upgrade.
  - a. Create a database for the Global Configuration Data (GCD).
  - b. Create or designate user accounts required for upgrading Content Engine.
  - c. Gather information about all your 3.5.x authentication realms.
  - d. If you have not already done so, install and configure the application server on the host machine on which Content Engine 4.0.0 is to be deployed.
  - e. Verify that all in-progress event actions have finished.
  - f. Verify that all publishing requests have been completed.
  - g. Delete custom subclasses of the PublishRequest object.
  - h. Verify that all pending content transactions, indexing requests, and fixed-content migration requests have completed.
  - i. Stop and disable all Content Engine 3.5.x services.
  - j. Back up the databases used by Content Engine.
  - k. If you use the full-text search feature (CBR) in 3.5.x, you must install and configure Content-Search Engine (Autonomy K2) and upgrade existing content-search indexes (Verity collections).
2. Install and deploy Content Engine 4.0.0 on an application server.
3. If your existing FileNet P8 domain includes a SnapLock fixed content device and version 4.0.0 of Content Engine will be on a UNIX platform, prepare the device for upgrade.

### To create the database for the GCD

---

Do the upgrade-applicable configuration procedures listed in one of the following tasks, depending on your database type, to create the GCD:

- [“Verify that DB2 Server Is Installed for IBM FileNet P8” on page 92](#)
- [“Verify that Microsoft SQL Server Is Installed for IBM FileNet P8” on page 81](#)
- [“Verify that Oracle Server Is Installed for IBM FileNet P8” on page 85](#)

**To create or designate user accounts required for upgrade**

Create or designate the following accounts, which are required to complete a Content Engine upgrade. Note that you may have to make some additional configurations to your existing security accounts. See [“Specify IBM FileNet P8 Accounts” on page 60](#) for further account-related details.

User Account	Description
Content Engine Upgrader Tool User	<p>The Windows operating system user who runs Upgrader Tool needs the following permissions:</p> <ul style="list-style-type: none"> <li>• Read access to the Content Engine 3.5.x directory (default location is <b>C:\Program Files\FileNet\Content Engine</b>).</li> <li>• Read access to the 3.5.x version of <b>sysinit.dat</b>.</li> <li>• Read/write access to the root directories for the file storage areas</li> </ul>
P8 4.0 GCD Administrator	<p>To do upgrades, at least one P8 4.0 GCD administrator must:</p> <ul style="list-style-type: none"> <li>• Be an object store administrator for all object stores to be upgraded. Use P8 3.5.x FileNet Enterprise Manager to ensure that this user has Full Control to each P8 3.5.2 object store.</li> <li>• Have Full Control for the P8 3.5.x and P8 4.0.x domains.</li> </ul>
P8 4.0 Content Engine Operating System User	<p>The operating system user that the P8 4.0 Content Engine software runs as. This user must also have read/write access to the root directories for the P8 3.5.x file stores.</p> <p>If version 3.5.x file stores are on a SnapLock fixed content device, then this user must also have read/write access to the device.</p>



User Account	Description
Database User - SQL Server	<p>For each object store you are going to upgrade, grant the associated 3.5.x database user at least the following additional roles and access permissions:</p> <p>SQL Server roles</p> <ul style="list-style-type: none"> <li>• System Administrators</li> <li>• Security Administrators</li> <li>• Database Creators</li> </ul> <p>Database access (for GCD and object stores)</p> <ul style="list-style-type: none"> <li>• public</li> <li>• db_owner</li> </ul> <p>Add to the SQL Server master database and grant these roles:</p> <ul style="list-style-type: none"> <li>• public</li> <li>• SqIDBCXAUser (so the SQL Server user can participate in distributed transactions with the JDBC driver)</li> </ul>
Database User - Oracle	<p>For each object store you are going to upgrade, grant the associated 3.5.x tablespace user at least the following additional roles and access permissions:</p> <ul style="list-style-type: none"> <li>• CREATE SESSION</li> <li>• CREATE TABLE</li> <li>• CREATE SEQUENCE</li> </ul>
Database User - DB2	<p>An operating system user on the database server with the following DB2 database permissions. In the case of a remote database, no equivalent users are needed on the Content Engine server.</p> <ul style="list-style-type: none"> <li>• Connect to the database</li> <li>• Create tables in the tablespace (CREATETAB)</li> <li>• Use the tablespace (USE OF) for User and User Temp tablespaces</li> </ul>

User Account	Description
K2 Security User	When upgrading full-text indexes (Verity collections), Upgrader Tool interacts with the Autonomy K2 server software, and requires a K2 user name and password.
K2 Operating System User	<p>The Autonomy (Verity) K2 software runs as this operating system user.</p> <ul style="list-style-type: none"> <li>• This user must also have read/write access to the root directories for the P8 3.5.x file stores.</li> </ul>

---

**To gather 3.5.x P8 domain authentication information**

Note down the type, Host name and Port of your 3.5.x authentication provider, and have the information available while installing your new Content Engine 4.0.0 system. For more information, see [“Gather 3.5.x authentication information and plan 4.0.x authentication.” on page 466.](#)

---

**To install and configure an application server for Content Engine Server**

1. Install one of the following IBM FileNet P8-supported application servers on the machine that will host version 4.0.0 of Content Engine Server:
  - WebSphere Application Server
  - WebLogic Server
  - JBoss Application Server
2. Configure the application server according to one of the following topics:
  - [“Configure an Application Server for Content Engine \(WebSphere\)” on page 106](#)
  - [“Configure an Application Server for Content Engine \(WebLogic\)” on page 110](#)
  - [“Configure an Application Server for Content Engine \(JBoss\)” on page 118](#)

---

**To verify that all in-progress event actions have finished**

Verify that all in-progress event actions have been processed by launching Enterprise Manager and doing the following for each object store:

1. In the list view, under the object store icon, right-click the Search Results folder, and choose New Search.
2. In the Content Engine Query Builder dialog box, choose QueueItem from the Select From Table list.
3. Retain all default settings and click **OK**. (Click **Yes** at the prompt for a WHERE clause.)
4. If any event items remain in the queue, you will see them in the Query Status dialog box. If no event items appear, then all events have been processed.

If any event items do remain in the queue, you will see them in the Query Status dialog box. To remove unwanted items, set up the same search again; but this time select the Delete check box in the Action tab of the Search dialog box, before clicking **OK**. Click **OK** again to confirm the deletion.

### **To verify that publishing requests have been completed**

---

1. Start Enterprise Manager.
2. In each object store to be upgraded, do the following:
  - a. Expand the Publishing folder and click **Queue** and choose View > All requests.
  - b. Verify that the queue (right pane) of publishing requests is empty. If the queue is not empty, do the following:
    - i. Wait until all publish requests in the In Queue state or In Work state are processed.
    - ii. If any publish requests are in the In Error state, contact your publishing administrator for the appropriate action to take (such as retrying after correcting the error or just deleting the item).
3. In the list view, under the object store icon, right-click the Search Results folder, and choose **New Search**.
4. In the Content Engine Query Builder dialog box, choose PublishRequest from the Select From Table list.
5. Retain all default settings and click **OK**, and then click **Yes** at the prompt for a WHERE clause.
6. Manually delete any publishing requests displayed in the Query Status window.

### **To delete custom subclasses of the PublishRequest object**

---

Before Upgrader Tool can upgrade an object store, you must delete any custom subclasses of the PublishRequest class, along with any instances of such subclasses, as shown in the following steps:

1. Start version 3.5.x of Enterprise Manager.
- Do [Step 2](#) and [Step 3](#) for each object store.
2. In the left pane of Enterprise Manager, open an object store, and navigate to Other Classes > Publish Request.
  3. In the right pane, delete any custom subclasses of Publish Request and any instances of such subclasses.
  4. Exit from Enterprise Manager.

### **To verify that all in-progress Content Engine transactions involving content files have finished**

---

Before upgrading Content Engine, you must verify that all pending content transactions, indexing requests, and fixed-content migration requests have completed. You can do this by running the Content Resource Manager utility, **RMU.exe**, as shown in the following steps:

**To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain**

**CAUTION** All transactions and requests must have completed before you upgrade Content Engine; otherwise the upgrade will fail.

1. On each machine where Content Engine 3.5.x is installed, do the following:
  - a. Start File Store Service if it is not already running.
  - b. Stop Object Store Service and Content Cache Service.
2. On a machine where Content Engine 3.5.x is installed, start **RMU.exe**, located (by default) at **C:\Program Files\FileNet\Content Engine**.
3. In the Content Resource Manager Utility window, do the following:
  - a. Navigate to Statistics > Transactions and note the values of *Total prepared <phase 1>*, *Total committed*, and *Total aborted*.
  - b. Navigate to Statistics > Can't Do Queue and note the value of *Current items in the can't do queue*.
  - c. Add the values of *Total committed* and the *Total aborted*. If the sum equals the value of *Total prepared <phase 1>*, and the value of *Current items in the can't do queue* is zero, then all content transactions have completed.
  - d. Navigate to Indexing Service > Index Control. If the value of *Current index queue files* is zero, then all pending indexing requests have completed.
  - e. Navigate to Fixed Content > Migration Queuing. If the value of *Current number in queue* is zero, then all fixed-content migration requests have completed.
4. If all activity checked in [Step 3](#) has completed, then stop File Store Service on each machine where Content Engine 3.5.x is installed, and continue at ["To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain" on page 492](#); otherwise, do the following:
  - a. Exit from the Content Resource Manager utility.
  - b. Wait a few minutes, and then return to [Step 1](#).

**To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain**

1. On each machine where version 3.5.x of Content Engine Server software is installed, log on with local administrator permissions.
2. Stop and disable the following services:
  - Apache2
  - Content Engine Content Cache Service
  - Content Engine File Store Service
  - Content Engine Object Store Service
  - FileNet Publishing HTML Plug-in Service
  - FileNet Publishing PDF Plug-in Service
  - Process Services Manager

- Wasp Server for Java

3. Right-click the Apache services monitor in the Windows system tray and click **Exit**.

#### **To verify the directory permissions on file stores**

---

If the application server on which Content Engine Server 4.0.0 is to be deployed is running on Windows, then give the following users write permission to the root directory where a 3.5.x file store is located. (This procedure is necessary to allow conversion to a 4.0.0 file storage area during upgrade.)

- The user running the Upgrader Tool
- The user running the application server (for WebSphere on Windows, this user is a service log-on account)

#### **To back up the database**

---

Use your existing database backup solution to back up the Content Engine 3.5.x object store databases.

#### **To unconfigure CBR from Content Search Engine 3.5.x object stores**

---

Before upgrading Content Engine to version 4.0.x, unconfigure and remove CBR indexes for all documents. As part of the Content Engine upgrade procedures, you will be instructed later in this guide to reinstall Content Search Engine and reindex all the documents in your upgraded version 4.0.x object stores.

To unconfigure and remove CBR indexes for all documents in existing version 3.5.x object stores, navigate in the P8 3.5.x Help to FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Unconfigure CBR.

#### **To upgrade Content Engine**

---

1. Continue at [“Install and Deploy Content Engine” on page 160](#) to run Content Engine Setup to install version 4.0.0 Content Engine Server. On a Windows machine, Content Engine Setup automatically installs Enterprise Manager.

**CAUTION** Do not install Enterprise Manager on any machine running the 3.5.x version—at least until *after* the upgrade of Content Engine to version 4.0.0 is complete. Otherwise, you will no longer be able to run the 3.5.x version of FileNet Enterprise Manager against any remaining 3.5.x object stores.

2. If you have NetApp fixed content devices, continue at [“To prepare NetApp SnapLock Volumes for the upgrade” on page 493](#); otherwise, go to [“Install Content Engine Software Updates” on page 497](#).

#### **To prepare NetApp SnapLock Volumes for the upgrade**

---

If each of the following conditions apply, you must do this procedure; otherwise continue at [“Install Content Engine Software Updates” on page 497](#):

- Version 4.0.0 of Content Engine is on a UNIX platform.

- Your existing FileNet P8 domain includes a SnapLock fixed content device.

Do this procedure for each of your NetApp filers (network-attached appliances for data storage) to allow access to your fixed content via NFS instead of, or in addition to, CIFS.

1. Check prerequisites

- 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 4.0.0 of IBM FileNet P8. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)

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

- Make sure you are licensed to use NFS to access your NetApp filers.

2. Set the security style of a NetApp storage volume to enable support for NFS clients

Each qtree (virtual subvolume of a storage volume) has exactly one of the security styles (scheme for setting security on files and directories in the qtree) shown in the following table:

Security Style	Description
UNIX	UNIX file permission attributes. Only NFS clients can create files and directories in a UNIX qtree.
NTFS	Windows access control lists. Only CIFS clients can create files and directories in an NTFS qtree.
Mixed	Both UNIX and NTFS security styles. Only one security style at a time is allowed. The current style is that of the last client to modify it.

Since all pre-4.0.0 versions of IBM FileNet P8 support only CIFS, all NetApp storage volumes used by Content Engine Server use NTFS security style, which you must change to UNIX or Mixed. Specify the Mixed style for qtrees that must service requests for both NFS and CIFS clients during the upgrade process; otherwise, specify UNIX style.

For each qtree, do the following steps to specify the security style:

- Access the Data ONTAP administrative console (see the ONTAP System Administration Guide for information on administrative access methods).
- Run the qtree command as in the following example, which sets the UNIX security style on the qtree `/vol/vol1/sa1` of NetApp filer `NAFiler` the security style to UNIX:

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

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

### 3. Map UNIX users and groups to Windows equivalents

Because version 3.5.x of Content Engine Server support CIFS, rather than NFS, all existing files on a NetApp volume have NTFS security attributes; and the users and groups with access rights to these files are defined by Windows.

To allow version 3.5.x files to remain accessible, you must create a mapping between the new UNIX account for version 4.0.0 of Content Engine Server and the old Windows account for version 3.5.x.

Each NetApp filer has its own configuration file, `/etc/usermap.cfg`, to map between Windows user names and equivalent UNIX user names. A UNIX user attempting to access a file having NTFS security attributes uses `usermap.cfg` to determine if a mapping exists between the UNIX account and an equivalent Windows account. If the mapping exists, the access checks on the target file will use the Windows account.

Each `usermap.cfg` entry has the following format:

```
[IP_qualifier:] Windows_name [direction] [IP_qualifier:] UNIX_name
```

The meaning of each element in the entry is shown in the following table:

Element	Meaning
<i>IP_qualifier</i>	Qualifies the name according to the source address of the requester
<i>Windows_name</i>	The name of the Windows user or group in domain name format (for example, <i>DomainName\UserName</i> ). The Windows name must be in the Windows domain that the NetApp filer is configured to use when authenticating Windows users.
<i>UNIX_name</i>	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 <code>/etc/passwd</code> (for users) or <code>/etc/group</code> (for groups).  If it is a group, it may be necessary to also define it in an NIS repository or an LDAP directory server, depending on how the filer is configured.  In either case the UID (UNIX user ID) must be identical to the UID of the user under which the Content Engine Server is executing.
<i>Direction</i>	The direction of the mapping, either <code>&lt;=</code> or <code>=&gt;</code> .  <code>&lt;=</code> : Maps <i>UNIX_name</i> to <i>Windows_name</i>  <code>=&gt;</code> : Maps <i>Windows_name</i> to <i>UNIX_name</i>

For example, the following steps define a mapping on a NetApp filer between `FNCE_OS_User` (the Windows user account under which version 3.5.x of Content Engine Server executes) and

FNCE\_UNIX\_User (the UNIX user account under which version 4.0.0 of Content Engine Server executes).

- a. Log on to the machine where version 3.5.x of the Content Engine is installed.
- b. Connect to the root volume on the target NetApp filer using the Administrator account.

By default, the NetApp filer root volume is accessible from a Windows client as a CIFS share named C\$, as in the following example (where *NAFiler* is the host name of the target NetApp filer at your site).

```
C\> net use n: \\NAFiler\C$ /user:NAFiler1\Administrator
```

- c. Edit **/etc/usermap.cfg** by adding the following stanza:

```
CEDomain\FNCE_OS_User => FNCE_UNIX_User
```

- d. Edit **/etc/passwd** by adding the following stanza:

```
FNCE_UNIX_User:CEServers:205:7100::/home/FNCE_UNIX_User:
```

4. Continue at [“Install Content Engine Software Updates” on page 497](#).



## Task 4: Install Content Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Content Engine.

### To install the Content Engine software updates

---

1. To download the latest software update, and to determine whether additional interim fixes are needed, contact your service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. P8CE-4.0.0-002 Service Pack (or later)
  - b. Any subsequent interim fixes (typically optional)
3. Continue at [“Upgrade Content Engine Data”](#) on page 498.

## Task 5: Upgrade Content Engine Data

Upgrading Content Engine data from version 3.5.x to 4.0.1 involves the following major steps (shown later in this task topic), to be done in the order shown:

1. Configure database connectivity from the application server for your version 3.5.x object stores.
2. Upgrade Content Engine data
  - a. Edit the upgrader utility file **CE401Upgrader.bat**.
  - b. Run Upgrader Tool from a command line or a graphical interface to upgrade Content Engine 3.5.x items, including the GCD, object stores, file storage areas, etc. to version 4.0.1.
3. Complete post-upgrade Content Engine configuration.
  - a. Manually create a CodeModules folder within the root folder of each upgraded object store.
  - b. Clear the read-only attribute for NTFS file storage areas.
  - c. (Optional) Move file storage areas and Content Search areas (collections) from Windows machines to UNIX machines.

### NOTES

- You need not check in checked-out documents before upgrading Content Engine.
- The Upgrader Tool must run on a Windows machine with at least 1.5 GB of available memory (as indicated by Windows Task Manager)
- Version 1.4.2 of the Java Runtime Environment must be on the machine where you are going to run Upgrader Tool.
- If the maximum heap size of the JVM is 1 GB or more, do not run Upgrader Tool on the same machine where any other major application is running (such as the database used by Content Engine) unless the machine has at least 2 GB of RAM.
- To upgrade Content Engine 3.5.x data—object stores, addons, file stores, fixed content devices (FCDs), etc.—to version 4.0.1, you will use Upgrader Tool, which you can run interactively, via a graphical user interface (GUI), or silently, via a command line interface (CLI). You can also switch between the two methods during the upgrade.
- If you are upgrading an IBM FileNet P8 environment containing a large number of object stores, IBM recommends the following approach:
  - When using Upgrader Tool, upgrade at most 20 object stores at a time.
  - If the application server where version 4.0.1 of Content Engine will be deployed is running on a 32-bit JVM, you should restrict each 4.0.1 FileNet P8 domain to no more than 50 object stores. If 3.5.x FileNet P8 domains contain more than 50 object stores, consider partitioning them into multiple 4.0.1 FileNet P8 domains during the upgrade process.
  - Ensure that 3.5.x object stores having the same basic set of system objects (for file stores, fixed file stores, etc.) are upgraded into the same 4.0.1 FileNet P8 domain.

- For non-English support relating to collections, refer to *IBM FileNet P8 Platform Installing Non-English Environments Technical Notice*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### **To configure application server database connectivity for existing object stores.**

---

1. Because version 4.0.x object stores require you to set up application server connectivity to your database, do the procedure in one of the following topics:
  - [“Configure Content Engine Application Server Database Connectivity \(WebSphere 5.1.x\)” on page 193](#)
  - [“Configure Content Engine Application Server Database Connectivity \(WebSphere 6.0.x\)” on page 206](#)
  - [“Configure Content Engine Application Server Database Connectivity \(WebLogic 8.1.x\)” on page 236](#)
  - [“Configure Content Engine Application Server Database Connectivity \(WebLogic 9.2.x\)” on page 241](#)
  - [“Configure Content Engine Application Server Database Connectivity \(JBoss 4.0.x\)” on page 247](#)
2. If you have Content Federation Services (CFS) fixed content devices, continue at [“To configure CFS fixed content devices for upgrading” on page 499](#); otherwise, skip to [“To edit the upgrader utility file” on page 499](#).

### **To configure CFS fixed content devices for upgrading**

---

If your version 3.5.x CFS fixed content devices use MS SQL Server databases, you need to perform the following steps; if your CFS fixed content devices use non-MS SQL Server database, then skip to [“To edit the upgrader utility file” on page 499](#).

1. For each CFS database, add the following roles to the CFS user:
  - System Administrators
  - Security Administrators
  - Server Administrators
  - Database Creators
2. For each CFS database, verify the following access permissions for the CFS database user:
  - public
  - db\_owner
3. Continue at [“To edit the upgrader utility file” on page 499](#).

### **To edit the upgrader utility file**

---

1. Navigate to the (default) location **C:\Program Files\FileNet\ContentEngineUpgrader**, which contains **CE401Upgrader.bat** and do one of the following, depending on whether you are going to upgrade via the command line interface (CLI) or the graphical user interface (GUI):

- (CLI) Change the line immediately after `CLI:` to the following (without carriage returns):

```
java -Xms512m -Xmx1024m -cp "%CLASSPATH%" %JAVA_OPTIONS% -Dwaspl.location=../wsi"
com.filenet.upgrader.ui.UpgradeUtility %1 %2 %3 %4
```

- (GUI) Change the line immediately after `GUI:` to the following (without carriage returns):

```
java -Xms512m -Xmx1024m -cp "%CLASSPATH%" %JAVA_OPTIONS% -Dwaspl.location=../wsi"
com.filenet.upgrader.ui.MainFrame
```

**NOTE** As shown in the above commands, the recommended minimum and maximum JVM heap sizes (-Xms and Xmx) are 512 MB and 1024 MB. If you are upgrading a system having only a few object stores that do not have many custom objects, you can reduce the minimum and maximum heap size arguments in these commands to 256 MB and 512 MB, respectively.

2. (Oracle only) Edit the file **CE401Upgrader.bat** to include the path to the Oracle JDBC Driver JAR file **ojdbc14.jar** in the system CLASSPATH environment variable.
3. (SQL Server only) Edit the file **CE401Upgrader.bat** to include the SQL Server 2005 JDBC Driver file **sqljdbc.jar** in the CLASSPATH system environment variable.
4. (DB2 Server only) Edit the file **CE401Upgrader.bat** to include the DB2 Server JDBC Driver files **db2jcc.jar**, **db2jcc\_license\_cisuz.jar**, and **db2jcc\_license\_cu.jar** in the CLASSPATH system environment variable.
5. Continue at one of the following methods:
  - [“Graphical User Interface to Upgrader Tool” on page 500](#)
  - [“Command Line Interface to Upgrader Tool” on page 507](#)

Even if you plan to run Upgrader Tool using the graphical user interface method, it is helpful to first read the command line interface method. Both methods involve the same basic steps:

1. Create an XML upgrade status file.
2. Run the upgrade utility, driving it from the XML upgrade status file.

## ***Graphical User Interface to Upgrader Tool***

The GUI version of Upgrader Tool follows the same sequence as the CLI version, but uses an interactive “shell,” in which each step is tied to the next, and in which you can check the accuracy of configuration parameters before attempting the actual upgrade.

### **To check Upgrader Tool prerequisites**

---

1. Log on as the Content Engine Upgrader Tool User (defined earlier in [“To create or designate user accounts required for upgrade” on page 488](#)) to the machine where Upgrader Tool is installed.
2. Choose Start > Programs > FileNet P8 Platform > Object Store Upgrader.
3. Choose File > New Upgrade 3.5.2 to 4.0.1....

4. At the *Open CE 3.5.2 GCD file* screen, specify the GCD file **sysinit.dat**, whose default version 3.5.2 location is **FileNet\Content Engine\sysconfig\sysinit** and click **Open**.

**NOTE** A limitation in JVM 1.4.2 causes it to display all file types in the screen, not just the ones that match **sysinit.dat**.

5. At the *Save upgrade status to xml file* screen, specify where to save **sysinit.dat** and click **Save**.
6. Before Upgrader Tool opens its first screen, it tests the following:
  - the location of the log4j (logger) configuration file, typically **log4j.properties**.
  - whether the custom UpgradeAppender is referenced by the **log4j.properties**.

In case of error, Upgrader Tool then opens the System Precheck for CE 3.5.2 Upgrade screen and indicates the results of these tests: A green dot indicates success; a red dot failure.

If any test failed, a red dot appears on the **OK** button of the screen and Upgrader Tool exits. Correct the failure, and retry this step.

If every test succeeds, continue at [“To create the XML upgrade status file” on page 501](#).

### To create the XML upgrade status file

---

1. The CE 3.5.2 to CE 4.0 Upgrader screen opens. If you are resuming an upgrade (from a previous effort to upgrade Content Engine), choose **File > Resume Upgrade...** and continue at [Step 2](#). Otherwise, skip to [Step 3](#) to start the process of creating the file.
2. In the *Open XML upgrade status file* screen, select the XML upgrade status file (for example, **upgrade.xml** in this procedure) and click **Open**. Continue at [“To configure JAAS, P8 domain, CBR, and FCD parameters” on page 501](#).
3. In the *Save upgrade status to xml file* screen specify the name of the XML upgrade file that Upgrader Tool will create (for example, **upgrade.xml** in this procedure) and click **Save**.
4. In the XML Upgrade File Creator screen, select or clear the CBR check box, depending on whether your Content Engine 3.5.2 installation has a Verity content-based retrieval (CBR) server; and then click **Start** to begin creating the XML upgrade status file.

Upgrader Tool sends logging information to the Log window of this screen. (If no output appears in the Log window, the **log4j.properties** file is probably not pointing to the UpgradeAppender.)

5. Once Upgrader Tool writes a message into the Log window to say that it has successfully created the XML upgrade file, click **Close** and continue at [“To configure JAAS, P8 domain, CBR, and FCD parameters” on page 501](#).

### To configure JAAS, P8 domain, CBR, and FCD parameters

---

Upgrader Tool displays the Upgrader screen, which contain up to five tabs (five, if you selected CBR and have fixed content devices) when creating the XML upgrade status file:

- JAAS/Domain Config
- CBR Config. (displayed only if you selected the CBR check box in [Step 4](#) of [“To create the XML upgrade status file” on page 501](#))

- Fixed Content Devices (displayed only if at least one FCD exists in the FileNet P8 domain)
- Upgrade
- Report

You will access the tabs in the order shown above. This procedure deals with the first three tabs.

For the JAAS/Domain Config and CBR Config tabs, you can check the validity of your settings by clicking the **Test** button on the tab. If all settings for a given tab are valid, the icon next to the **Test** button turns green. If any settings are invalid, you must correct them and then click **Test** again.

Upgrader Tool places your valid settings into the XML upgrade status file. When all your settings in a given tab are valid, Upgrader Tool moves to the next tab.

1. Specify the parameter values in the JAAS/Domain Config. tab, as shown in the following table, and then click **Test** to validate your settings.

Parameter	Description
UserID	P8 4.0 GCD Administrator in <a href="#">Step</a> of <a href="#">"To create or designate user accounts required for upgrade" on page 488</a>
Password	Password for UserID
Name	Name of the FileNet P8 4.0.x domain.
URI	URI to IBM FileNet P8 application. Example: <code>http://&lt;ApplicationServerName&gt;:&lt;port&gt;/wsi/FNCEWS40DIME</code>

2. If you selected CBR when creating the XML upgrade status file, specify the parameter values from your Autonomy K2 installation in the CBR Config. tab, as shown in the following table, and then click **Test**. Autonomy K2 security account information is required. For more information on the accounts required, see ["To create or designate user accounts required for upgrade" on page 488](#) for details on which accounts to assign and the permissions to enable.

Parameter	Description
Username	K2 Security User
Password	K2 Security User password
Style Set Alias	FileNet_FileSystem_PushAPI
Index Server Names	K2 Index Server names
Brokers	K2 Broker Server names

Parameter	Description
User Domain	Domain on which the K2 services run
User Group	K2 Security Group
Server Host Name	Name of host where K2 Master Administration Server is installed
Server Port	K2 Master Administration Server port
K2 Server Names	K2 Server names

To specify multiple Index Server names, K2 Server names, or K2 Broker Server names, press **<Enter>** after typing each value to bring the cursor to a new line. Do not use any other delimiters (such as commas or spaces) to separate your values.

- The left-hand panel of the Fixed Content Devices tab, displays all the FCDs referenced in the XML upgrade status file.

**NOTE** Hitachi, Tivoli, and IBM System Storage™ DR550, fixed content device types supported in version 3.5.x of Content Engine, are not supported in the initial release of version 4.0.0. If you need to migrate content from these devices to version 4.0.0, you must first move the content to a fixed content device supported in version 4.0.0.

Specify the parameter values for each FCD in turn, and then click **Accept**. Upgrader Tool will update the XML upgrade status value accordingly and then automatically select the next FCD to be configured.

For descriptions of these parameters, refer to IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > FileNet P8 Domain > How to... > View/modify P8 domain properties > Fixed Content Devices \(General tab\)](#).

## Prepare Items for Upgrade

The Upgrade tab consists of three panes:

- Upgrade Tree (left side of tab)
- Properties (upper-right side of tab)
- Log (lower-right side of tab)

The node labels for object stores and databases are preceded by a yellow icon, to indicate that Upgrader Tool requires additional configuration information for these nodes. Note also that this color-coded status propagates upward within the tree to all parent nodes as well.

Selecting an item (node) in the Upgrade Tree pane causes the Tree Node Description to display node-specific information in several tabs. For example, selecting an object store causes the Properties pane to display three tabs: JNDI Settings, Event Actions, and Description.

The P8CoreDomainAddons nodes and the P8CoreObjectStoreAddons nodes in the Upgrade Tree pane each have two lists of AddOns, for versions 3.5.2 and 4.0.1 of Content Engine, which you

can see in the Properties pane. Note that the addons which are presented in both versions 3.5.x and 4.0.0 are shaded in both lists.

To prepare an item for upgrade, select its node, and set the parameter values in each tab of the Properties pane, and save your settings by clicking **Save Settings**. The following procedures describe how to prepare each type of item for upgrade.

**To prepare object store nodes for upgrade**

---

For each object store in the Upgrade Tree pane, do the following steps to prepare it for upgrade:

1. Select the object store node in the Upgrade Tree pane.
2. In the JNDI Settings tab, specify values for the following parameters:
  - Local JNDI Name (that is, the non-XA value)
  - Global JNDI Name (that is, the XA value)

Click **Save Settings**. If the associated database does *not* need to be configured (no yellow Needs Info icon next to the Database node), then yellow icon next to the object store will be removed.

3. If the associated Database node needs info, select the Database node under the object store, select the Database Settings tab, and do the following:
  - a. Specify the database username, password, database type, class, and URL, and then click **Test**. The following table shows examples of URLs for each database type:

Database Type	URL
DB2	jdbc:db2:<database_name>
MS SQL Server	jdbc:sqlserver://<DB_hostname>:1433
Oracle	jdbc:oracle:thin:@<DB_hostname>:1521:<Oracle_instance_service_name>

- b. Select the Object Store node, and then the Description tab. Click **Take Offline**. (This button toggles between **Take Offline** and **Take Online**.) Upgrader Tool then checks for conflicting object store class and property names between the old and new versions of Content Engine.

If it finds no conflicts, Upgrader Tool takes the object store offline and appends the updated (offline) state of the object store to its name in the Upgrade Tree.

**NOTE** To allow all ongoing activity to finish, you should normally wait several minutes after taking the object store offline before actually upgrading it.



4. If it does find conflicts, Upgrader Tool displays the conflicting class names and property names in a window, and the object store remains online. Click **OK** and then do the following for each conflict:
  - a. Resolve the conflict (using Enterprise Manager).
  - b. Click **Reset** (to toggle the node from its error condition back to its ReadyToUpgrade state) and then redo [Step 3](#).

### To prepare file store and fixed file store nodes for upgrade

---

If version 4.0.0 of Content Engine is deployed on a UNIX machine, do the following steps for each file store and fixed file store to adjust directory paths accordingly. Autonomy K2 security account information is required. For more information on the accounts required, see [“To create or designate user accounts required for upgrade” on page 488](#) for details on which accounts to assign and the permissions to enable.

**NOTE** Certain non-English locales are not supported and will not be upgraded. Refer to *IBM FileNet P8 Platform Installing Non-English Environments Technical Notice*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

1. Select a file store (or FCD) in the Upgrade Tree pane.
2. In the File Store Settings tab set the root path to its UNIX equivalent and click **Save Setting**.
3. If no file store or FCD contains full-text (CBR) indexes (Verity collections) to upgrade, then its preparation is complete; otherwise continue at [Step 4](#).
4. Locate the directories containing collections and set permissions to allow access to the following users:
  - Content Engine Operating System User.
  - K2 Operating System User.
5. For each directory containing 3.5.x content-search indexes (Verity collections) that you want to upgrade to 4.0.0 index areas (K2 collections), make collections readable by Verity by opening the **Verity.cfg** file and entering the following information:
  - alias: Path number that increments for each path you list.
  - mapping: full path to the collections directory.
  - dirmode: permission value that you must set to *wr* (write and read).
  - a. Open the Verity configuration file in a text editor:
 

**C:\ProgramFiles\filenet\contentengine\verity\k2\common\verity.cfg**
  - b. Modify the next available alias settings by entering the information listed above for each collections directory you will upgrade.

For example, if the next available settings are number 6 and you want to upgrade collections on *myserver*, located in *FileStores\myfilestore\index*, you would change *alias6*, *mapping6*, and *dirmode6* to the following:

```
alias6=path1
mapping6=\\myserver\FileStores\myfilestore\index
dirmode6=wr
```

To add another directory, *myotherserver\collections\index* for example, you would modify settings for number 7 as follows:

```
alias7=path2
mapping7=\\myotherserver\collections\index
dirmode7=wr
```

## Perform the Upgrade

For an item to be qualified for upgrading, two conditions must be met:

- The check box for the corresponding node in the Upgrade Tree pane must be selected.
- The item must have a status of ReadyToUpgrade.

Each selected check box maps to a *yes* value of *SelectedForUpgrade* within the XML upgrade status file for the corresponding item.

When you click **Start** to initiate an upgrade, the icon next to each item in the Upgrade Tree pane changes color as its state changes. The color code for these states is indicated in the Upgrade Key at the bottom of the Upgrade Tree pane. The icon for any item whose upgrade fails turns red.

As the item upgrade proceeds, its corresponding Status value in the XML upgrade status file changes to reflect the current state of the upgrade. The contents of the Log pane also show how the upgrade progresses.

If you click **Stop**, the upgrade will stop after completing the current step for the item being upgraded, leaving a Status value of *UpgradeFailed* in the XML upgrade status file. Before trying again to upgrade an item whose previous upgrade attempt failed, do the following:

1. Fix the error in the item before retrying the upgrade.
2. With its check box selected, click **Reset** to reset the Status value of the item to ReadyToUpgrade.

You can also click **Stop** if the log indicates a condition that makes it pointless to try to upgrade any more of the items selected for upgrade.

There is an order constraint on upgrading the items: You must upgrade *CommonGCD* *before* upgrading any object stores, as indicated in the following procedure.

### To upgrade items

---

1. In the Upgrade Tree pane, select the check box for the *CommonGCD* item and clear the check boxes for all other nodes in the pane.

2. Click **Start** to initiate the upgrade and wait until the log or the color of the icon indicates that the upgrade is complete.
3. In the Upgrade tree pane, clear the check box for the CommonGCD item and select the check boxes for all the object stores to be upgraded.
4. Click **Start** to initiate the upgrade and wait until the log or the color of the icon indicates that the upgrade is complete.
5. In the Upgrade Tree pane, select any other items to be upgraded, and clear the check boxes for all those items already upgraded.
6. Click **Start** to initiate the upgrade and wait until the log or the color of the icon indicates that the upgrade is complete.

As the upgrade proceeds, the success or failure of each item being upgraded is captured in a viewable report, which you can access in the Report tab. This report is the same as that generated by running the command-line version of Upgrader Tool.

**NOTE** (DB2 only) If a Transaction Log Full exception (SQL ErrorCode -964) occurs during the upgrade, it is recommended that you increase the maximum number of secondary log files (LOGSECOND) and then re-do this step. If you encounter the exception after the log file counts have been doubled, contact your IBM FileNet service representative.

7. Click **Save** to save the report to disk (in HTML format).
8. If Upgrader Tool successfully upgrades every item continue at [“Complete Post-Upgrade Content Engine Configuration” on page 513](#). Otherwise, do the following:
  - a. Use the information (exception code and stack trace) in the command-line output or the **log4j** file to correct the error.
  - b. Select the failed item and click **Reset** on the Description tab, which will change the status of the item from UpgradeFailed to ReadyToUpgrade.
  - c. Return to [Step 1](#).

## **Command Line Interface to Upgrader Tool**

In this section you will create an XML upgrade status file and run the Upgrader Tool using the command-line interface (CLI) method.

### **To create an XML upgrade status file**

---

This procedure interrogates the GCD of the Content Engine 3.5.x installation to produce an XML file containing representations of the items that can be upgraded.

The XML file contain placeholders, which you will need to manually edit, for system settings that cannot be derived from the GCD, such as JAAS settings. Each placeholder is indicated by the character string '###'.

1. Log on as a local administrator to the machine where Upgrader Tool is installed.
2. Navigate to the (default) location **C:\Program Files\FileNet\ContentEngine\Upgrader**, which contains **CE352To40Upgrader.bat**.

3. Note the path **<GCD\_Path>** to the GCD file **sysinit.dat** on the Content Engine server machine.
4. Designate a path **<XML\_Path>** for the XML upgrade status file to be generated in [Step 5](#).
5. Run the following command to create the XML upgrade status file **upgrade.xml**:

```
CE352To40Upgrader.bat -i"<GCD_Path>/sysinit.dat" -o"xml_upgrade.xml"
```

6. Manually edit **xml\_upgrade.xml** as required for your site, as follows:
  - Replace each occurrence of the string '###' with information that is appropriate for your site.
  - The passwords you specify must be in plain text. Upgrader Tool will encrypt these passwords, as well as any other sensitive data, such as fixed content device parameters.
  - A 'yes' value of SelectedForUpgrade means you want to upgrade the item; a 'no' value means you do not want the item to be considered for upgrade.
  - The Status attribute values have the meanings shown in the following table:

Status Value	Description
NeedsInfo	Additional configuration information is needed. Set this field value to ReadyToUpgrade only after you have supplied the required information.
ReadyToUpgrade	The item will be upgraded if its SelectedForUpgrade value is yes.
UpgradeStarted	Upgrader Tool has started upgrading the item.
UpgradeFinished	Upgrader Tool has successfully upgraded the item.
UpgradeFailed	Upgrader Tool has failed to upgrade the item.
Unsupported	Devices (FCDs only) not supported by Upgrader Tool

- An item will be upgraded only if its SelectedForUpgrade value is yes and its Status value is 'ReadyToUpgrade'.
  - If you do not want to include Content Search Engine as part of your Content Engine upgrade, then you need not install the underlying Autonomy (Verity) K2 software, and delete the part of **xml\_upgrade.xml** that starts with **<Verity>** and ends with **</Verity>**.
7. If no file store or FCD contains full-text (CBR) indexes (Verity collections to upgrade, then its preparation is complete; otherwise continue at [Step 8](#).
  8. Enter the information below for the K2 variables in the XML file. Autonomy K2 security account information is required. For more information on the accounts required, see ["To create or](#)

designate user accounts required for upgrade” on page 488 for details on which accounts to assign and the permissions to enable.

Parameter	Description
UserName	K2 Security User
UserPassword	K2 Security User password
UserDomain	Domain on which the K2 services run
UserGroup	K2 Security Group
AdminServerHost Name	Name of the host on which the K2 Master Administration server is installed
AdminServerPort	K2 Master Administration Server port
StyleSetAlias	FileNet_FileSystem_PushAPIFileNet_FileSystem_PushAPI
K2ServerNames	K2 Server names
IndexServerNames	K2 Index Server names
Brokers	K2 Broker Server names

To specify multiple Index Server names, K2 Server names, or K2 Broker Server names, press **<Enter>** after typing each value to bring the cursor to a new line. Do not use any other delimiters (such as commas or spaces) to separate your values.

9. Locate the directories containing collections and set permissions to allow access to the following users:
  - Content Engine Operating System User
  - K2 Operating System User
10. For each directory containing 3.5.x content-search indexes (Verity collections) that you want to upgrade to 4.0.0 index areas (K2 collections), make collections readable by Verity by opening the **Verity.cfg** file and entering the following information:
  - alias: Path number that increments for each path you list.
  - mapping: full path to the collections directory.
  - dirmode: permission value that you must set to *wr* (write and read).
  - a. Open the following Verity configuration file in a text editor:
 

**C:\ProgramFiles\filenet\contentengine\verity\k2\common\verity.cfg**
  - b. Modify the next available alias settings by entering the information listed above for each collections directory you will upgrade.

For example, if the next available settings are number 6 and you want to upgrade collections on *myserver*, located in *FileStores\myfilestore\index*, you would change *alias6*, *mapping6*, and *dirmode6* to the following:

```
alias6=path1
mapping6=\\myserver\FileStores\myfilestore\index
dirmode6=wr
```

To add another directory, *myotherserver\collections\index* for example, you would modify settings for number 7 as follows:

```
alias7=path2
mapping7=\\myotherserver\collections\index
dirmode7=wr
```

11. Continue at [“To run Upgrader Tool using the CLI” on page 510](#).

### To run Upgrader Tool using the CLI

---

You will now run Upgrader Tool, **CE352To40Upgrader.bat**, with the file **upgrade.xml** you generated in [“To create an XML upgrade status file” on page 507](#) as input to drive the actual upgrade.

Before upgrading an object store, Upgrader Tool takes the object store offline. After upgrading an object store, Upgrader Tool updates the corresponding Status value in **upgrade.xml**.

1. Log on as the Content Engine Upgrader Tool User (defined earlier in [“To create or designate user accounts required for upgrade” on page 488](#)) on the machine where Upgrader Tool is installed.
2. Navigate to the (default) location **C:\Program Files\FileNet\ContentEngine\Upgrader**, which contains **CE352To40Upgrader.bat**.
3. (Optional) To see the available options, run Upgrader Tool from a command line, as follows:

```
CE352To40Upgrader -h
```

Notice from the command output that Upgrader Tool allows you to do the following:

- Specify the amount of time (in seconds) for it to wait after taking an object store offline before upgrading it (using the **-d** option).
  - Generate an HTML report of the upgrade (using the **-r** option).
4. Run Upgrader Tool, specifying options shown in [Step 3](#).  
As it attempts to upgrade each item, Upgrader Tool sends a status message to the command line and to a log4j logging system. If it fails in upgrading an item in the GCD, Upgrader Tool moves on to the next item; if it fails in upgrading an object store, Upgrader Tool will halt.
  5. If Upgrader Tool successfully upgrades every item in [Step 4](#), then continue at [“Complete Post-Upgrade Content Engine Configuration” on page 513](#). Otherwise, do the following:
    - a. Use the information (exception code and stack trace) in the command-line output or the **log4j** file to correct the error.

- b. Edit **upgrade.xml** by replacing any Status value of UpgradeStarted or UpgradeFailed to ReadyToUpgrade.
- c. If the failure occurred due to an error *after* the database upgrade has completed, then restore your database from backup.
- d. Return to [Step 4](#).

## Task 6: Install Content Search Engine Software Updates

Install any service packs, fix packs or/or interim fixes required for Content Search Engine.

### **To install the Content Search Engine software updates**

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. Content Search Engine 4.0.1 Service Pack
  - b. Any subsequent interim fixes (typically optional)



## Task 7: Complete Post-Upgrade Content Engine Configuration

Do the procedures in this task to complete the upgraded Content Engine configuration.

### Check Completion Status of Asynchronous Upgrade Events

When the Upgrader Tool shows that an object store has been upgraded, there will still be two asynchronous upgrade events running in the background against the object store:

- Security Upgrade
- MIME Type Upgrade

Although Content Engine is functional while these asynchronous upgrade events are running, system performance may be diminished. To determine whether an asynchronous upgrade event has completed, run the following SQL query against the object store database:

```
select async_upgrade_state from DDState
```

This query will return one row, and the following table shows how to interpret the value of `async_upgrade_state`:

async_upgrade_state value	Security Upgrade state	MIME Type state
0	Completed	Completed
1	Running	Completed
2	Completed	Running
3	Not completed	Not completed

### Create CodeModules Folders for Upgraded Object Stores

For each upgraded object store, do the following procedure to manually create a corresponding CodeModules folder. Without these folders, you will not be able to create event actions via Enterprise Manager.

#### To create a CodeModules folder for an object store

---

1. Log on to Enterprise Manager as an administrative user.
2. Navigate to the root folder of the object store.
3. Within the root directory of the object store, create a folder named *CodeModules*.
4. Set the *IsHiddenContainer* property value of the CodeModules folder to *true*.

## Clear the Read-Only Attribute for NTFS File Storage Areas

The following procedure is needed to allow or retain secure deletes of NTFS file storage areas in IBM FileNet P8 4.0.0.

**NOTE** The time required to clear the read-only attribute can be many hours, or even several days, if the file storage areas contain millions of content elements (files) or are on slow machines. You can still access files while this procedure is in progress; however, any attempt to securely delete a file whose read-only attribute is not yet cleared will fail. In which case, the delete operation will be re-queued until the read-only bit has been cleared.

### To clear the read-only attribute for NTFS file storage areas

---

For each NTFS file storage area you have upgraded, clear its read-only attribute, as follows:

1. In a command-line window, change your current directory to the root directory of the file store.
2. Run the following command, which will clear the read-only attribute for all the files in the current directory and in all its subdirectories:

```
attrib -r /s
```

**NOTE** You do not need to wait for the attrib command to complete on one file storage area before running the command on another file storage area. For help with the attrib command, run the command `help attrib`.

## Move File Storage Areas from Windows to UNIX

If you upgraded Content Engine from its Windows-based environment so that it is deployed on a UNIX-based application server, you can now migrate (move) your pre-existing file storage areas from Windows machines to UNIX machines.

This process does not change the path value for the file storage areas that you move.

For example, suppose a file (file storage area or collection) ABC on your upgraded system uses path `/opt/mount00/ABC`, where directory `mount00` is an NFS mount to Windows server folder `C:\Storage`.

To continue using the path `/opt/mount00/ABC` after moving the file to a UNIX machine, you will unmount `/opt/mount00`, copy ABC to its new location on a UNIX machine, and then remount `/opt/mount00` to point to the new location. Thus the file storage areas move to UNIX while their paths stay the same.

The procedure for moving file storage areas uses the following conventions (substitute your own values in their place):

- `/opt` is a partition on a UNIX machine to which you will move file storage areas.
- `/opt/mount01` is a mount point on a Content Engine machine that originally points to a share on a Windows machine where file storage areas and collections are located (a collection is in a subdirectory of a files storage area).

### **To move file storage areas and from Windows to UNIX machines**

---

1. Shut down Content Engine on each application server where it is deployed.
2. Shut down Content Search Engine on each machine where it is installed.
3. On each UNIX machine (which need not be a machine where Content Engine is deployed) to which you will be moving file storage areas, do the following:
  - a. Log on as a user with write access permission for the directory to which you will move file storage areas (**/opt**, for example).
  - b. For each mount point (**/opt/mount01**, for example) on the Content Engine machine that points to file storage areas to be moved to this UNIX machine, do the following:
    - i. Create or designate an existing destination directory for the file storage areas to be moved to this UNIX machine (**/opt/migrated01**, for example).
    - ii. Copy the content of the mount point (**/opt/mount01**) to the destination directory (**/opt/migrated01**).
    - iii. On the Content Engine machine, unmount **/opt/mount01** and remount it to point to **/opt/migrated01**.
    - iv. On the Content Search Engine machine, unmount **/opt/mount01** and remount it to point to **/opt/migrated01**.
4. Start each Content Search Engine.
5. Start each Content Engine.

### ***Uninstall Version 3.5.x of Content Engine***

After upgrading Content Engine (including object stores) to version 4.0.x, you can (optionally) uninstall version 3.5.x on each machine where it is installed, as follows:

#### **To uninstall version 3.5.x of Content Engine**

---

1. Open the Windows Control Panel.
2. Double-click the **Add/Remove Programs** icon.
3. Highlight FileNet Content Engine in the list of currently installed programs and click **Remove** to launch the Content Engine Setup wizard.
4. Click **Yes** to confirm you want to remove the Content Engine installation.

## Complete Content Search Engine Upgrade and Create Collections

After upgrading Content Engine and moving object stores for upgrades, you must create new collections and, if upgrading, remove old index areas.

### To Complete a Content Search Engine Upgrade and Create New or Upgraded Collections

---

1. Launch Enterprise Manager and log in as the GCD Administrator.
2. Access each 3.5.x Index Area and set the status to **Closed**. For details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > View/change index area status](#).

**NOTE** The 3.5.x Index Areas must be closed first and deleted later because they cannot be deleted until after a re-index has completed.

3. Create new index areas. For details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Create Verity index area](#).
4. Reindex to create new collections. For details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Reindex](#).
5. Delete the old 3.5.x index areas. For details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Content Engine Administration > Content-based retrieval > How to... > Remove Verity Index Area](#).

## Task 8a: Upgrade Process Engine (UNIX)

### NOTES

- Verify that you have reconciled the Process Engine user security information. See [“Reconcile the Process Engine user security information.” on page 465.](#)
- IBM recommends that you complete the upgrade of the entire P8 platform and verify functionality before upgrading database software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to upgrade Process Engine but can be re-enabled after the upgrade is complete. See [“To turn off Oracle Password Complexity Verification” on page 521](#) and [“To re-enable Oracle Password Complexity Verification” on page 526](#) for details.
- Determine when to execute pre-installation SQL scripts for Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle in an xterm window.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts” on page 647](#) for detailed information on the scripts and modes of execution.

- Verify that the raw partitions used for the SEC database have been expanded to 64MB.
- Purge the event logs and statistic records in the Process Engine database. Clear the event logs and statistics records that are no longer required for Process Engine tracker items or workflow milestones. Clearing these items reduces the amount of time required to upgrade the Process Engine database.
- If you are using Process Analyzer, verify that all Process Engine events have been transmitted to Process Analyzer and that the Process Analyzer events have been published. If the events have not been transmitted and published, either before purging from the Process Engine database, or before the Process Engine database upgrade, they will not be available to the Process Analyzer.

**WARNING** Do not start the Process Task Manager until specifically told to do so. When you start Task Manager, the database is accessed, which should not be done until the database has been upgraded using the procedures in [“Complete Post-Upgrade Process Engine Configuration” on page 539.](#)

**To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x**

The steps in this procedure apply only to customers who are using Process Analyzer. If you are not using Process Analyzer, proceed to [“To purge event logs and statistics records” on page 520](#). If you are using Process Analyzer you must execute these procedures before proceeding to [“To purge event logs and statistics records” on page 520](#).

The Process Analyzer gets its data from the Process Engine database. All generated events must be transmitted from Process Engine to Process Analyzer before the upgrade. The following steps must be taken to verify that transmission is complete. Some of these steps must be taken on the Process Engine database, some on the Process Engine server using the vwtool utility, and some on the Process Analyzer VMAEDM database.

You will query the Process Analyzer VMAEDM database for a date/time value. Date/time values are stored in the Process Engine and Process Analyzer databases in different formats. When a value has been acquired from the Process Analyzer database, you will use vwtool to convert the value to the appropriate format for the Process Engine database. You will then execute a query on the Process Engine database using the converted date/time value as one parameter in the SQL query.

1. Stop Process Engine applications. These applications include any applications that are generating events or running workflows.
2. Keep both Process Engine and Process Analyzer running until all the events from Process Engine are transmitted to Process Analyzer.
3. Execute the following sub-steps to acquire a date/time value from the Process Analyzer database and convert it to an appropriate format for the Process Engine database.
  - a. On the Process Analyzer VMAEDM database, execute the following SQL query:

```
select InstallDate from X_SchemaInfo
```

This returns a date and time string, such as 09/11/2006 16:23:59. This string must be converted.

- b. On the Process Engine server, start the vwtool utility to convert the Process Analyzer data/time string to Process Engine format.
- c. At the vwtool prompt, type `convert`, as in:

```
<vwtool:26>convert
```

Then press **Enter**. The following choice list displays:

```
t - Time number to string
s - String to time number
e - Error tuple to three part
p - Three part error to error tuple
l - Log event type number to string
i - User id to user name
n - User name to user id
```

To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x

- d. At the Choice? prompt, type the following:

s

and press **Enter**. This converts a string to a time number and returns the following information to indicate what the current date/time mask is, as in:

Current System Mask: mm/dd.yyy hh:tt:ss

Time Mask (CR=system mask):

- e. Press **Enter** to accept the default mask.
- f. When prompted to enter the time string (CR="), type the value that was returned from the SQL query executed on Process Analyzer VMAEDM in step a. (for example, 09/11/2006 16:23:59). Your input must match the format of the current system mask from step d above. Then press **Enter**.

A string value is returned for the date/time entered, which you must make note of for the next set of queries. For example:

Time...[0x4505F00F].....**1158017039** => '09/11/2006 16:23:59'

- g. Get a list of all regions on the disk by typing the following at a vwtool prompt:

regions

- h. When prompted, respond by typing:

d

- i. For every region, type the following at a vwtool prompt:

reg X

where X is the region number

- j. Type the following at the vwtool prompt:

config

- k. Locate and make note of the physical table name associated with every event log.

- l. On the Process Engine database, execute a SQL query (such as the following example) to verify that no untransmitted events remain in the Process Engine database. You will query for the number of records in every physical table associated with event logs, using the names you acquired in the previous step. Following is an example of the query syntax:

```
Select count(*) from f_sw.<physical table name> where F_AEXmitStat = 1 and
F_TimeStamp > <PAInstallDate>
```

where :

The *physical table name* was acquired in step k.

The *PAInstallDate* is the number returned in step f.

Note that the physical table name must be preceded by f\_sw. in the query.

The query result must be 0. If the queries do not return 0, not all events have been collected, in which case Process Engine and Process Analyzer must keep running until all the events are transmitted and the queries return 0.

---

**To verify that all Process Analyzer 3.5.x events have been published**

---

Process Analyzer must publish all the events in its VMAEDM database. Verify that all events have been published through the Microsoft Query Analyzer. Query for the number of rows in the F\_Events table in the VMAEDM database with PAJobId = 0.

The following is an example of the query on the Process Analyzer database.

Using VMAEDM:

```
Select count(*) from F_Events where PAJobId = 0
```

The above query should return 0. If the query returns anything other than 0, then not all events have been published. In that case, you must leave Process Analyzer running until the query returns 0.

---

**To purge event logs and statistics records**

---

Before you upgrade Process Engine, use the vwlog utility to reduce the number of event log and statistics records in the database. This step is optional, but eliminating some of these records can significantly reduce the amount of time necessary for the upgrade to complete. Note that purging these records can take a significant amount of time, so plan this activity accordingly.

The following are some examples of vwlog syntax. Do not use the -P option when you purge logging records if you are using Process Analyzer.

```
vwlog -X -r 100 (this command removes the statistics from isolated region 100)
```

The following command will remove all log records from the isolated region. Use this only if all workflows that have terminated and you no longer need tracking or milestone information.

```
vwlog -L -r 100 (this command removes all log records from isolated region 100)
```

There are multiple optional parameters for the vwlog utility, allowing selection of log records for deletion meeting a number of conditions, such as log records for terminated workflows, for tracker related records, and more. See the IBM FileNet P8 help topics under [FileNet P8 Administration > Process Engine Administration > Administrative tools > vwlog](#) for additional information.

**NOTE** If Process Analyzer is installed, you must complete the steps detailed in [“To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x” on page 518](#) and [“To verify that all Process Analyzer 3.5.x events have been published” on page 520](#) before you purge event logs.

---

**To stop all Process Engine-related services and applications**

---

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 software is shut down:

```
killfnsw -D -A -y -S
```

5. On AIX, execute the following:

```
slibclean
```

#### To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the upgrade process. Turn off this Oracle feature and do not re-enable it until you have upgraded Process Engine.

#### To update Process Engine (UNIX)

---

1. Verify that there is a current system backup.
2. Log on to the server as the root user.
3. Access the Process Engine software package and launch the appropriate **P8PE-4.0.0-  
<operating\_system>.bin/.exe** Setup program
4. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement.

In this screen...	Perform this action...
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation was installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation website. You can use multi-part root folders (for example, <code>/docs/ecm_help</code>) if your application server supports them.</p>
Specify Oracle Version	<p>Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and the server.</p>
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the scripts in an xterm window.</li> <li>• I want to run the scripts silently using operating system authentication.</li> </ul>
Specify the Oracle SYS Password	<p>Enter the Oracle SYS password.</p> <p><b>NOTE</b> This prompt appears only if you indicated that you want to run the scripts with a prompted password.</p>
Choose an Application Server	<p>From the drop-down list, select the application server type and version on which the associated Content Engine servers are deployed.</p>

In this screen...	Perform this action...
Content Engine API Configuration	<p>Configure the Content Engine API, as follows:</p> <ul style="list-style-type: none"> <li>• Transport Method Select WSI from the drop-down list.</li> <li>• Content Engine Client Software URL: Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example: for WebLogic <code>http://hqcomp1:7001/wsi/FNCEWS40DIME</code> for WebSphere <code>http://hqcomp2:9080/wsi/FNCEWS40DIME</code> for JBoss: <code>http://hqcomp3:8080/wsi/FNCEWS40DIME</code> Do not modify the remainder of the string from the default values. <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.</li> </ul>
Please Read the information Below	Review the information on the installation process. Note that you can check the progress of the installation in the <b>/fnsw/local/logs/wizard</b> file. Click <b>Next</b> when you are ready to proceed.
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.
xterm Window	Enter the SYS password for Oracle.  <b>NOTE</b> This window displays if you chose to run the SQL scripts in an xterm window.
Complete the Setup wizard	Click <b>Finish</b> to complete the Process Engine installation.

In this screen...	Perform this action...
The wizard requires that you log out and log back in.	Click <b>Finish</b> , log off and log back in as fnsw.

- To check the progress of the installation, monitor the **/fnsw/local/logs/wizard** file for current activity.
- Check the following log files and correct any errors or failures indicated before proceeding to the next procedure:

Log	Location
Process Engine logs	<b>/fnsw/local/logs</b> and all its subdirectories
IS mini-installer logs	<b>/fnsw/tmp_installer/ MINI_INSTALLER_ERROR.LOG</b>

### To set file permissions

On the Process Engine server, execute `fn_setup -d /` to set file permissions on files and directories installed by Process Engine Setup.

### To edit the `ims_start` file (HP only)

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

Change:

```
nohup /usr/ccs/sbin/dlidd32 2>&1 >/dev/null
```

to

```
nohup /usr/ccs/sbin/dlidd32 +a 0x70000000 2>&1 >/dev/null
```

### To restore any custom modifications for root and fnsw users.

Process Engine Setup 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 Process Engine Setup:

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

## To enable ports

---

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons. By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet uses several ports higher than 32768. See [“IBM FileNet P8 Port Numbers” on page 643](#) for details on which ports IBM FileNet uses.

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

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

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

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

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

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

Enter the following line:

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

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

Enter the following line:

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

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

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

### **To re-enable Oracle Password Complexity Verification**

---

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

Proceed to [“Install Process Engine Software Updates” on page 537](#).

## Task 8b: Upgrade Process Engine (Windows)

### NOTES

- Verify that you have reconciled the Process Engine user security information. See [“Reconcile the Process Engine user security information.” on page 465.](#)
- IBM recommends that you complete the upgrade of the entire P8 platform and verify functionality before upgrading database software.
- If password complexity verification for Oracle databases has been enabled, it must be disabled to upgrade Process Engine but can be re-enabled after the upgrade is complete. See [“To turn off Oracle Password Complexity Verification” on page 532](#) and [“To re-enable Oracle Password Complexity Verification” on page 536](#) for details.
- Determine when to execute pre-installation SQL scripts for Oracle databases. These scripts must be run either:
  - manually, before running the Process Engine Setup program.
  - or
  - automatically, from the Process Engine Setup program, allowing setup to prompt for the sys password for Oracle in an xterm window.
  - or
  - automatically, from the Process Engine Setup program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

See [“Process Engine SQL Scripts” on page 647](#) for detailed information on the scripts and modes of execution.

- Purge the event logs and statistic records in the Process Engine database. Clear the event logs and statistics records that are no longer required for Process Engine tracker items or workflow milestones. Clearing these items reduces the amount of time required to upgrade the Process Engine database.
- If you are using Process Analyzer, verify that all Process Engine events have been transmitted to Process Analyzer and that the Process Analyzer events have been published. If the events have not been transmitted and published, either before purging from the Process Engine database, or before the Process Engine database upgrade, they will not be available to the Process Analyzer.

**WARNING** Do not start the Process Task Manager until specifically told to do so. When you start Task Manager, the database is accessed, which should not be done until the database has been upgraded using the procedures in [“Complete Post-Upgrade Process Engine Configuration” on page 539.](#)

### **To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x**

---

The steps in this procedure and [“To verify that all Process Analyzer 3.5.x events have been published” on page 530](#) apply only to customers who are using Process Analyzer. If you are not using Process Analyzer, proceed to [“To purge event logs and statistics records” on page 530.](#) If you are

To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x

using Process Analyzer you must execute these procedures before proceeding to [“To purge event logs and statistics records” on page 530](#).

The Process Analyzer gets its data from the Process Engine database. All generated events must be transmitted from Process Engine to Process Analyzer before the upgrade. The following steps must be taken to verify that transmission is complete. Some of these steps must be taken on the Process Engine database, some on the Process Engine server using the vwtool utility, and some on the Process Analyzer VMAEDM database.

You will query the Process Analyzer VMAEDM database for a date/time value. Date/time values are stored in the Process Engine and Process Analyzer databases in different formats. When a value has been acquired from the Process Analyzer database, you will use vwtool to convert the value to the appropriate format for the Process Engine database. You will then execute a query on the Process Engine database using the converted date/time value as one parameter in the SQL query.

1. Stop Process Engine applications. These applications include any applications that are generating events or running workflows.
2. Keep both Process Engine and Process Analyzer running until all the events from Process Engine are transmitted to Process Analyzer.
3. Execute the following sub-steps to acquire a date/time value from the Process Analyzer database and convert it to an appropriate format for the Process Engine database.

- a. On the Process Analyzer VMAEDM database, execute the following SQL query:

```
select InstallDate from X_SchemaInfo
```

This returns a date and time string, such as 09/11/2006 16:23:59. This string must be converted.

- b. On the Process Engine server, start the vwtool utility to convert the Process Analyzer date/time string to Process Engine format.
- c. At the vwtool prompt, type `convert`, as in:

```
<vwtool:26>convert
```

Then press **Enter**. The following choice list displays:

```
t - Time number to string
s - String to time number
e - Error tuple to three part
p - Three part error to error tuple
l - Log event type number to string
i - User id to user name
n - User name to user id
```

- d. At the Choice? prompt, type the following:

```
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/yyyy hh:tt:ss
```



To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x

Time Mask (CR=system mask):

- e. Press **Enter** to accept the default mask.
- f. When prompted to enter the time string (CR="), type the value that was returned from the SQL query executed on Process Analyzer VMAEDM in step a. (for example, 09/11/2006 16:23:59). Your input must match the format of the current system mask from step d above. Then press **Enter**.

A string value is returned for the date/time entered, which you must make note of for the next set of queries. For example:

Time...[0x4505F00F].....**1158017039** => '09/11/2006 16:23:59'

- g. Get a list of all regions on the disk by typing the following at a vwtool prompt:

```
regions
```

- h. When prompted, respond by typing:

```
d
```

- i. For every region, type the following at a vwtool prompt:

```
reg X
```

where X is the region number

- j. Type the following at the vwtool prompt:

```
config
```

- k. Locate and make note of the physical table name associated with every event log.
- l. On the Process Engine database, execute a SQL query (such as the following example) to verify that no untransmitted events remain in the Process Engine database. You will query for the number of records in every physical table associated with event logs, using the names you acquired in the previous step. Following is an example of the query syntax:

```
Select count(*) from f_sw.<physical table name> where F_AEXmitStat = 1 and
F_TimeStamp > <PAInstallDate>
```

where :

The *physical table name* was acquired in step k.

The *PAInstallDate* is the number returned in step f.

Note that the physical table name must be preceded by f\_sw. in the query.

The query result must be 0. If the queries do not return 0, not all events have been collected, in which case Process Engine and Process Analyzer must keep running until all the events are transmitted and the queries return 0.

### **To verify that all Process Analyzer 3.5.x events have been published**

---

Process Analyzer must publish all the events in its VMAEDM database. Verify that all events have been published through the Microsoft Query Analyzer. Query for the number of rows in the F\_Events table in the VMAEDM database with PAJobId = 0.

The following is an example of the query on the Process Analyzer database.

Using VMAEDM:

```
Select count(*) from F_Events where PAJobId = 0
```

The above query should return 0. If the query doesn't return 0, not all events have been published. You must leave Process Analyzer running until the query returns 0.

### **To purge event logs and statistics records**

---

Before you upgrade Process Engine, use the vwlog utility to reduce the number of event log and statistics records in the database. This step is optional, but eliminating some of these records can significantly reduce the amount of time necessary for the upgrade to complete. Note that purging these records can take a significant amount of time, so plan this activity accordingly.

The following are some examples of vwlog syntax. Do not use the -P option when you purge logging records if you are using Process Analyzer.

```
vwlog -X -r 100 (this command removes the statistics from isolated region 100)
```

The following command will remove all log records from the isolated region. Use this only if all workflows that have terminated and you no longer need tracking or milestone information.

```
vwlog -L -r 100 (this command removes all log records from isolated region 100)
```

There are multiple optional parameters for the vwlog utility, allowing selection of log records for deletion meeting a number of conditions, such as log records for terminated workflows, for tracker related records, and more. See the IBM FileNet P8 help topics under [FileNet P8 Administration > Process Engine Administration > Administrative tools > vwlog](#) for additional information.

**NOTE** If Process Analyzer is installed, you must complete the steps detailed in [“To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x” on page 527](#) and [“To verify that all Process Analyzer 3.5.x events have been published” on page 530](#) before you purge event logs.

### **To create the Process Engine ODBC data source and test the connection (SQL Server only)**

---

The ODBC data source is required for both local and remote databases.

1. From the Start Menu, navigate to Program > Administrator Tools and launch Data Source (ODBC).
2. Click the **System DSN** tab, then click **Add**.
3. Select **SQL Server** as the driver to use for the new data source and click **Finish**.
4. Enter a name and description for the data source.
5. Choose the SQL Server to connect to from the drop-down list of servers. If only a server name appears in the list, the connection will be with the default instance. If there are named instances

in the database, the name will appear as <server>/<instance name>. Choose the server and instance name being used by the Process Engine 3.5.x software. Click **Next**.

6. Choose SQL Server authentication and check the box to indicate that default settings for additional configuration options should be obtained by connecting to the SQL Server. Indicate the login ID and password to connect with. This login ID does not need to be an administrator, and it is only used to connect to the database to get the default values for the remaining settings to configure the data source.
7. Click **Next**.
8. Change the default database to the Process Engine database being used by the Process Engine 3.5.x software.
9. Clear *Use ANSI null, paddings, and warning* and select *Use ANSI quoted identifiers*. Click **Next**.
10. Select *Perform translation for character data* and click **Finish**.
11. Verify the settings for the data source configuration and click **Test Data Source**. If the test is successful click **OK**. Otherwise resolve the problem before you continue.
12. Double-click **SQL Server** on the **Connection Pooling** tab.
13. Select **Don't pool connection to this driver** and click **OK**.
14. Click **OK** at the ODBC Data Source Administrator window to finish configuration of the data source.

---

**To stop all Process Engine-related services and applications**

1. Log on as fnsw.
2. Set the PPM and any routers to manual startup if they are currently configured to autostart.
3. Stop the following components if they are running:
  - Process Simulator
  - Process Analyzer
  - Custom applications that require a router
  - Component Manager
  - Routers - for Application Engines, Content Engine and custom applications
  - Content Engine
  - Pooled Process Manager (PPM)
  - Process Service
4. Enter the following at a command prompt after the FileNet software is shut down:

```
killfnsw -D -y -S
```

5. In the Windows Services dialog box, stop the following services if they are present and running, and set them to manual startup.
  - IMS ControlService (Process Engine)
  - Process Engine Services Manager (Process Engine and Content Engine)
  - VMAE Publisher service (Process Analyzer)
  - AEEEngine service (Process Analyzer)

### To turn off Oracle Password Complexity Verification

---

Process Engine does not support Oracle Password Complexity Verification during the upgrade process. Turn off this Oracle feature and do not re-enable it until you have upgraded Process Engine.

### To reset the fnsw password to the default value

---

The fnsw user's password must be reset to BPMtemp1pzwd before running Process Engine setup for the upgrade. After you change the password for the fnsw user, you must also use the Windows Services tool to update the Log On tab for the IMS ControlService and the Process Engine Services Manager accordingly. If the password is not changed before the upgrade, the upgrade will fail. After successfully upgrading Process Engine, reset the password again.

### To update Process Engine (Windows)

---

1. Verify that there is a current system backup.
2. Log on as fnsw.

**WARNING** Setup will run a number of steps but to complete the installation the computer must be restarted. If you are installing on a Windows 2003 R2 server, you must create an empty file named **fn\_newinit** in the **c:\temp** directory before restarting the computer.

3. Access the Process Engine software package and launch the appropriate **P8PE-4.0.0-Win.exe** Setup program.

**NOTE** If you intend to run this installation from a remote server, either interactively or silently (that is, via a command script), be aware that Process Engine Setup has a 64-character path limitation and an 8.3 filename limitation. This limitation applies to the IMSInst subdirectory where the IS mini-installer setup.exe is located. For example, the path where the IS mini-installer resides:

**\\englibcm\working\Disks\ProcessEngineDisks\pe400.182\p8\ProcessEngine.Windows\IMSInst**

when expressed in 8+3 format is:

**\\englibcm\working\disks\PROCES~1\pe400.182\p8\PROCES~1.WIN\IMSInst**

This is 67 characters long.

4. Complete the Process Engine Setup screens, as follows:

In this screen...	Perform this action...
Welcome to Process Engine Setup	Click <b>Next</b> on the Welcome screen to proceed with the installation.
License Agreement	Review and accept the license agreement.
(Running BPM Components Detected)	An untitled screen might come up indicating which, if any, BPM software components have been detected. Click <b>Next</b> to stop the software.
Specify the Documentation URL	<p>Enter the Documentation URL, which is where the IBM FileNet P8 Platform Documentation was installed. Your entry must be in the following format:</p> <p><code>http://&lt;docserver:port#&gt;/&lt;ecm_help&gt;</code></p> <p>where:</p> <p><i>docserver</i> is the name of the Java web server.  <i>port#</i> is the port number.  <i>ecm_help</i> is the root folder of the documentation website. You can use multi-part root folders (for example, <b>/docs/ecm_help</b>) if your application server supports them.</p>
Specify Execution Mode for Oracle Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the Oracle SYS user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually.</li> <li>• I want to run the SQL scripts with a prompted password.</li> <li>• I want Setup to run the SQL scripts silently using operating system authentication.</li> </ul>
Specify the Oracle SYS Password	Enter the Oracle SYS password.
<p><b>NOTE</b> This prompt appears only if you indicated that you want to run the scripts with a prompted password.</p>	

In this screen...	Perform this action...
Specify Execution Mode for SQL Server Scripts	<p>A series of SQL scripts must be executed. You could have already run the scripts manually before starting Process Engine Setup. If you did not run them manually, you need to indicate how you want to run them now. You can run them as the SQL Server Administrator sa user or as another user who can be authenticated through the operating system.</p> <ul style="list-style-type: none"> <li>• I have already run the pre-install scripts manually</li> <li>• I want to run the SQL scripts with a prompted password</li> <li>• I want Setup to run the SQL scripts silently using operating system authentication</li> </ul> <p>Validation of the SQL Server connection will also occur if you choose to run the scripts now, either with a prompted password or using operating system authentication. No validation of the database connection will be done if you indicate that you have already run the SQL scripts manually.</p>
Specify the SQL Server System Administrator Password  <b>NOTE</b> This prompt only appears if you indicated that you want to run the SQL scripts with a prompted password.	Enter the SQL Server sa password.
Specify SQL Server Config Parameters	Enter the ODBC Data Source name created in <a href="#">“To create the Process Engine ODBC data source and test the connection (SQL Server only)”</a> on page 530.
Specify Oracle Version	<p>Oracle9i or Oracle 10g versions are supported. Indicate which version of Oracle software to use on the database server for Process Engine.</p> <p><b>NOTE</b> The Oracle versions must be the same on the client and the server.</p>
Specify SQL Server Version	Indicate whether you will be using SQL Server 2000 or SQL Server 2005 for the Process Engine database.
Choose an Application Server	From the drop-down list, select the application server type and version on which the associated Content Engine servers are deployed.

In this screen...	Perform this action...
Content Engine API Configuration	<p>Configure the Content Engine API, as follows:</p> <ul style="list-style-type: none"> <li>Transport Method Select WSI from the drop-down list.</li> <li>Content Engine Client Software URL: Replace the sample server name and port number (<i>CEserver.example.com:7001</i>) with the host name of the Content Engine application server to which Process Engine will connect. The port number depends on the application server type. For example: for WebLogic <code>http://hqcomp1:7001/wsi/FNCEWS40DIME</code> for WebSphere <code>http://hqcomp2:9080/wsi/FNCEWS40DIME</code> for JBoss: <code>http://hqcomp3:8080/wsi/FNCEWS40DIME</code> Do not modify the remainder of the string from the default values. <b>NOTE</b> To change the Content Engine server host name later, or to connect to a different Content Engine server, edit the <b>WcmApiConfig.properties</b> file. For information, see the IBM FileNet P8 help topic <a href="#">FileNet P8 Administration &gt; Application Engine Administration &gt; Key configuration files and logs</a>.</li> </ul>
Please Read the Summary Information Below	Verify your selections, and click <b>Install</b> to install Process Engine.
Select Software Components to Start	Select the BPM software components to start when installation is complete. You might not want to automatically start the Process Engine Services Manager if you intend to change the fnsw password and make corresponding changes to the services.
Complete the Setup	Click <b>Finish</b> to complete the Process Engine installation.
Specify a system reboot now	Choose the option to restart the computer now and click <b>Finish</b> . <b>WARNING</b> If you are installing on a Windows 2003 R2 server, you must create an empty file named <b>fn_newinit</b> in the <b>c:\temp</b> directory before restarting the computer.

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

Log	Location
Process Engine logs	C:\Program Files\FileNet\PE\PE400_setup.log
IS mini-installer logs	<Windir>/mini_installer.log
Output files from SQL Script validation (only if the SQL scripts were executed from Process Engine Setup)	C:\Program Files\FileNet\PE

6. Start the following services:
- IMS ControlService
  - Process Engine Services Manager
7. Reset the fnsw user's password. After you change the password for the fnsw user, you must also use the Windows Services tool to update the Log On tab for the IMS ControlService and the Process Engine Services Manager accordingly.

**To re-enable Oracle Password Complexity Verification**

---

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

Proceed to ["Install Process Engine Software Updates"](#) on page 537.



## Task 9: Install Process Engine Software Updates

Install any service packs, fix packs and/or interim fixes required for Process Engine.

### To install the Content 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 service representative.
2. Open the readmes for the following software updates and perform the installation procedures provided:
  - a. Process Engine 4.0.2 Service Pack
  - b. Any subsequent fix pack (P8PE-4.0.2-001 or later)
  - c. Any subsequent interim fixes (typically optional)

## Task 10: Install the Latest Content Engine Client Files on Process Engine Servers

Install any Content Engine client file updates that are available.

### **To install the Content Engine client files**

---

1. To download the latest software updates, and to determine which of these updates may be required for use with other components and expansion products, contact your service representative.
2. Open the readme for the P8CE-4.0.0-002 (or later) Fix Pack and perform the installation procedures provided for the Content Engine Java Client files.

## Task 11: Complete Post-Upgrade Process Engine Configuration

You must perform the following additional procedures to complete the upgrade of all Process Engine data and objects.

Do not start the procedures in this task unless you have installed Process Engine 4.0.2 (or later) Service Pack.

The format of the event log files changes in the 4.0 release. As a part of this upgrade, several SQL scripts must be executed for SQL Server and DB2 databases only.

**WARNING** Throughout this procedure there are multiple software restarts. Execute all restarts as documented. Do not start or restart Process Task Manager or other IBM FileNet software unless specifically told to do so. Starting Process Task Manager accesses the database, which should not be done until the database has been upgraded.

To execute steps on a UNIX operating system, the terminal must support X Windows and the DISPLAY environment variable must be set.

SQL Server Client software is required on the Process Engine server to execute a number of SQL scripts documented in this topic if the database is a remote SQL Server. The SQL Server Client software can be removed from the Process Engine server after the Process Engine database has been successfully upgraded to the 4.0 release.

### To update the Process Engine database objects

After you have updated the Process Engine software, you must update the Process Engine database objects.

1. Log on as fnsw on UNIX or a local administrator on Windows.
2. (Windows only) Ensure that the Process Engine services are started:
  - IMS ControlService
  - Process Engine Services Manager
3. (Windows only) Enable the redirection of log messages to the Image Services error log. This redirection logs messages to the Image Services error log as well as to the default Windows Event Log. When you enable this redirection, you can monitor the progress of the database object upgrade in a command window.

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

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

4. Restart the Process Engine software on Windows and UNIX platforms as follows:

At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

5. If you are using a SQL Server database, proceed to step 6. If you are using a DB2 database, proceed to step 7. If you are using an Oracle database, proceed to step 8.
6. (SQL Server only) Edit and run the **\fnsw\mssql\vwssql35to40\_pre1.bat** file on the database server. Database schema changes will be made to existing event log database tables.

- a. Save the file to the same directory as **vwssql35to40\_pre1a.bat**.
- b. Change the values in the file as appropriate for your system. The content of the **vwssql35to40\_pre1.bat** file looks like this:

```
osql /U sa /P /n /d VWdb /h-1 /i vwssql35to40_pre1.sql
```

Change the values for your system to:

```
osql /D <DSN> /U <sa> /P <sa> /n /d <VWdb> /h-1 /i vwssql35to40_pre1.sql  
/o pre1a.log
```

where:

/D indicates the following variable is your ODBC data source name (DSN)

/U indicates the following variable is the administrator user name in the Process Engine database

/P indicates the following variable is the administrator user's password in the Process Engine database

/d indicates the following variable is the Process Engine database name

Optionally, you can add an output file /o **pre1a.log**. Otherwise, all output goes only to the screen.

- c. Run the **vwssql35to40\_pre1a.bat** file.
- d. Proceed to step 8.
7. (DB2 only) Edit and run the **\fnsw\DB2\vwdb2\_35to40\_pre1.bat** file on Windows or **\fnsw\DB2\vwdb2\_35to40\_pre1.sh** on UNIX platforms. Database schema changes will be made to existing event log database tables.

- a. Copy the file to the same directory as **vwdb2\_35to40\_pre1a.bat** or **vwdb2\_35to40\_pre1a.sh**.
- b. Change the values in the file as appropriate for your system.

Change the values for your system to:

```
db2 connect to <database_name> user <PE runtime user> using < password>
```

where:

<database\_name> is your Process Engine DB2 database name

<PE runtime user> is the Process Engine runtime user (f\_sw)

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

- c. Run the **vwdb2\_35to40\_pre1a.bat** under from the DB2 command line processor on Windows, or run **vwdb2\_35to40\_pre1a.sh** from a command prompt on UNIX.

d. Proceed to step 8.

8. Initiate the database Process Engine schema changes by executing the following command:

```
vwtool
```

Choose Yes when a message is presented indicating that an upgrade is required.

You are prompted to initiate tracing to capture the changes made to a trace file. Reply to the prompt as desired. IBM recommends capturing trace information on upgrades of development systems and on production systems if there are failures during this update.

When vwtool starts, it automatically checks the Process Engine database level, updates the schema accordingly, and creates two additional scripts.

For SQL Server:

```
fnsw\mssql\vwssql35to40_post1.sql
```

```
fnsw\mssql\vwssql35to40_post2.sql
```

For DB2:

```
fnsw\DB2\vwdb2_35to40_post1.sql
```

```
fnsw\DB2\vwmdb2_35to40_post2.sql
```

**WARNING** If the upgrade fails at this point, you must restore the Process Engine database backup.

See the IS error log to monitor the progress of the updates and ensure that no errors occur.

Check the logs to verify that messages similar to the following are captured:

```
2006/10/17 16:23:43.261 <fnsw> VW/Process (14952) ... [INFO]
```

```
VW: Database upgrade successful to version 46, please follow instructions to perform the next step.
```

```
2006/10/17 16:23:43.303 <fnsw> VW/Process (14952) ... [INFO]
```

```
VW: Must restart software to complete upgrade procedure
```

Ignore messages designated as SERIOUS if they are in combination with a successful message for that upgrade, especially if all process IDs are the same for all the errors and INFO messages.

9. Exit vwtool when you get a message that the procedure is complete.
10. Do a backup of the Process Engine database. This backup can serve as a checkpoint, should an error occur later that requires a database restore.
11. Restart the Process Engine software. At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

To update the Process Engine security

Before updating the Process Engine security, ensure:

- Your directory server is up and correctly configured.
- Content Engine 4.0 is up.

**NOTE** When the Process Task Manager starts, a message will be presented indicating that routers must be migrated. Routers will be migrated in [“To migrate routers and update isolated regions” on page 543](#) as a part of completing Process Engine configuration changes. The message can be ignored now.

1. Update the information on the Process Task Manager Security / General tab.

a. Launch Process Task Manager, as follows:

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

UNIX: Enter `vwtaskman` at the command prompt.

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 click the **Security** tab to configure the General settings.

Provide a service username and password, an administrator group name, and an optional configuration group name. See the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Process Engine > Configure the Process Engine > Security](#) for details on the user and groups.

**NOTE** The service username should be entered as a short name, not a distinguished name.

d. Click **Apply**.

**NOTE** If you get an error applying security settings, click **Close** on the message, correct the problem if noted, and repeat step 1.

e. Click OK to close the dialog indicating that you must run vwtool.

f. Exit Process Task Manager.

2. Run vwtool and choose **Yes** when a message is presented indicating that an upgrade is required.

This step moves all existing user environment records from the 3.5.x format to the 4.0.0 format.

a. IBM recommends that you reply Yes to the prompt to initiate tracing to capture the changes made to a trace file.

**WARNING** 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 un-migrated users are no longer valid Process Engine users - that is, they have no Process Engine work, then you can override these errors and complete this

part of the upgrade. Once you choose to override these errors, there is no way to recover the user environment records for any users not migrated to 4.0.0. Any Process Engine work for un-migrated users is lost. Choose Yes to ignore errors and force the completion of the upgrade only after you have resolved any outstanding problems.

As user environment records are moved, information is logged to the Image Services errorlog. When vwtool finishes, it will display messages on the screen indicating whether or not the migration was successful, how many users were migrated, and how many users were not migrated.

If all users did not successfully migrate to the 4.0.0 format, you will need to look at the messages in the errorlog and resolve the problems. Examples of the types of resolution required could include the need to fix a problem with the configuration of the Content Engine and its application server's access to the directory server or the need to create users in the underlying LDAP store.

- b. Address any errors that occurred execute vwtool again. You might be prompted to ignore issues related to the user environment record upgrades.
3. Check the errorlog to verify that the database version number has been updated to 49. This update will happen only after either successful migration of all environment records or all errors have been intentionally overridden.
4. Restart the Process Engine software, as follows:

At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

### **To migrate routers and update isolated regions**

---

You must convert all routers to connection points, and assign passwords to any existing isolated regions, using the following procedure.

1. On the Process Engine and Application Engine servers, use Process Task Manager to view the general properties of each Process Router. For each, make a note of the following properties:
  - Process Router name
  - Process Engine
  - Isolated region
2. Delete each Process Router. You can do this immediately or wait until the corresponding connection points are created. The Process Routers displayed in Process Task Manager are not used and have no effect on the system other than to cause a warning message on Process Task Manager startup. Once all Process Routers have been deleted and Process Task Manager has been restarted, the Process Router node no longer appears.
3. To delete a Process Router, do the following:
  - a. Select the Process Router you want to delete.
  - b. Select Delete from the Action menu.
4. Start Enterprise Manager 4.0.0.

5. Navigate to the PE Region IDs node and start the wizard.
6. Specify an isolated region for each unique Process Engine / isolated region combination. (If you have more than one Process Router pointing to the same Process Engine / isolated region combination, you will identify only one Process Engine Region ID.)

For this Process Engine Region ID property...	Use the value from this Process Router property...
DNS name or IP address of the Process Engine	Process Engine
Region number	Isolated region

7. Assign a password for each region as you create it. Make note of the password you assign. You will need to enter that password in the steps that follow for assigning the password to regions in Process Task Manager. The passwords must match.
8. Navigate to the PE Connection Points node, start the wizard, and create new connection points for each region.
9. Close Enterprise Manager.
10. Start Process Task Manager on the Process Engine server.

(Windows) Navigate to Start > Programs > FileNet P8 Platform > Process Engine > Process Task Manager.

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

11. Select Process Engine in the feature pane.
12. Click the **Security** tab, then **Region Passwords** tab and the **Add** icon on the right-most side of Process Task Manager to add an isolated region and password. Click **Apply**.

**NOTE** The password must match the password that you entered when creating a Process Engine Region in [Step 7 on page 544](#).

13. After all parameters have been entered, click **Apply** and restart the Process Service when prompted.

### To update email notification

---

If you're using email notification, add a language pack for the Default Authoring Locale. Unless this is done, email notification will be disabled. Verify that the Default Authoring Locale is correct (it has defaulted to the operating system's locale). See the IBM FileNet P8 help topic [FileNet P8 Administration > Process Engine Administration > Workflow administration tasks > Coordinating workflow design > Enable Email notification](#) for information on adding a language pack for the Default Authoring Locale.



### To update every isolated region

---

Execute a transfer of the upgrade.cdl file in every isolated. Take the following steps.

1. Restart the Process Engine software, as follows:

At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

2. Start vwtool and do the following substeps to update each existing isolated region and recreate views for the existing Process Engine tables (logs, rosters, and queues):

- a. Get a list of all regions on the disk by typing the following at a vwtool prompt:

```
regions
```

Make note of all region numbers.

- b. When prompted, respond by typing:

```
d
```

- c. Exit from vwtool.

- d. Change directories to the location of the upgrade.cdl file.

```
/fnsw/local/sd (UNIX)
```

```
\fnsw_loc\sd (Windows)
```

- e. Initiate a transfer on every working isolated region by entering the following command.

```
vwtfcr -o upgrade.cdl -r X
```

where *X* is the isolated region number

At the prompt, log on as a user who is a member of the PEAdministrators group.

### To run post-upgrade scripts (SQL Server and DB2 only, optional)

---

Two post-upgrade scripts can optionally be executed on Process Engine SQL Server and DB2 databases. The first script copies records from archived event log tables to the new 4.0 version of the tables. The second script will delete the archived log tables, saving database space.

For SQL Server database, execute steps 1 and 2. For DB2 databases, execute steps 3 and 4. For Oracle databases, proceed to [“To backup the database and restart software” on page 547](#).

1. (SQL Server only) Edit and run the **lfnsw\mssql\vwssql35to40\_post1.bat** file.
  - a. Save the file to the same directory as **vwssql35to40\_post1a.bat**.
  - b. Change the values in the file as appropriate for your system. The contents of the **vwssql35to40\_post1.sql** file looks like this:

```
osql /U sa /P /n /d VWdb /h-1 /i vwssql35to40_post1.sql
```

Change the values for your system to:

```
osql /D <DSN> /U <sa> /P <sa> /n /d <VWdb> /h-1 /i vwssql35to40_post1.sql /o post1a.log
```

where:

/D indicates the following variable is your ODBC data source name (DSN)

/U indicates the following variable is the administrator user name in the Process Engine database

/P indicates the following variable is the administrator user's password in the Process Engine database

/d indicates the following variable is the Process Engine database name

/o indicates the following variable is the optional post1a.log output file. If you choose to eliminate this entry, output displays to the screen only.

- c. Run the **vwssql35to40\_post1a.bat** file.

2. (SQL Server only) Delete archived event log tables by editing and running **vwmmql35to40\_post2.bat**.

- a. Change the values in the file as appropriate for your system. The contents of the **vwssql35to40\_post2a.sql** file looks like this:

```
osql /U sa /P /n /d VWdb /h-1 /i vwssql35to40_post2.sql
```

Change the values for your system to:

```
osql /D <DSN> /U <sa> /P <sa> /n /d <VWdb> /h-1 /i vwssql35to40_post2.sql /o post2a.log
```

where:

/D indicates the following variable is your ODBC data source name (DSN)

/U indicates the following variable is the administrator user name in the Process Engine database

/P indicates the following variable is the administrator user's password in the Process Engine database

/d indicates the following variable is the Process Engine database name

/o indicates the following variable is the optional post2a.log output file. If you choose to eliminate this entry, output displays to the screen only.

- b. Run the **vwmssql35to40\_post2a.bat** file.
- c. Proceed to step 5.
3. (DB2 only) Connect to the DB2 database and run the **\fnsw\DB2\vwdb2\_35to40\_post1.sql** file.
  - a. Start the DB2 command line processor and logon to the Process Engine database with the Process Engine runtime user (f\_sw).
  - b. Enter the following command in that window.

```
db2 connect to <database_name> user <PE runtime user> using < password>
```

where:

<database\_name> is your Process Engine DB2 database name

<PE runtime user> is the Process Engine runtime user (f\_sw)

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

- c. Run the **vwdb2\_35to40\_post1.sql** file by executing the following in the DB2 command line processor.

```
db2 -tvf vwdb2_35to40_post1.sql
```

4. (DB2 only) Connect to the DB2 database and run the **\fnsw\DB2\vwmdb2\_35to40\_post2.sql** file.
  - a. Start the DB2 command line processor and logon to the Process Engine database with the Process Engine runtime user (f\_sw).

```
db2 connect to <database_name> user <PE runtime user> using < password>
```

where:

<database\_name> is your Process Engine DB2 database name

<PE runtime user> is the Process Engine runtime user (f\_sw)

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

- b. Run the **vwmdb2\_35to40\_post2.sql** file by executing the following in the DB2 command line processor.

```
db2 -tvf vwdb2_35to40_post2.sql
```

- c. Proceed to [“To backup the database and restart software” on page 547.](#)

### **To backup the database and restart software**

---

5. Back up the Process Engine database. While this backup is not required, it is recommended to provide a checkpoint that can be used later if a restore is needed in the context of this upgrade.
6. (Windows only) Start the following services and set them back to automatic startup.
  - IMS ControlService
  - Process Engine Services Manager

7. (Windows only) Disable the redirection of log messages to the Windows Event Log. Change the LogToFiles value from 1 to 0 for the following registry key:

HKEY\_LOCAL\_MACHINE > SOFTWARE > FileNET > IMS > CurrentVersion

8. Restart the Process Engine software, as follows:

At a Windows command prompt, or UNIX command line, type the following:

```
initfnsw -y restart
```

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

---

Execute the following steps to configure the largest available contiguous free memory block. If you fail to perform this procedure, the system will not allocate shared memory at some point during normal execution and will cease to function correctly.

1. Start vwtool at a command prompt.
2. Log on using the service username you provided when completing the steps in [“To update the Process Engine security” on page 542](#).
3. Use the processmap command to find the largest contiguous free memory area, as in:

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

vwtool returns the following:

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

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

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

Address	Attrib	Size	Owner
=====	=====	=====	=====
00000000	Free	65536	
00010000	Private	12288	
00013000	Free	53248	
00020000	Private	4096	

(pages of memory addresses omitted here)

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

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

```
Largest FREE block found : 453873664 bytes at address 0x4B577000  

Rounded up to a 64K boundary, free block address 0x4B580000
```

In the example above, 0x4B580000 is the address we want. In some cases you might only see the line referencing the largest free block because the value is already at a 64K boundary.

4. Run the Windows regedit command to create a key for IS StartShmAddress, using the address found in step 2. Run regedit on the Process Engine server and go to:

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

5. Create a new DWORD value. Name it:

```
StartShmAddress
```

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

```
Value name = StartShmAddress
```

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

From the example above the value will be 4B580000.

Base is hexadecimal.

7. Click **OK**.
8. Exit from regedit.
9. Restart the Process Engine software.
10. Verify the setting by executing the following at a command prompt:

```
ipc_tool -A
```

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

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

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

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

Shared Memory Address Manager Information

Address	Shm id	Creator
---------	--------	---------

Enter <space> to continue, 'q' to quit:

0	0x4b580000	FNSHM_464d0000	Shared address manager
---	------------	----------------	------------------------

**NOTE** The First shared memory address is 0x4B580000, the value from this example.

1	0x4c580000	FNSHM_464a0000	FileNet server software
---	------------	----------------	-------------------------

Total Image Services shared memory allocated: 16 MB

(This does not include segment #0)

11. Exit `ipc_tool`. If the shared memory address is correct, proceed to the next installation task. If the value is not correct, verify steps 1 - 8 above before proceeding.

## Task 12: Upgrade Application Engine

This task includes Application Engine upgrade instructions for WebSphere, WebLogic, and JBoss (UNIX and Windows).

### Before you upgrade Application Engine

---

- Verify that Content Engine and Process Engine have been upgraded.
- Review the Application Engine details of the [“Upgrade Overview” on page 460](#).
- Review the steps needed to retain the AE 3.5.x configuration.

You will need these if, for any reason, you need to back out of this installation.

- Verify that you have recorded *all necessary settings* as outlined in [“Upgrade Checklists” on page 470](#).

(Best practices for backup) If you want to retain your existing Application Engine settings you must record all necessary settings before you start the upgrade installation.

- As part of the upgrade you will create backup copies of all important 3.5.x configuration files in [“Backup, undeploy, and remove the Workplace web application from the J2EE application server.” on page 552](#).
  - During the upgrade, a number of existing configuration files will be moved to the newly created `<AE_install_path>Config\AE` directory. See Table 1, “Configuration files that will be moved,” on page 554.
  - As part of the upgrade, the installation program automatically creates a backup of your existing AE 3.5.x configuration files, appending the suffix `.old` to the filenames. See Table 2, “Configuration files that will be backed up,” on page 554.
  - The existing version 3.5.x **Actions.xml** and **web.xml** files will be merged with the 4.0.x versions during the upgrade installation.
  - In addition, all comments added to the **Actions.xml** file will be lost during the merge.
- If you have the IBM FileNet P8 eForms expansion product installed, uninstall it.

See the “Removing Software” topic in the *IBM FileNet P8 eForms Installation Guide*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### To upgrade Application Engine

---

1. Log on to the Application Engine server:

UNIX - log on as root.

Windows - log on as a member of the local Administrators group or a user with equivalent permissions.

2. Verify that the Process Router is stopped.

**NOTE** Although the AE router has been deleted as part of the PE upgrade, a local, disconnected instance might be running on your server.

a. Launch the Process Task Manager from **<AE\_install\_path>/Router**.

- UNIX:

```
./routercmd.sh
```

- Windows:

```
routercmd.bat
```

b. Stop the router, if running.

**NOTE** This is the router configured and started as part of the FileNet P8 Platform 3.5.x installation. For more information, see the *FileNet P8 Platform 3.5.x Installation and Upgrade Guide* task “Start the Process Router.” To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21.](#)

c. Exit the Process Task Manager.

3. (Windows only) Using the Windows Task Manager, verify that no javaw.exe processes are running.

If the applications in the preceding steps stopped correctly, no Application Engine related javaw.exe processes should be running. If needed, use Windows Task Manager or some other process management tool to stop all javaw.exe processes with the directory path **<AE\_install\_path>\\_jvm\bin**.

4. Backup, undeploy, and remove the Workplace web application from the J2EE application server.

During this step you will create a backup copy of the deployed web application containing all customized files and all configuration files stored in the **WEB-INF** directory.

**NOTE** Even though the installer automatically creates a backup of your existing AE 3.5.x configuration files IBM recommends backing up the directories below to get a complete backup of your deployed system. You might need to use these files in [Step 5 “Copy modified files to the installed Workplace directory.” on page 553](#) and [Step 6 “Upgrade and configure the Application Engine software.” on page 554](#) below to retain your 3.5.x settings.

- WebSphere 5.x

i. Make a backup copy of the deployed Workplace directory (**<deploy\_backup>**):

```
<WAS_HOME>/installedApps/<node_name>/app_engine_war.ear/app_engine.war
```

ii. Stop the Workplace application from the admin console.

iii. Uninstall the Workplace application from **Enterprise Applications**.

iv. Save the changes and stop the WebSphere server.

v. Delete the temp Workplace directory (default: **app\_engine\_war**) from:

```
<WAS_HOME>\WebSphere\AppServer\temp\<Node_name>\<Server>\app_engine_war
```



- WebSphere 6.x
    - i. 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**
    - ii. Stop the Workplace application from the admin console.
    - iii. Uninstall the Workplace application from **Enterprise Applications**.
    - iv. Save the changes and stop the WebSphere server.
    - v. Delete the temp Workplace directory (default: **app\_engine\_war**) from:  
**<WAS\_HOME>/profiles/default/temp/<node\_name>/<instance>/app\_engine.war**
  - WebLogic
    - i. From the WebLogic Administration Console, stop the Workplace web application module.
    - ii. Make a backup copy of the Workplace folder (**<deploy\_backup>**):  
**<AE\_install\_path>/Workplace**
    - iii. Delete the Workplace Web Application Module.
    - iv. Stop the WebLogic server.
  - JBoss
    - i. Stop the JBoss server.
    - ii. Make a backup copy of the deployed Workplace directory (**<deploy\_backup>**):  
**<JBoss\_HOME>/server/default/deploy/Workplace.war**
    - iii. Delete the deployed Workplace directory:  
**<JBoss\_HOME>/server/default/deploy/Workplace.war**
    - iv. Delete the temporary Workplace directory for JBoss:  
**<JBoss\_HOME>\server\default\work\jboss.web\localhost\Workplace**
5. Copy modified files to the installed Workplace directory.

During the upgrade, the configuration files listed in [Table 1 on page 554](#) will be moved from the following directory:

**<AE\_install\_path>/Workplace/WEB-INF**

to

**<AE\_install\_path>/Config/AE**

**NOTE** If you have made modifications to any of these files directly in the deployed Workplace directory (**app\_engine.war** for WebSphere or **Workplace.war** for JBoss) in your FileNet P8 3.5.x environment you must copy modified version of these files from the **<deploy\_backup>**

directory to the installed directory **<AE\_install\_path>/Workplace/WEB-INF** before you run the installer. :

**Table 1: Configuration files that will be moved**

- |                                |                             |
|--------------------------------|-----------------------------|
| • actions.xml                  | • icons.properties          |
| • bootstrap.properties         | • InfoPages.xml             |
| • ClassFilter.xml              | • PagingConfiguration.xml   |
| • ConfigurableLabels.xml       | • PolicyProcessors.xml      |
| • containericons.properties    | • PrimaryViews.xml          |
| • content_redir.properties     | • PropertiesPages.xml       |
| • customobjecticons.properties | • SimpleSearch.xml          |
| • download_redir.properties    | • SystemsPropertiesView.xml |
| • fnsoap.xml                   |                             |

**NOTE** As part of the upgrade the installation program automatically creates a backup of your existing AE 3.5.x configuration files, appending the suffix .old to the filenames.

**Table 2: Configuration files that will be backed up**

- |                                 |                                      |
|---------------------------------|--------------------------------------|
| • actions.xml.old               | • UpdateActions.xml.old              |
| • ClassFilter.xml.old           | • UpdateClassFilter.xml.old          |
| • ConfigurableLabels.xml.old    | • UpdateConfigurableLabels.xml.old   |
| • containericons.properties.old | • UpdateProps.xml.old                |
| • content_redir.properties.old  | • UpdateSystemPropertiesView.xml.old |
| • icons.properties.old          | • UpdateWeb.xml.old                  |
| • SystemsPropertiesView.xml.old | • web.xml.old                        |

The backed up files are located in:

**<AE\_install\_path>/backup-4\_0\_1**

- Upgrade and configure the Application Engine software.

**NOTE** As part of this procedure you might be required to install software updates and client files for CE and PE. You do NOT need to deploy/redeploy the Application Engine web application (Workplace) as part of those installations. For this upgrade you will deploy the web application, as part of [Step iii “Deploy Application Engine.” on page 555](#) below, after you have completed all upgrade tasks.

a. Install and configure the Application Engine.

i. Do [“Install Application Engine” on page 344.](#)

**NOTE** The default install path for version 3.5.x is different than the default path for version 4.0.x. During the upgrade, your version 3.5.x install path will be used. The default install path for version 3.5.x is:

- UNIX - **`/opt/FileNet/`**
- Windows - **`C:\Program Files\FileNet\`**

ii. Configure Application Engine.

Do one of the following:

- [“Configure Application Engine \(WebSphere\)” on page 356.](#)
- [“Configure Application Engine \(WebLogic\)” on page 373.](#)
- [“Configure Application Engine \(JBoss\)” on page 381.](#)

b. Manually copy custom data.

If your custom data was *not* retained in the upgrade you must manually copy the data from the version 3.5.x backup files (**`<deploy_backup>`**) to the upgraded files.

**NOTE** You must also copy any data you want to retain if you have made any custom additions 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 above.

c. Install required updates and deploy Application Engine.

i. Do [“Install Application Engine Software Updates” on page 383.](#)

**CAUTION** Do NOT deploy/redeploy the application as part of this step. You will deploy the Application Engine web application (Workplace) as part of [Step iii “Deploy Application Engine.” on page 555](#) below.

ii. Do [“Install the Latest Content Engine and Process Engine Client Files on Application Engine Servers” on page 384.](#)

**CAUTION** Do NOT deploy/redeploy the application as part of this step. You will deploy the Application Engine web application (Workplace) as part of [Step iii “Deploy Application Engine.” on page 555](#) below.

iii. Deploy Application Engine.

Do one of the following:

- [“Deploy Application Engine \(WebSphere\)” on page 385](#)
- [“Deploy Application Engine \(WebLogic\)” on page 391](#)
- [“Deploy Application Engine \(JBoss\)” on page 394](#)

7. Continue with [“Complete Post-Upgrade Application Engine Configuration” on page 556.](#)

# Task 13: Complete Post-Upgrade Application Engine Configuration

## To complete post-upgrade Application Engine configuration

---

1. Edit the **taskman.login.config** file.
  - a. Make a backup copy of **taskman.login.config**.  
This file is located at:  
**<AE\_install\_path>/Router/taskman.login.config**
  - b. Edit **taskman.login.config**.  
Locate the following section:

```
FileNetP8
{
  weblogic.security.auth.login.UsernamePasswordLoginModule required debug=false;
};
```

Replace it with this section:

```
FileNetP8
{
  com.filenet.api.util.WSILoginModule required debug=false;
};
```
  - c. Save and close the file.
2. (WebLogic only) If you upgraded from WebLogic 8.1.x to WebLogic 9.x, you must set the case-sensitivity for the server.
  - a. From the WebLogic Administration Console, click **<my\_domain>**.
  - b. Click **Lock and Edit**.
  - c. Select the **Security** tab.
  - d. Click **Advanced**.
  - e. Set **Web App Files Case Insentive** to false.
  - f. Click **Save**.
  - g. Click **Activate Changes**.
  - h. Close the administration console.
3. Sign in to Workplace to test your connection.
  - a. On any computer, open a browser and type:  
`http://<ApplicationEngineServerName>:<port#>/Workplace`
  - b. Enter a user name and password, and click **Sign in**.

- c. If the Bootstrap Preferences page is displayed, you must set the preferences.

Follow the instructions in [“Set Application Engine Bootstrap Preferences” on page 403](#) to reset your bootstrap properties. Use the notes you made of your bootstrap settings in [“Application Engine” on page 476](#) to complete this step.

4. Set the Process Engine Connection Point.

The Process Router from version 3.5.x has been replaced with the Process Engine Connection Point.

- a. In Workplace click **Admin**.
- b. Click **Site Preferences**.
- c. Under General Settings > Tasks, select a **Process Engine Connection Point** from the drop down list.
- d. Click **Apply** and then **Exit**.
- e. Confirm that Application Engine is communicating with Process Engine.
  - i. In Workplace, click **Tasks**.
  - ii. Verify that the Tasks page displays.
- f. Sign out of Workplace.

5. (IBM FileNet P8 systems using Image Services Integration only) Verify that you are running a supported version of ISRA.

See the *IBM FileNet P8 Hardware and Software Requirements* document for details on supported ISRA versions. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

For more information on how to upgrade ISRA and configure Image Services Integration, see your ISRA documentation and [“Enable Application Engine to Use ISRA” on page 451](#).

## Upgrade Add-On Components

You must upgrade all of the IBM FileNet P8 Platform core components before upgrading the add-on components listed in this topic. See [“Upgrade Core Components” on page 477](#) for the tasks required to upgrade the IBM FileNet P8 Platform core components.

The IBM FileNet P8 Platform add-on components can be upgraded in any order, except that you must upgrade Process Analyzer before upgrading Process Simulator.

### **To upgrade optional IBM FileNet P8 components**

---

- Upgrade Enterprise Manager (standalone). Do [Task 14 on page 559](#).
- Upgrade Workplace Application Integration and File Tracker. Do [Task 15 on page 560](#).
- Upgrade IBM FileNet Publishing components. Do [Task 16 on page 562](#).
- Upgrade custom applications. Do [Task 17 on page 563](#).
- Upgrade Server-Side Scripts and COM Objects. Do [Task 18 on page 610](#).
- Upgrade ISRA Servlet. Do [Table 19 on page 617](#).
- Install service packs or fix packs for optional add-on components. Do [Task 20 on page 622](#).

## Task 14: Upgrade Enterprise Manager (Standalone)

To upgrade a standalone Enterprise Manager from version 3.5.x to version 4.0.0, you need only run version 4.0.0 Content Engine Setup (or its silent counterpart), as documented in [“Install and Deploy Content Engine” on page 160](#).

**CAUTION** You cannot install version 4.0.0 of Enterprise Manager on a machine where version 3.5.x of the component is already installed.

## Task 15: Upgrade Application Integration and File Tracker

Upgrading to Workplace Application Integration 4.0.0 involves installing the new version on top of your existing Workplace Application Integration 3.5.x install. The upgrade installation program detects the earlier version and notifies you before proceeding with the upgrade.

During the upgrade, the installer detects the Microsoft applications that were previously integrated with Workplace and retains that configuration. For example, if you had Microsoft Outlook integrated with Workplace in the 3.5.x release, Microsoft Outlook is upgraded to 4.0.0 during the upgrade installation.

You cannot change your installed components during the upgrade. You can add or remove a Microsoft application from Application Integration after the upgrade using your Add/Remove Programs application.

Verify your computer meets the platform requirements documented in the *IBM FileNet P8 Hardware and Software Requirements*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

### *Upgrade considerations*

The procedures for upgrading Application Integration and File Tracker depend on which versions are currently installed on the client machine.

- If you upgrade from Application Integration version 3.5.1-001 or earlier, the upgrade to 4.0.0 will install File Tracker for you.
- If you upgrade from Application Integration version 3.5.1-002 or higher, then these conditions apply:
  - If you only had Application Integration installed (without File Tracker), then the upgrade to 4.0.0 will install File Tracker for you.
  - If you only had File Tracker installed (without Application Integration), then see [“To upgrade Workplace File Tracker” on page 561](#).
  - If you had both Application Integration and File Tracker installed, then the order that you installed these determines how you upgrade.
    - If you installed File Tracker 3.5.1-002 or higher before you installed Application Integration, you must upgrade File Tracker before you upgrade Application Integration. See [“To upgrade Workplace File Tracker” on page 561](#), and then see [“To upgrade Workplace Application Integration” on page 561](#).
    - If you installed Application Integration 3.5.1-002 or higher, and later used File Tracker, then you only need to upgrade Application Integration. See [“To upgrade Workplace Application Integration” on page 561](#)

**NOTE** If you upgrade from Application Integration version 3.5.1-002 or earlier, there is a change in the behavior of entry templates for adding email. Beginning with version 3.5.1-003, entry templates for adding an email are no longer restricted to using the email class or subclass.



### To upgrade Workplace Application Integration

---

1. Log onto the client machine with Application Integration installed using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Click **Download Application Integration for Microsoft Office**. The File Download dialog box for your system appears. Do one of the following:
  - Click **Open** to run the program from its current location.
  - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the **ApplicationIntegration.exe** file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the upgrade installer.

**NOTE** If you have Workplace Application Integration 3.5.x currently installed, you are prompted about the impending upgrade. Click **Yes** to upgrade to the current version you are installing or click **No** to end the installation.

5. After the install is complete, click **Finish** to complete the upgrade process.

### To upgrade Workplace File Tracker

---

1. Log on to the client machine with File Tracker installed using an account that has Administrator privileges.
2. Sign in to Workplace.
3. Click **Author**, and then click **General Tools**.
4. Scroll down and click **Download File Tracker** and do one of the following:
  - Click **Open** to run the program from its current location.
  - Click **Save**. In the Save As dialog box, find a location on your machine in which to download and save the **FileTracker.exe** file locally, and then click **Save**. Once the file is saved to your hard drive, double-click the file to run the upgrade installer.
5. After the install is complete, click **Finish** to complete the upgrade process.

### To verify your Workplace Application Integration upgrade

---

1. Start Microsoft Word.
2. From the File menu, click FileNet P8, point to Open Document, and then select Select Item. The Logon dialog box opens.
3. Log on using any valid domain account.
4. Click **Options** to view the Workplace address. The version number appears below the address.
5. Close all dialog boxes and close Microsoft Word.

## Task 16: Upgrade IBM FileNet Publishing Components

If you have IBM FileNet Rendition Engine software currently installed, you can upgrade to the latest software version. For instructions, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide* at [FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

## Task 17: Upgrade Custom Applications

This release includes changes to the following development tools that might impact custom applications:

- Content Engine COM API
- Content Engine Web Service (CEWS)
- Content Java API
- Process Java API
- Web Application Toolkit
- Workplace Application Integration Toolkit

If you have custom applications that were written using any of these APIs, Toolkits, or CEWS, then review the appropriate sections below for code changes that might be required.

### ***Content Engine COM API***

This section describes changes necessary to your 3.5.x Content Engine COM API applications so they can run in a 4.0 environment. The 4.0 release includes a Content Engine COM API Compatibility Layer, which is a client-side API that allows you to upgrade applications written using the 3.5.x Content Engine COM API. The compatibility layer is designed to maximize support for the 3.5.x COM API in a 4.0 environment, providing both source and binary compatibility for Visual Basic 6, C++, and script applications, particularly those using non-administrative interfaces. However, because of differences in platform technologies, parts of the COM API are not supported (such as those interfaces that are useful to administrative applications). In addition, you should be aware that because the compatibility layer is an emulation layered on top of the CE 4.0 .Net API, it does not provide performance parity with the natively implemented 3.5.x COM API. Therefore, some applications with stringent performance requirements might need to seek an alternative migration strategy (probably, to rewrite directly to the CE 4.0 .Net API).

The compatibility layer is provided only as a backward compatibility layer to support existing applications, and can be installed using the Content Engine installer. You can also install the compatibility layer on a client system. The COM Compatibility Layer requires Microsoft .NET Framework (2.0) and Web Services Enhancements (3.0) for all COM clients.

Documentation for the compatibility layer is only available on the [IBM Information Management support page on www.ibm.com](#). The documentation consists of the COM API help that was available for 3.5.x, with an updated reference section that reflects the compatibility layer. Note that the Guide portion of the help, as well as the Code Examples, have not been updated for 4.0.0.

### **New Content Engine 4.0 Object Types**

New object types introduced by Content Engine 4.0 (for example, Site, Index Area, Virtual Server) are not directly supported by the COM Compatibility Layer. However, it is possible in some circumstances to operate on these objects to a limited extent as GenericObject (see below).

## COM API Object Types Not Supported

The following COM API object types are not supported by Content Engine 4.0:

- CbrEngineType
- Computer
- ContentCacheService
- ContentManagerService
- FileStore
- ObjectStoreService
- Transient

These object types and their associated collection interfaces are defined in the COM Compatibility Layer, but there are no implementation classes. Any navigational properties (of supported object types) that, in the existing COM API, delivered a singleton or collection of these object types (such as `ObjectStore.FileStores`) return a "Value not set" error using the COM Compatibility Layer.

## New Content Engine 4.0 Properties on Supported Object Types

The COM Compatibility Layer does not provide type-safe accessors for properties defined solely in the Content Engine 4.0 API. However, the COM API Properties collection does expose such properties. These properties are also exposed in the PropertyDescriptions collection of the ClassDescription for the object. Therefore, for example, you cannot retrieve one of these properties in the following way:

```
doc.CompoundDocumentState
```

But can retrieve it, as follows:

```
doc.Properties.Item("CompoundDocumentState").Value
```

For a navigational property that delivers a singleton or collection of a Content Engine 4.0-only object type, retrieving the property value through the Properties collection yields either a `GenericObject` or a `ReadOnlyObjectSet` collection of `GenericObjects`.

Properties falling into this category include the Document class additions for compound document support and `PropertyDescription/Template/DefinitionXyz.PersistenceType`, which is added in place of `IsPersistent`.

## COM API Properties Not Supported

Content Engine 4.0 does not support a number of properties that existed in the 3.5.x Content Engine COM API (although in many cases these properties were deprecated). These unsupported properties are handled in the exact opposite way to new 4.0 properties; the type-safe accessors remain and behave as described below, but the properties are absent from the Properties collection (and from the PropertyDescriptions collection of the ClassDescription for the object).

The behavior of the type-safe accessors fall into two categories. In many cases, the information that was delivered by the 3.5.x property exists in another form or can be synthesized from state

held in the 4.0 object, and in such cases the compatibility layer performs such translation/synthesis to produce a compatible value for the property. Properties in this "emulated" category include:

- various.CreatePending, DeletePending, UpdatePending, deprecated in a previous release; use the PendingOperation property instead
- various.OIID, deprecated in a previous release; obsolete
- various.ObjectType, deprecated in a previous release; use the IsOfClass method instead
- various.ObjectStore, deprecated in a previous release; obsolete
- various.InstanceType, deprecated in a previous release; use the IsOfClass method instead
- ObjectStore.DatabaseType, deprecated in a previous release
- PropertyDescription/Template/Definition.IsPersistent, renamed to PersistenceType
- Subscription.EnableOn\*, deprecated in a previous release; use SubscribedEvents instead

For the remaining unsupported properties, there is no means of synthesizing a compatible value, so the type-safe accessor returns either a "value not set" error or, in some cases, an empty collection. Properties in this category include:

- ObjectStore.DatabaseConnectionString, the 4.0 Content Engine uses a completely different mechanism for specifying the database for an object store, and that mechanism does not allow compatible values for this property to be synthesized
- ObjectStore.DatabaseName, deprecated in a previous release
- ObjectStore.DatabaseServerName, deprecated in a previous release
- ObjectStore.EnumBatchSize, DefaultQueryBatchSize, DefaultQueryRowLimit, MaxQueryRowLimit, see ["Query-related ObjectStore Properties" on page 572](#)
- StoragePolicy.StorageRepositoryType, see ["StorageRepositoryType Property Not Supported" on page 573](#)
- QueryOperatorDescriptions, deprecated in a previous release; for internal use only
- Realm.Users and Realm.Groups, use FindUsers or FindGroups instead

Additional information on removed properties is provided in sections below.

## Content Engine 4.0 Methods Not Supported

New methods added by Content Engine 4.0 to objects supported by the COM Compatibility Layer are not exposed in the COM Compatibility Layer.

Methods supported by the COM API but not by the Content Engine 4.0 API are exposed in the COM Compatibility Layer interfaces, but attempting to call these methods will return a "Not supported" error. The following is a list of these methods:

### Export/Import-related

- ClassDescription.ExportSchema

- ObjectStore.ImportManifest
- ObjectStore.Export

### Content-Based Retrieval-related

- ObjectStore.ConfigureCbrEngine
- ObjectStore.UnconfigureCbrEngine
- ObjectStore.ConfigureCbrStore
- ObjectStore.UnconfigureCbrStore
- ObjectStore.Reindex

### File Store-related

- ObjectStore.CreateFileStore
- ObjectStore.ReindexFileStore

### OLEDB/ADO-related

- ObjectStore.GetADODConnection

### Miscellaneous

- ObjectStore.FilterClassDescriptions
- ObjectStore.FilterPropertyDescriptions

If your existing code uses either of these properties, you must write your own filtering logic based on iterating through the ObjectStore.ClassDescriptions collection or using ObjectStore.GetObject("ClassDescription", "<classname>").

### Validation of Property Values

When using the COM Compatibility Layer in Content Engine 4.0, property values are validated on the server rather than on the client side. The effect is that invalid property values are detected when the objects are saved, not when the property values are set.

The principal functional impact is that code like the following will fail differently with the COM Compatibility Layer than under 3.x.

```
Obj.IntProp = <value greater than maximum permitted>  
Obj.Save
```

In 3.x, a failure will be reported at the first statement, whereas with the COM Compatibility Layer the failure will not occur until the second.

### Default Method Short-Cutting

When using Visual Basic and VBScript to access the 3.x Content Engine COM API, a short-cut syntax is supported where a default method invocation can be omitted from certain forms of a

statement. Examples of the short-cut forms are shown below, with the corresponding non-short-cut form immediately following (the last example contains two short-cuts):

1. `Set prop = obj.Properties("DocumentTitle")`  
`Set prop = obj.Properties.Item("DocumentTitle")`
2. `Set cd = objStore.RootClassDefinitions("Document")`  
`Set cd = objStore.RootClassDefinitions.Item("Document")`
3. `obj.Properties("DocumentTitle") = "My Doc Title"`  
`obj.Properties.Item("DocumentTitle").Value = "My Doc Title"`

All of these short-cut forms are fully supported by the Content Engine COM Compatibility Layer for VBScript clients, and (1) is also fully supported for Visual Basic clients. However, there are some limitations for Visual Basic clients when using forms (2) and (3), depending on how the target object variable (`objStore` or `obj` in the examples above) is declared. If the target variable is DIMmed as `Object`, forms (2) and (3) both work. However, if the target variable is DIMmed as an FNCE type (such as `FNCE.Document`), form (3) does not work at all, and form (2) only works if the collection property is defined within that type. An example of how form (2) would not work is:

```
Dim event As ObjectChangeEvent
...
MsgBox event.ModifiedProperties(1)
```

This does not work because `ObjectChangeEvent` does not include the `ModifiedProperties` property. However, if the DIMmed type is changed to `UpdateEvent`, which does include the `ModifiedProperties` property, it works.

In cases where the short-cut syntax does not work, the statement will fail with an error indicating the "wrong number of arguments or invalid property assignment." You can resolve these limitations in your application in one of the following ways:

- DIM as `Object`
- Make the short-cut method call explicit.
- For form (3), use the more economical equivalent:  
`obj.DocumentTitle = "My Doc Title"`

Note that there are variety of circumstances beyond the simple examples given above where this shortcutting limitation can be encountered. As a general rule, any time you get the "wrong number of arguments ..." error, it is likely that you will need to adopt one of the remedies described above.

## VB Error Handling

The Visual Basic runtime maps certain HRESULTs into its own runtime errors. For example, `DISP_E_TYPEMISMATCH` is normally converted to runtime error 13, `E_INVALIDARG` to runtime error 5, and `DISP_E_BADINDEX` to runtime error 9. This affects the value of `Err.Number` in any error handling code, as well as the error message displayed. For these cases, the Content Engine COM Compatibility Layer honors the HRESULT returned, but this has proven unsuccessful in inducing Visual Basic to compatibly map to the same runtime errors. The result is that Visual Basic code that handles specific errors falling into this category must be modified to check for the unmapped HRESULT rather than for the runtime error number.

## Property Side-effects

In previous releases, certain property-setting operations and methods had an immediately observable side-effect on other properties. In the COM API Compatibility Layer, there is no such immediate side-effect for those properties that are reflected directly through from the Properties collection of the underlying object. In these cases, a save and refresh is required before the side-effects are observable.

The most notable example of this is the Name property. For most objects, Name is synthetic; it reflects either the value of the designated name property for the object (such as DocumentTitle), or the String version of the object ID if there is no designated name property. In previous releases, the effect was that an update to the designated name property was immediately observable in the value of the Name property. However, in Content Engine 4.0, the synthesis of Name takes place on the server, and on the client the Name property is indistinguishable from "normal" (unsynthesized) properties; so any changes to the corresponding name property only become visible in Name after they are saved.

## Managed Client Applications Using Catclient

The same COM interoperability features of the .NET runtime that allow the COM API Compatibility Layer to masquerade as the CE 3.5.x COM API also operate in the opposite direction, allowing managed clients (VB.Net or C#) to call COM APIs. Thus it is possible to write a managed application that uses the CE 3.5.x COM API.

However, such applications do not run when the COM API Compatibility Layer is substituted in place of the CE 3.5.x COM API; they fail when instantiating FNCE.EntireNetwork, with what appears to be a registration error. This is believed to be because of a limitation of the .NET runtime preventing transitioning through the COM interoperability layer twice (once calling out from the client application, once calling in to the COM API Compatibility Layer).

Therefore, such managed clients are not supported by the COM API Compatibility Layer, and would need to be rewritten to use the native 4.0 Content Engine .NET API.

## Error Reporting

A new error handling framework has been developed for Content Engine 4.0. The COM Compatibility Layer maps the .NET exceptions generated using the new framework into standard HRESULT return codes. Exceptions specific to Content Engine 4.0 are returned as "E\_FAIL".

In most cases, the error codes returned are unchanged from the 3.5.x Content Engine COM API. However, in some cases error codes and possibly error semantics have changed. Applications highly dependent on specific error returns might require some modification.

## Security

### Permissions Property for Callers without Appropriate Access

In 3.x, attempting to access the Permissions property of an object for which the caller did not have READ\_ACL, WRITE\_ACL, or WRITE\_OWNER rights would result in an error. In 4.0, the returned Permissions list will be empty. When you encounter an empty Permissions list, you must consult



the object's AccessMask to see if the list is empty because the caller lacks READ\_ACL, WRITE\_ACL, or WRITE\_OWNER rights. (This applies to any object with a Permissions property.)

### Extensible Authentication Framework

The Extensible Authentication Framework was introduced in FileNet P8 3.5 to facilitate integration with Single-Sign-On (SSO) solutions that would not allow the transmission of a username and password to the FileNet system. FileNet P8 4.0 introduces a standards-based methodology for authentication using the J2EE Java Authentication and Authorization (JAAS) framework.

If you have implemented the Extensible Authentication Framework in your 3.5 application, for the 4.0 release you must write a JAAS LoginModule using the FileNet Web Service Extensible Authentication Framework. Note that this approach does not require any modifications to existing applications, as the existing client-side logic to submit custom values in the username and password fields will remain in place.

You write the custom JAAS LoginModule on the server, conforming to FileNet's 4.0 Web Service Extensible Authentication Framework (WS-EAF). This custom LoginModule will extract the username and password values from the UsernameToken in the incoming WS-Security header, and get the same custom values that were used in the 3.5.2 EA solution, and then invoke whatever custom logic is required to authenticate based on those parameters. For more information, see the [IBM FileNet P8 help topics under Developer Help > Web Service Extensible Authentication Framework Developer's Guide > Introduction](#).

### Directory Service Lookups

In FileNet P8 3.5.x, LDAP directory lookups could be done using a system configured account, or the user context. In FileNet P8 4.0, directory service lookups based on the user context are no longer supported. If you have an existing application that enumerates users, more results might be returned in the 4.0 release (since lookups done using the user context would have returned only those users and groups to which the user had access).

### Logon Names for Windows Active Directory

In 3.5.x, a logon (principal) name that contained an embedded @ character could successfully log on to a Content Engine server. In 4.0, logon names are no longer allowed to have any @ characters embedded in them because anything after the @ symbol is considered to be the domain name.

### Pre-Windows 2000 Login Format

In 3.5.x, the Content Engine server supported a pre-Windows 2000 login format (domain\login-name) through Microsoft-specific authentication APIs. The functionality is no longer directly available through the J2EE application server-based authentication modules. You are limited to what the application server (or the independent authentication provider that it uses) supports in terms of authentication login-name formats.

## OLEDB/ADO Data Provider

Content Engine 4.0 does not include an OLEDB Provider, which in the 3.5.x release was the mechanism provided for Windows applications to access the search (query) capabilities of Content Engine (either directly or through ADO). Applications requiring this functionality with Content Engine 4.0 must seek alternative means:

- Content Engine 4.0 includes a [JDBC provider](#) (JDBC is the Java equivalent of OLEDB). Many off-the-shelf reporting tools (such as Business Objects' Crystal Reports) are available in both an OLEDB version and a JDBC version, so where such a tool is being used, a simple option is to switch from the OLEDB version to the JDBC version.
- Otherwise, your application must be ported to use the Content Engine 4.0 Java API (or .NET API), which provides full access to [search functionality](#).

## Microsoft IIS (Internet Information Services) WebDAV Provider

Content Engine 3.5.x used IIS and ISAPI to implement the WebDAV provider on the Content Engine host. Content Engine 4.0 no longer relies on IIS, and instead implements a Java-based WebDAV provider on the Application Engine host. Existing applications calling the WebDAV provider will need to be reconfigured to reference the new location on the Application Engine host.

## Publishing

Publishing via the Content Engine COM API Compatibility Layer is not available in the 4.0 release.

## COM API Extensions

The Content Engine 3.5.x COM API Extensions--which exposed an API to create custom application server publishing applications (for example, a rendition engine or an administrative application)--are no longer used, and are not supported in the COM Compatibility Layer.

## CEEDocument Class

The CEEDocument class is no longer used, and is not supported in the COM Compatibility Layer. Content Engine 4.0 API publishing-related methods are on the Document class itself.

## Events and Subscriptions

### Event Actions

Content Engine 3.5.x supported scripts and COM objects for event actions. In the 4.0 release, event actions must be Java-based. Any EventAction objects that your application uses must be modified to use a Java class instead of a script or a COM object. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. (For information on how to create a Java-based event action, see [Work with event actions](#) in the *4.0 Help for Content Engine Administration*.) In addition, changes in the following EventAction properties might require modifications in your application:

- **ProgId** - In 3.5.x, this property contained a string (such as VBScript) that specified the type of code to be executed when a Content Engine event occurred. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- **ScriptText** - In 3.5.x, this property contained the script text of event procedures that executed when an action-related object's events were raised. In 4.0, the text in this property is ignored and should be left empty.

### Priority Property Deprecated

In 3.5.x, you could specify a value used to order the firing of subscriptions in response to an event via the **Priority** property on a **Subscription** object. However, the server did not use this property, and therefore the property has been deprecated in 4.0.

### Workflow Subscription Filter Expressions

In 3.5.x, a **Process Engine** filter expression could be supplied with a workflow subscription using the **Expression** property. This type of filter expression was evaluated by the **Process Engine** after the workflow event was launched. Because the filter expression evaluation occurred well after the original transaction was complete, this capability was inefficient and often inaccurate. In the 3.5.x documentation, IBM recommends that you use the **FilterExpression** property instead, as the **FilterExpression** property is evaluated within the transaction that initiates the event (before an event is queued).

In 4.0, the **Expression** property is no longer supported. During the upgrade to **Content Engine 4.0**, if the **Expression** property contains a value, then the value will be moved to the **FilterExpression** property (but only if the **FilterExpression** property was not populated with a value). If both properties were populated with a value, then the upgrade tool will not move the contents of the **Expression** property but will produce a report of these cases.

Note that the **Expression** property supported certain SQL syntax that is not supported in the **FilterExpression** property. Specifically, with the exception of the **LIKE** operator, all other operators requiring a database query (such as the **EXISTS** operator) are not supported in the **FilterExpression** property. For a list of unsupported operators, see [FilterExpression Property](#) in the *4.0 Content Engine Java and .NET Developer's Guide*.

If your application uses the **Expression** property, you must modify your code to use the **FilterExpression** property instead. In addition, if any 3.5.x **Expression** values included a SQL operator that is not supported in the **FilterExpression** property, you must resolve these instances by either programmatically setting the value of the **FilterExpression** property to a string that includes only supported operators or by using **Enterprise Manager** to correct the value.

### IsolatedRegion Property Deprecated

In the 4.0 release, the **IsolatedRegion** property on the **ClassWorkflowSubscription** and **InstanceWorkflowSubscription** classes has been renamed to **IsolatedRegionNumber**. The old property (**IsolatedRegion**) is deprecated, but existing code that accesses this property via the type-safe accessor will continue to work.

## Query

### Content Based Retrieval

Content Engine 3.5.x integrated the Verity VDK for CBR. For Content Engine 4.0, Autonomy K2 (the former Verity product) is integrated for CBR. The following changed search behaviors are a result of integrating Autonomy K2 as the CBR engine:

- When an individual property is identified for CBR, only the specified property is searched.
- You can now do a simultaneous search for text in properties that is different for text in content.
- Search no longer supports identification of a content element.

In addition, the following Content Engine COM API methods are no longer supported as they are administrative functions that are no longer relevant to the 4.0 management of CBR:

- ObjectStore.ConfigureCbrEngine
- ObjectStore.UnconfigureCbrEngine
- ObjectStore.ConfigureCbrStore
- ObjectStore.UnconfigureCbrStore

### Query-related ObjectStore Properties

The following query configuration-related properties on an ObjectStore are no longer supported:

- EnumBatchSize
- DefaultQueryBatchSize
- DefaultQueryRowLimit
- MaxQueryRowLimit

In the 4.0 release, the functionality supplied by these properties has been replaced with properties on the new ServerCacheConfiguration class:

- QueryPageMaxSize
- QueryPageDefaultSize
- NonPagedQueryMaxSize

Any existing code that accessed the ObjectStore properties will no longer function. Either rewrite your code using one of the 4.0 Content Engine APIs, or manually set the new ServerCacheConfiguration properties using Enterprise Manager.

### QueryOperatorDescriptions Property Not Supported

In a previous release, the QueryOperatorDescriptions property on PropertyDescription\* classes was deprecated. In the 4.0 release, attempting to retrieve this property will return a DMARC\_NOT\_SUPPORTED error.

## Storage

### StoragePolicy Class

Content Engine 4.0 has only the StoragePolicy class, whereas the COM API also has subclasses DatabaseStoragePolicy and FileStoragePolicy. Although DatabaseStoragePolicy and FileStoragePolicy are exposed in the COM compatibility layer, objects retrieved from the server (for example, through ObjectStore.StoragePolicies) will always be of the base class, and neither a DatabaseStoragePolicy nor a FileStoragePolicy instance can be created. This will generally have no impact on applications that simply enumerate the available policies and allow one to be selected as the policy for a new document. However, the following will occur:

- An IsOfClass test for either of the two subclasses will always return false.
- An attempt to create a storage policy will not be allowed.

### File Stores

In Content Engine 4.0, a File Store is a File Storage Area, and a Hybrid File Store is a Fixed Storage Area. StorageArea is a new abstract class representing a physical content storage location. It has three concrete subclasses: DatabaseStorageArea, FileStorageArea (replacing FileStore) and FixedStorageArea (replacing a hybrid FileStore). FixedStorageArea may can be further subclassed for fixed-content device-specific subclasses (such as IS or Snaplock). You can access the collection of StorageAreas in a non-typesafe manner through objStore.Properties.Item("StorageAreas").Value, which returns an ROObjectSet of GenericObjects.

### StorageRepositoryType Property Not Supported

In the 4.0.0 release, because a storage policy no longer (in general) references a single storage area, it therefore cannot be said to identify a particular storage repository type. Therefore, the StorageRepositoryType property on the StoragePolicy class and subclasses is no longer supported. Attempting to retrieve this property will return a DMARC\_NOT\_SUPPORTED error.

### Moving Document or Annotation Content

The MoveContent method on the Document and Annotation classes includes a StoragePolicy parameter. In the 4.0 release, you can pass in either a StoragePolicy or a StorageArea object (which is new in the 4.0 release). If you pass in a StoragePolicy object, the Content Engine COM Compatibility Layer will move the content to the first storage area selected by the filter for that policy.

### Fixed Content Devices

Content Engine 3.5.x supported Hitachi, Tivoli, and IBM IICE Fixed Content Devices on Windows platforms. The Hitachi and Tivoli Fixed Content Devices are not supported in Content Engine 4.0. If a Content Engine 3.5.x object store uses one of these devices, the 4.0 upgrade cannot succeed unless the content on these devices is first moved to a file store or to a Fixed Content Device supported by Content Engine 4.0. In addition, the following Content Engine COM API classes are no longer supported:

- `FSBFixedContentDevice`
- `TivoliFixedContentDevice`

### Changed FixedContentDevice Properties

In previous releases the `FixedContentDevice (FCD) ConfigurationParameters` property was one XML string defining all of the FCD parameters used to configure the various FCD devices. This XML string has been split into separate properties. Each of the FCD-specific properties are on the specific type of FCD (for example, the IS FCD has IS-specific parameters). Additionally, some of the remaining FCD property names have been renamed to be more accurate.

### Auditing

In previous releases, attempting to modify properties for which the user did not have appropriate access rights was recorded as an audit failure by the Content Engine. However, in the 4.0 release the Content COM API will immediately generate an error on the client under these circumstances, and therefore will not record the failed attempt in the audit log. (Only operations executed on the Content Engine server can be audited; failed or successful client operations are not audited by the server.)

### Document Lifecycles

#### Document Lifecycle Actions

Content Engine 3.5.x supported Microsoft ActiveScript technology for document lifecycle actions. In the 4.0 release, lifecycle actions must be Java-based. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. Any `DocumentLifecycleAction` objects that your application uses must be modified to use a Java class instead of a script. In addition, note changes in the following `DocumentLifecycleAction` properties:

- `ProgId` - In 3.5.x, this property contained a string (such as VBScript) that specified the type of code to be executed when a Content Engine event occurred. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- `ScriptText` - In 3.5.x, this property contained the script text of event procedures that executed when an action-related object's events were raised. In 4.0, the text in this property is ignored and should be left empty.

#### Cannot Delete Default DocumentLifecyclePolicy

In previous releases, you could delete any `DocumentLifecyclePolicy` object as long as the `DocumentLifecyclePolicy` was no longer being referenced by one or more `Document` instances. In the 4.0 release, you cannot delete a `DocumentLifecyclePolicy` if it is being referenced by one or more `Document` instances or if it is the default policy for a `Document` class.

## Document Classification

### COM-based Plug-ins Not Supported

Content Engine 3.5.x supported COM-based plug-ins for document classification actions. In the 4.0 release, classification actions must be Java-based. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. Any DocumentClassificationAction objects that your application uses must be modified to use a Java class instead of a COM class. For information on how to create a document classification action, see [Developing a custom classifier](#) in the *4.0 Help for Content Engine Administration*.) In addition, note changes in the following DocumentClassificationAction properties:

- **ProgId** - In 3.5.x, this property contained a Class Id or ProgId that represents the COM class. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- **ScriptText** - In 3.5.x, this property contained a string containing a script language program for implementing the classification behavior. If a value for ScriptText was set, then the ProgId property had to point to a Windows Scripting Host capable of interpreting that script. In 4.0, the text in this property is ignored and should be left empty.

### XMLPropertyMappingScript Algorithm

In 3.5.x, the algorithm used by the XMLPropertyMappingScript object was as follows:

1. Look for a Processing Instruction named FileNetDocClass. If it exists, use its value.
2. Else, look for a DTD with an external specifier. If it exists, use the value of the SYSTEM attribute (which usually contains a URL).
3. Else, use the name of the root element.

In 4.0, the algorithm is as follows:

1. Look for a Processing Instruction named FileNetDocClass. If it exists, use its value.
2. Else, get the expanded QName of the root element in the form of `{RootElementNamespaceURI}RootElementLocalName`. (The system looks for this in the XMLDocumentType property.)

If your existing application uses a FileNetDocClass Processing Instruction node in the XML content, no change is required for 4.0. Otherwise, you need to modify the value of the XMLDocumentType property of the XMLPropertyMappingScript object to use the new format:

```
{RootElementNamespaceURI}RootElementLocalName
```

For example, if the root element is:

```
<claim xmlns="http://www.filenet.com/claimSchema">
```

Then the expanded QName is:

```
{http://www.filenet.com/claimSchema}claim
```



## Content Type

In previous releases, when an application did not supply a content type value for a ContentReference or ContentTransfer object, the 3.x Content Engine used file extension mappings in the Windows registry along with Microsoft APIs to examine document content. Content Engine 4.0 identifies the content type based on a fixed listing of file extensions.

Because Content Engine 4.0 uses a set listing of file extensions to identify the document content type, the resulting (identified) content type can be different from that identified by Content Engine 3.5.x. To ensure that an object's content type is identified properly, explicitly specify the ContentType property for existing Content Engine 3.5.x applications when creating or updating a content element.

## VersionSeries Security

In previous releases, if a user attempted to retrieve a VersionSeries object for which he did not have access rights to the current version, the VersionSeries object would be returned (security checking was bypassed because the VersionSeries object was fabricated on the client side). In the 4.0 release, if the user does not have access rights to the current version, the COM Compatibility Layer will return a DMARC\_NOT\_FOUND error.

## Addons

Content Engine 3.5.x supported Microsoft ActiveScript technology for pre-install/post-install scripts for AddOns. In Content Engine 4.0, only JavaScript scripting is supported. If your application creates AddOn objects and uses any pre-import or post-import scripts written in VBScript, you must:

- Rewrite those scripts in JavaScript.
- In your existing code to create a AddOn instance, if your code set the ScriptType property to VBScript, change the property's value to JScript.

## Exception When Attempting to Modify Root Folder Properties

In previous releases, attempting to move a root folder generated a NOT\_SUPPORTED error. In the 4.0 release, an ACCESS\_DENIED error is generated. (Note that the ACCESS\_DENIED error is also thrown in 4.0 if you attempt to change a root folder's name or delete a root folder.)

## Exception When Attempting to Update the List of PropertyDefinition Objects for a ClassDefinition

In previous releases, attempting to update the list of PropertyDefinition objects for a ClassDefinition resulted in an ACCESS\_DENIED error. In 4.0, a NOT\_SUPPORTED error is generated instead.

## GUID Strings

In the 3.5.x release, operations that convert a GUID from a String to the internal form ignored any data after the closing } (brace). For example, the bolded text after the closing brace below would have been ignored:



```
{1321358E-DD13-4913-A44C-B57659A05C1E}, brown, filenet;comp:t3: //localhost:7001/  
FileNet/Engine,null
```

In 4.0, this is no longer the case; any applications relying on this behavior must be updated.

## Reusing Dependent Objects

In the 3.5.x release, you could reuse dependent objects and dependent object lists. For example, in 3.5.x you could create or retrieve the Permission objects for one object and then directly use those same Permission objects on another object. In the 4.0 release, dependent objects and dependent object lists are not reusable. If you attempt such reuse, a runtime error will be generated.

## Retrieving an ObjectStore's ClassDescription

In previous releases, if you retrieved the ClassDescriptions property for an ObjectStore, the ClassDescription for that ObjectStore would be included in the collection. In 4.0, this is no longer the case. However, as in 3.5.x, you can continue to retrieve the ClassDescription for an ObjectStore by retrieving the ClassDescription property on the ObjectStore object.

## Retrieving Realm Users and Groups

In previous releases, you could retrieve the users and groups in a Realm via the Users and Groups properties. However, because these properties returned exhaustive collections of all users and groups in the realm, IBM recommended that you use the FindUsers and FindGroups methods instead. In the 4.0 release, retrieving the Users and Groups properties of the Realm object returns empty collections instead of populated lists and you must use the Findxxx methods instead.

## Operating on Deleted Objects

In 3.5.x, it was possible to continue retrieving and setting properties on an object after the object had been successfully deleted with a non-pending Delete() or after committing (with Save() or SaveBatch()) a delete made pending with idmChangeSetPending. In 4.0 this is no longer the case, and attempting such operations will have unpredictable results.

## Type Property on Domain Class Not Supported

In the 4.0 release, the Type property on the Domain class is no longer supported. The Type property was meant to indicate the type of FileNet P8 domain (private, public, or workgroup), but public was the only type the Content Engine COM API ever supported.

## GCD Changes

In previous releases, the GCD (Global Configuration Data) was stored in the local file system. For the 4.0 release, the GCD is generally stored in a separate database instance, leveraging the RDBMS used by Content Engine.

During the upgrade operation from FileNet P8 3.5.x, the DateCreated and Creator properties for the following objects are replaced with the date of the upgrade and the name of the user performing the upgrade, respectively:

- ObjectStores
- MarkingSets
- FixedContentDevices

The DateCreated and Creator properties are not changed by the upgrade operation for objects that reside within an object store (such as Documents, Folders, and CustomObjects).

### PropertyDescription DateTime Values

In 3.x releases, if you specified a date earlier than 1 Jan 1753 (which was the constraint specified by the PropertyMinimumDateTime property on the PropertyDescriptionDateTime class) for a custom property, the Content Engine server silently changed the date to 1 Jan 1753. In the 4.0 release, the PropertyDescriptionDateTime class has been updated to reflect the database-specific limits on date/time values. System and custom property values exceeding the database constraints will generate an error. If your 3.5.x application relied on the above-mentioned PropertyMinimumDateTime behavior, you should modify your application to either pre-validate the date/time values, or to catch the error.

### Content Engine Web Service

This section describes changes necessary to your 3.5.x Content Engine Web Service API applications so they can run in a 4.0 environment. Note that the 3.5.x Content Engine Web Service API continues to support Microsoft .NET Framework 1.1, Web Services Enhancements (WSE) 2.0, and DIME attachments only.

Documentation for the 3.5.x Content Engine Web Service API is installed as part of the 4.0 FileNet P8 Documentation installation. You can navigate to the documentation, as follows: **FileNet P8 Documentation > Developer Help > Content Engine Development > Web Service Developer's Guide**.

Content Engine 4.0 introduces a number of new object types and also adds some properties to existing object types. These are visible to, and can be manipulated by, web service client applications. 4.0 also removes some object types and properties of existing object types, invalidating any web service client code that relies upon those objects or properties. The following sections give details of the removals.

### Classes Not Supported

The following classes are no longer supported by Content Engine 4.0:

- CbrEngineType
- Computer
- ContentCacheService
- ContentManagerService
- FileStore
- ObjectStoreService

- Transient

## Properties Not Supported

Content Engine 4.0 does not support a number of properties that existed in the 3.5.x Content Engine (although in some cases these properties were deprecated in the 3.5.x release). These unsupported properties include the following:

- `various.OIID`, in previous releases, this was retrievable only via queries; use **This** instead.
- `PropertyDescription/Template/Definition.IsPersistent`, renamed to `PersistenceType`.
- `Subscription.EnableOn*`, deprecated in a previous release; use `SubscribedEvents` instead.
- `ObjectStore.DatabaseConnectionString`, the 4.0 Content Engine uses a completely different mechanism for specifying the database for an object store. For more information on this mechanism, see the `jdbc.driver` class in the Content Engine Java API Reference documentation.
- `ObjectStore.DatabaseName`, deprecated in a previous release.
- `ObjectStore.DatabaseServerName`, deprecated in a previous release.
- `ObjectStore.EnumBatchSize`, `DefaultQueryBatchSize`, `DefaultQueryRowLimit`, `MaxQueryRowLimit`, see [“Query-related ObjectStore Properties” on page 572](#).
- `StoragePolicy.StorageRepositoryType`, see [“StorageRepositoryType Property Not Supported” on page 573](#).
- `QueryOperatorDescriptions`, deprecated in a previous release; for internal use only.
- `Realm.Users` and `Realm.Groups`, use the `PrincipalSearch` option of `ExecuteSearch` to obtain lists of users and/or groups meeting specified criteria.

Additional information on removed properties is provided in sections below.

## Security

### Extensible Authentication Framework

The Extensible Authentication Framework was introduced in FileNet P8 3.5 to facilitate integration with Single-Sign-On (SSO) solutions that would not allow the transmission of a username and password to the FileNet system. FileNet P8 4.0 introduces a standards-based methodology for authentication using the J2EE Java Authentication and Authorization (JAAS) framework.

If you have implemented the Extensible Authentication Framework in your 3.5 application, for the 4.0 release you must write a JAAS `LoginModule` using the FileNet Web Service Extensible Authentication Framework. Note that this approach does not require any modifications to existing applications, as the existing client-side logic to submit custom values in the username and password fields will remain in place.

You write the custom JAAS `LoginModule` on the server, conforming to FileNet's 4.0 Web Service Extensible Authentication Framework (WS-EAF). This custom `LoginModule` will extract the username and password values from the `UsernameToken` in the incoming WS-Security header,

and get the same custom values that were used in the 3.5.2 EA solution, and then invoke whatever custom logic is required to authenticate based on those parameters. For more information, see the [IBM FileNet P8 help topics under Developer Help > Web Service Extensible Authentication Framework Developer's Guide > Introduction](#).

## Directory Service Lookups

In FileNet P8 3.5.x, LDAP directory lookups could be done using a system configured account, or the user context. In FileNet P8 4.0, directory service lookups based on the user context are no longer supported. If you have an existing application that enumerates users, more results might be returned in the 4.0 release (since lookups done using the user context would have returned only those users and groups to which the user had access).

## User Credentials

Web service clients supply credentials for authentication through a Security header in each SOAP request, typically one containing a UsernameToken. Content Engine 4.0 uses a different authentication framework to 3.5, with consequences for the form of credentials that may be provided. This affects particularly the format of user names that may be used in the Username element of a UsernameToken, as described below.

## Logon Names for Windows Active Directory

In 3.5.x, a logon (principal) name that contained an embedded @ character could successfully log on to a Content Engine server. In 4.0, logon names are no longer allowed to have any @ characters embedded in them because anything after the @ symbol is considered to be the domain name.

## Pre-Windows 2000 Login Format

In 3.5.x, the Content Engine server supported a pre-Windows 2000 login format (domain\login-name) through Microsoft-specific authentication APIs. The functionality is no longer directly available through the J2EE application server-based authentication modules. You are limited to what the application server (or the independent authentication provider that it uses) supports in terms of authentication login-name formats.

## Events and Subscriptions

### Event Actions

Content Engine 3.5.x supported scripts and COM objects for event actions. In the 4.0 release, event actions must be Java-based. Any EventAction objects that your application uses must be modified to use a Java class instead of a script or a COM object. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. (For information on how to create a Java-based event action, see the [IBM FileNet P8 help topics under Content Engine Administration > Events and subscriptions > How to... > Work with event actions](#) in the *4.0 Help for Content Engine Administration*.) In addition, changes in the following EventAction properties might require modifications in your application:

- **ProgId** - In 3.5.x, this property contained a string (such as VBScript) that specified the type of code to be executed when a Content Engine event occurred. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- **ScriptText** - In 3.5.x, this property contained the script text of event procedures that executed when an action-related object's events were raised. In 4.0, the text in this property is ignored and should be left empty.

### Priority Property Deprecated

In 3.5.x, you could specify a value used to order the firing of subscriptions in response to an event via the Priority property on a Subscription object. However, the server did not use this property, and therefore the property has been deprecated in 4.0.

### Workflow Subscription Filter Expressions

In 3.5.x, a Process Engine filter expression could be supplied with a workflow subscription using the Expression property. This type of filter expression was evaluated by the Process Engine after the workflow event was launched. Because the filter expression evaluation occurred well after the original transaction was complete, this capability was inefficient and often inaccurate. In the 3.5.x documentation, IBM recommended that you use the FilterExpression property instead, as the FilterExpression property is evaluated within the transaction that initiates the event (before an event is queued).

In 4.0, the Expression property is no longer supported. During the upgrade to Content Engine 4.0, if the Expression property contains a value, then the value will be moved to the FilterExpression property (but only if the FilterExpression property was not populated with a value). If both properties were populated with a value, then the upgrade tool will not move the contents of the Expression property but will produce a report of these cases.

Note that the Expression property supported certain SQL syntax that is not supported in the FilterExpression property. Specifically, with the exception of the LIKE operator, all other operators requiring a database query (such as the EXISTS operator) are not supported in the FilterExpression property. For a list of unsupported operators, see [FilterExpression Property](#) in the *4.0 Content Engine Java and .NET Developer's Guide*.

If your application uses the Expression property, you must modify your code to use the FilterExpression property instead. In addition, if any 3.5.x Expression values included a SQL operator that is not supported in the FilterExpression property, you must resolve these instances by either programmatically setting the value of the FilterExpression property to a string that includes only supported operators or by using Enterprise Manager to correct the value.

### IsolatedRegion Property Deprecated

In the 4.0 release, the IsolatedRegion property on the ClassWorkflowSubscription and InstanceWorkflowSubscription classes has been renamed to IsolatedRegionNumber. Any existing code that accesses the old name (IsolatedRegion) must be rewritten to use the new name.

## Query

### Content Based Retrieval

Content Engine 3.5.x integrated the Verity VDK for CBR. For Content Engine 4.0, Autonomy K2 (the former Verity product) is integrated for CBR. The following changed search behaviors are a result of integrating Autonomy K2 as the CBR engine:

- When an individual property is identified for CBR, only the specified property is searched.
- You can now do a simultaneous search for text in properties that is different for text in content.
- Search no longer supports identification of a content element.

### Query-related ObjectStore Properties

The following query configuration-related properties on an ObjectStore are no longer supported:

- EnumBatchSize
- DefaultQueryBatchSize
- DefaultQueryRowLimit
- MaxQueryRowLimit

In the 4.0 release, the functionality supplied by these properties has been replaced with properties on the new `ServerCacheConfiguration` class:

- QueryPageMaxSize
- QueryPageDefaultSize
- NonPagedQueryMaxSize

### Query Results

If your application uses the `ExecuteSearch` method to perform a query that retrieves all information from a Content Engine 4.0 table (using `SELECT *`), a much larger result set will be returned than in previous Content Engine releases. This is because previously, many system properties were marked as not selectable. In Content Engine 4.0, all properties are selectable. Some of the newly selectable properties can be quite large; in particular, any properties of type `ListOfObject` (such as `Permissions` or `PropertyDefinitions`). This extra data can cause an increase in response times for queries. For any application that traverses all properties in a result set, this can cause a further increase in response time, and could change behavior if new, unexpected data types are encountered. Queries that specify an explicit list of columns in their `SELECT` clause will not be affected.

### Unevaluated Properties in Query Results

If your application uses the `ExecuteSearch` method to perform a query that retrieves information from a Content Engine 4.0 table, some new properties that are returned from the query might be of type `Unevaluated`. (A query against a Content Engine 3.5.x table would never have returned an unevaluated object.) An object-valued property can be in an unevaluated state when the server

cannot determine whether a value exists for that property. Any code that traverses all properties of a query result row might need to be updated to handle unevaluated properties.

## Storage

### StoragePolicy Class

Content Engine 4.0 has only the StoragePolicy class, whereas Content Engine 3.5.x also has subclasses DatabaseStoragePolicy and FileStoragePolicy. In the 4.0 release, objects retrieved from the server (for example, when retrieving the StoragePolicies property of an ObjectStore) will always be of the base class, and neither a DatabaseStoragePolicy nor a FileStoragePolicy instance can be created. This will generally have no impact on applications that simply enumerate the available policies and allow one to be selected as the policy for a new document

### File Stores

In Content Engine 4.0, a File Store is a File Storage Area, and a Hybrid File Store is a Fixed Storage Area. StorageArea is a new abstract class representing a physical content storage location. It has three concrete subclasses: DatabaseStorageArea, FileStorageArea (replacing FileStore) and FixedStorageArea (replacing a hybrid FileStore). FixedStorageArea can be further subclassed for fixed-content device-specific subclasses (such as IS or Snaplock).

### StorageRepositoryType Property Not Supported

In the 4.0.0 release, because a storage policy no longer (in general) references a single storage area, it does not identify a particular storage repository type. Therefore, the StorageRepositoryType property on the StoragePolicy class and subclasses is no longer supported.

### Moving Document or Annotation Content

In previous releases, the targetPolicyId attribute for a MoveContentAction element specified the ID of a StoragePolicy object. In the 4.0 release, you must specify the ID of a StorageArea object.

### Fixed Content Devices

Content Engine 3.5.x supported Hitachi, Tivoli, and IBM II CE Fixed Content Devices on Windows platforms. However, Hitachi and Tivoli Fixed Content Devices are not supported in Content Engine 4.0. If a Content Engine 3.5.x object store uses one of these devices, the 4.0 upgrade cannot succeed unless the content on these devices is first moved to a file store or to a Fixed Content Device supported by Content Engine 4.0. In addition, the following Content Engine classes are no longer supported:

- FSBFixedContentDevice
- TivoliFixedContentDevice



## Changed FixedContentDevice Properties

In previous releases the FixedContentDevice (FCD) ConfigurationParameters property was one XML string defining all of the FCD parameters used to configure the various FCD devices. This XML string has been split into separate properties. Each of the FCD-specific properties are on the specific type of FCD (for example, the IS FCD has IS-specific parameters). Additionally, some of the remaining FCD property names have been renamed to be more accurate.

## Retrieving an ObjectStore's ClassDescription

In 3.5.x, you could return the ClassDescription object for the ObjectStore class via ObjectStore.GetObjects() or Domain.GetObjects(). In 4.0, only Domain.GetObjects() will work; ObjectStore.GetObjects() will return an error.

In addition, in previous releases if you retrieved the ClassDescriptions property for an ObjectStore, the ClassDescription for that ObjectStore would be included in the collection. In 4.0, this is no longer the case. However, as in 3.5.x, you can continue to retrieve the ClassDescription for an ObjectStore by retrieving the ClassDescription property on the ObjectStore object.

## ObjectStore DatabaseType property

In 3.x, this property was a string, giving the name of the OLEDB provider for the type of database. In 4.0, this is replaced by an integer property of the same name, with a value indicating the database type as follows:

- 1 - Microsoft SQL Server
- 2 - Oracle
- 3 - IBM DB2

## Content Type

In previous releases, when an application did not supply a content type value for a ContentReference or ContentTransfer object, the 3.x Content Engine used file extension mappings in the Windows registry along with Microsoft APIs to examine document content. Content Engine 4.0 identifies the content type based on a fixed listing of file extensions.

Because Content Engine 4.0 uses a set listing of file extensions to identify the document content type, the resulting (identified) content type can be different from that identified by Content Engine 3.5.x. To ensure that an object's content type is identified properly, explicitly specify the ContentType property for existing Content Engine 3.5.x applications when creating or updating a content element.

## Creating an Addon

Any 3.5.x CE WS application that creates an AddOn requires code changes to work with a CE 4.0 server. In particular, you must set the ImportData property (of type ContentData) instead of setting the XMLManifest property (of type SingletonString). Note that these code changes are required only for code that creates an AddOn and not for code that installs an AddOn. In addition, the PreImportScript and PostImportScript properties were of type String, but in the 4.0 release are of type ContentData.



## IsPersistent Property

Previous releases included the `IsPersistent` property on `PropertyDefinition` and `PropertyDescription` classes. In the 4.0 release, this property has been renamed to `PersistenceType` for these classes. (For the `ClassDefinition` and `ClassDescription` classes, the property name remains `IsPersistent`). Although the server will convert `IsPersistent` to `PersistenceType` during `PropertyTemplate` creation, any existing code that attempts to retrieve the `IsPersistent` property for `PropertyDefinition` and `PropertyDescription` objects must be changed to retrieve the `PersistenceType` property.

## Document Lifecycles

### Document Lifecycle Actions

Content Engine 3.5.x supported Microsoft ActiveScript technology for document lifecycle actions. In the 4.0 release, lifecycle actions must be Java-based. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. Any `DocumentLifecycleAction` objects that your application uses must be modified to use a Java class instead of a script. In addition, note changes in the following `DocumentLifecycleAction` properties:

- `ProgId` - In 3.5.x, this property contained a string (such as VBScript) that specified the type of code to be executed when a Content Engine event occurred. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- `ScriptText` - In 3.5.x, this property contained the script text of event procedures that executed when an action-related object's events were raised. In 4.0, the text in this property is ignored and should be left empty.

### Cannot Delete Default DocumentLifecyclePolicy

In previous releases, you could delete any `DocumentLifecyclePolicy` object as long as the `DocumentLifecyclePolicy` was no longer being referenced by one or more `Document` instances. In the 4.0 release, you cannot delete a `DocumentLifecyclePolicy` if it is being referenced by one or more `Document` instances or if it is the default policy for a `Document` class.

## Document Classification

### COM-based Plug-ins Not Supported

Content Engine 3.5.x supported COM-based plug-ins for document classification actions. In the 4.0 release, classification actions must be Java-based. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. Any `DocumentClassificationAction` objects that your application uses must be modified to use a Java class instead of a COM class. For information on how to create a document classification action, see [Developing a custom classifier](#) in the *4.0 Help for Content Engine Administration*.) In addition, note changes in the following `DocumentClassificationAction` properties:

- **ProgId** - In 3.5.x, this property contained a Class Id or ProgId that represents the COM class. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- **ScriptText** - In 3.5.x, this property contained a string containing a script language program for implementing the classification behavior. If a value for ScriptText was set, then the ProgId property had to point to a Windows Scripting Host capable of interpreting that script. In 4.0, the text in this property is ignored and should be left empty.

### XMLPropertyMappingScript Algorithm

In 3.5.x, the algorithm used by the XMLPropertyMappingScript object was as follows:

1. Look for a Processing Instruction named FileNetDocClass. If it exists, use its value.
2. Else, look for a DTD with an external specifier. If it exists, use the value of the SYSTEM attribute (which usually contains a URL).
3. Else, use the name of the root element.

In 4.0, the algorithm is as follows:

1. Look for a Processing Instruction named FileNetDocClass. If it exists, use its value.
2. Else, get the expanded QName of the root element in the form of {RootElementNamespaceURI}RootElementLocalName. (The system looks for this in the XMLDocumentType property.)

If your existing application uses a FileNetDocClass Processing Instruction node in the XML content, no change is required for 4.0. Otherwise, you need to modify the value of the XMLDocumentType property of the XMLPropertyMappingScript object to use the new format:

```
{RootElementNamespaceURI}RootElementLocalName
```

For example, if the root element is:

```
<claim xmlns="http://www.filenet.com/claimSchema">
```

Then the expanded QName is:

```
{http://www.filenet.com/claimSchema}claim
```

### GUID Strings

In the 3.5.x release, operations that convert a GUID from a String to the internal form ignored any data after the closing } (brace). For example, the bolded text after the closing brace below would have been ignored:

```
{1321358E-DD13-4913-A44C-B57659A05C1E}, brown, filenet;cemp:t3: //localhost:7001/  
FileNet/Engine;null
```

In 4.0, this is no longer the case; any applications relying on this behavior must be updated.

## Type Property on Domain Class Not Supported

In the 4.0 release, the Type property on the Domain class is no longer supported. The Type property was meant to indicate the type of FileNet P8 domain (private, public, or workgroup), but public was the only type the Content Engine COM API ever supported.

## GCD Changes

In previous releases, the GCD (Global Configuration Data) was stored in the local file system. For the 4.0 release, the GCD is generally stored in a separate database instance, leveraging the RDBMS used by Content Engine.

During the upgrade operation from FileNet P8 3.5.x, the DateCreated and Creator properties for the following objects are replaced with the date of the upgrade and the name of the user performing the upgrade, respectively:

- ObjectStores
- MarkingSets
- FixedContentDevices

The DateCreated and Creator properties are not changed by the upgrade operation for objects that reside within an object store (such as Documents, Folders, and CustomObjects).

## Content Java API

This section describes changes that might be necessary to your 3.5.x Content Java API applications so they can run in a 4.0 environment. The 4.0 release includes a Content Java API Compatibility Layer, which is a client-side API that allows you to upgrade and maintain applications written using the 3.5.x Content Java API. The compatibility layer is designed to maximize support for the 3.5.x Content Java API in a 4.0 environment. However, in a few cases platform technologies diverged too greatly for an interface to be supported.

The compatibility layer is provided only to support existing applications, and can be installed during Content Engine installation. Documentation for the compatibility layer is only available on the [IBM Information Management support page on www.ibm.com](http://www.ibm.com). The documentation consists of reference help (javadocs) updated to reflect the compatibility layer as well as the existing 3.5.x Developer's Guide (not updated for 4.0.0).

## JDK

The JDK 1.3.x development environment is no longer supported; JDK 1.4.x is the supported development environment. Refer to the *IBM FileNet P8 Hardware and Software Requirements* for the IBM FileNet P8 4.0 release for complete information. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

## Removed from the API

The following API members, which were previously deprecated, have been removed from the API. FileNet P8 4.0.0.

## Fields

- All of the INHERITANCE\_TYPE\_\* Permission constants, use INHERITABLE\_DEPTH instead.
- The PRINCIPAL\_ID\_SPECIAL, RIGHT\_ARCHIVE, and RIGHT\_DEPLOY Permission constants.
- All of the ENABLE\_ON\_\* Property constants (these were constants for properties that were on the Subscription and WorkflowSubscription classes), use SUBSCRIBED\_EVENTS instead.
- All of the WCM\_\* Property constants.

## Methods

- ClassDescription.getAuditDefinitions(), use ClassDescription.getAuditDefinitions(inclInherited) instead.
- All forms of the getGroups and getGroupsXML methods on the Domain and Realm interfaces, use Realm.findGroups or Realm.findGroupsXML instead.
- All forms of the getUsers and getUsersXML methods on the Domain and Realm interfaces, use Realm.findUsers or Realm.findUsersXML instead.
- Permission.getInheritanceType and Permission.setInheritanceType, use Permission.getInheritableDepth and Permission.setInheritableDepth instead.
- EntireNetwork.getObjectStores and EntireNetwork.getObjectStoresXML, use Domain.getObjectStores or Domain.getObjectStoresXML instead.
- SecurityGrantee.getParentGroups(getAll), use SecurityGrantee.getParentGroups() instead.
- ObjectFactory.getSession(appId, credTag, userId, password, domain), use ObjectFactory.getSession(appId, credTag, userId, password) instead.
- The getUserAccess(collectionType, userId, domain) and getUserAccessXML() methods on the CustomObject, Document, Folder, Link, and SecurityPolicy interfaces, use the getUserAccess(collectionType, userId) or the getUserAccess() methods instead.
- ObjectStore.installFeatureAddOn(addOnInst, flags), use ObjectStore.installFeatureAddOn(addOnInst) instead.

## Deprecated

This section summarizes what has been deprecated in this release; sections below provide more information on these deprecations.

## Properties

The following properties (and their associated Property.<propertyname> constants) have been deprecated in this release. These properties are retained in the API for binary compatibility, but if used have no effect.

- ObjectStore.DefaultQueryBatchSize
- ObjectStore.DefaultQueryRowLimit

- `ObjectStore.EnumBatchSize`
- `ObjectStore.MaxQueryRowLimit`
- `Expression`
- `IsolatedRegion`
- `Priority`
- `StorageRepositoryType`

## NOTES

- Although the `Property.STORAGE_REPOSITORY_TYPE` constant was deprecated in this release, this should have no impact on existing applications since the classes on which this property existed were not exposed via the Content Java API.
- Although the `ObjectStore`, `ObjectType`, `OIID`, and `QueryOperatorDescriptions` properties are no longer supported in the Content Engine COM API, they continue to work in the Content Java API Compatibility Layer as the compatibility layer synthesizes these properties locally.

## Methods

In previous releases, requests from the application server to the Content Engine server traveled over an HTTP connection, and the Content Java API included `Session` methods that allowed you to supply additional HTTP header information. Due to differences in transport architectures between Content Engine 3.x and 4.0, the following `Session` methods have been deprecated:

- `Session.setProxyHost`
- `Session.getProxyHost`
- `Session.setProxyPort`
- `Session.getProxyPort`
- `Session.setTransportHeaders`

These methods are retained in the API for binary compatibility, but if used have no effect.

## Error Reporting

Content Engine 4.0 introduces a new error handling framework designed for the J2EE environment. A single unchecked exception class is used, with an associated class identifying the exception more specifically. The Content Java API Compatibility Layer maps 4.0 Content Engine API exceptions to appropriate 3.5.x Content Java API exceptions. In most cases, the same exception classes are thrown under the same circumstances in 4.0. There will be cosmetic differences in exception message text and in stack traces.

## Security

### Configurable JAAS Login Label

The Content Java API Compatibility Layer performs a JAAS login to accomplish authentication. In the 4.0 release, the default JAAS login label is "FileNetP8" and is case-sensitive. If appropriate, and as shown in the example below, the usual JAAS fallback label of "other" can also be used if there is no "FileNetP8" in your JAAS configuration:

```
other {
    weblogic.security.auth.login.UsernamePasswordLoginModule
        required debug=false;
};
```

The JAAS login label is configurable. To use a name other than the default "FileNetP8", you can edit the **WcmApiConfig.properties** configuration file with a text editor, setting or changing the value of the optional, case-insensitive `JAASConfigurationName` keyword. (Note that the format of the remote server URLs has changed in FileNet P8 4.0, as shown in the example below.) This `keyword=value` pair is shown in the contents of the **WcmApiConfig.properties** file below:

```
RemoteServerUrl=cemp:iiop://CEServer:2809/FileNet/Engine
RemoteServerUploadUrl=cemp:iiop://CEServer:2809/FileNet/Engine
RemoteServerDownloadUrl=cemp:iiop://CEServer:2809/FileNet/Engine
```

```
jaasconfigurationName=MyLoginLabel
```

```
CredentialsProtection=Clear
CredentialsProtection/UserToken=Symmetric
CryptoKeyFile/UserToken=C:\\Program Files\\FileNet\\Authentication\\
UTCryptoKeyFile.properties
```

Specifying the single value of "!" for the `JAASConfigurationName` keyword indicates that a JAAS login context is already in effect. Using "!" means that the Java compatibility layer will not perform the login step.

### Extensible Authentication Framework

The Extensible Authentication Framework was introduced in FileNet P8 3.5 to facilitate integration with Single-Sign-On (SSO) solutions that would not allow the transmission of a username and password to the FileNet system. FileNet P8 4.0 introduces a standards-based methodology for authentication using the J2EE Java Authentication and Authorization (JAAS) framework.

If you have implemented the Extensible Authentication Framework in your 3.5 application, you have two approaches you can take for the 4.0 release:

#### To implement a JAAS based solution

---

This is the preferred approach for Java environments (assuming your solution is in an SSO environment using an industry-standard solution, such as TAM or SiteMinder). This approach has the advantage of leveraging commodity SSO vendor components for all authentication plumbing.

Note, however, that this approach has the following considerations:

- It only works over the EJB transport.
- It requires authentication code in all client applications to be reworked, in addition to replacing server-side Extensible Authentication logic.

**To write a JAAS LoginModule using the FileNet Web Service Extensible Authentication Framework**

- If your client application is not running within a J2EE application server, this approach tends to be more complicated.
- If your client is a thin client, then a reverse proxy server is probably needed.
- If your client is Workplace, then in the 4.0.0 timeframe, you are limited to configurations that we have qualified for use with Workplace (such as SiteMinder + WebLogic 8.1 + Apache reverse proxy and TAM + WAS 6.0.2 + WebSeal reverse proxy).

If you use this approach, you must:

1. Install the SSO vendor's Login modules on the Content Engine server.
2. Install the SSO vendor's policy server in your environment (unless it is already in place).
3. Install the SSO vendor's LoginModules on all client machines.
4. Modify all applications to remove the legacy 3.5.2 Extensible Authentication logic that put custom credentials in the username and password fields.
5. Modify all applications so that they run over the EJB transport, and they perform a JAAS login using the SSO Login modules, prior to calling the server.

**To write a JAAS LoginModule using the FileNet Web Service Extensible Authentication Framework**

This is the preferred approach if a non-standard SSO environment is in place (or if web services transport is needed). This approach has the following considerations:

- It does not require any modifications to existing applications, as the existing client-side logic to submit custom values in the username and password fields will remain in place.
- There are no concerns if the client is not in a J2EE application server environment.
- It requires custom development work rather than leveraging commodity SSO vendor components.
- It will not work for Workplace in the 4.0 release, as we don't support Workplace over the web service transport in 4.0.

If you use this approach, you must write a custom JAAS LoginModule on the server, conforming to FileNet's 4.0 Web Service Extensible Authentication Framework (WS-EAF). This custom LoginModule will extract the username and password values from the UsernameToken in the incoming WS-Security header, and get the same custom values that were used in the 3.5.2 EA solution, and then invoke whatever custom logic is required to authenticate based on those parameters. For more information, see the [IBM FileNet P8 help topics under Developer Help > Web Service Extensible Authentication Framework Developer's Guide > Introduction](#).

**Setting a Marking Property**

In 3.5.x, if you attempted to set a property on a Marking object to which you did not have Add rights (Permission.RIGHT\_ADD\_MARKING), the Content Java API would throw a BadPropertyValueException. In 4.0.0, an InsufficientPermissionException is thrown.

## Directory Service Lookups

In FileNet P8 3.5.x, LDAP directory lookups could be done using a system configured account, or the user context. In FileNet P8 4.0, directory service lookups based on the user context are no longer supported. If you have an existing application that enumerates users, more results might be returned in the 4.0 release (since lookups done using the user context would have returned only those users and groups to which the user had access).

## Logon Names for Windows Active Directory

In 3.5.x, a logon (principal) name that contained an embedded @ character could successfully log on to a Content Engine server. In 4.0, logon names are no longer allowed to have any @ characters embedded in them because anything after the @ symbol is considered to be the domain name.

## Pre-Windows 2000 Login Format

In 3.5.x, the Content Engine server supported a pre-Windows 2000 login format (domain\login-name) through Microsoft-specific authentication APIs. The functionality is no longer directly available through the J2EE application server-based authentication modules. You are limited to what the application server supports (or the independent authentication provider that it uses) in terms of authentication login-name formats.

## SecurityParent and SecurityFolder Properties

The SecurityParent property (on the Document and CustomObject classes) has been deprecated. If the SecurityParent property is set (by an existing application, for example), the new SecurityFolder property is assigned the value of the SecurityParent property. When upgrading to this release:

- Objects with the SecurityParent property set use the tail (Folder) of the SecurityParent ReferentialContainmentRelationship object to set the SecurityFolder property to the SecurityParent value. Any access control entries in the object that were persisted in the object's ACL with a source type of PARENT are removed.
- Objects having superseded versions that do not have a SecurityParent property set, will have a SecurityFolder property set to the SecurityParent property value for the version that had a security parent.

## Changed Handling for UpdateSecurityEvent

If an object is subscribed to the UpdateSecurityEvent event and the security for an object changes, the system generates an UpdateSecurityEvent event. In prior releases, modifying the security of a folder could also modify the security of the documents contained in that folder. The applicable access control entry would be physically copied, as inherited, from the parent folder to the child document. Because this modified a document's security descriptor, an UpdateSecurityEvent event was generated for each child document.

The new dynamic security inheritance model does not copy the access control entries. Therefore, the system no longer generates the UpdateSecurityEvent for each impacted (inheriting) object. In



the case of a parent folder and multiple child documents, only a single UpdateSecurityEvent event is generated for the security modification performed on the folder.

An update to any object-valued property designated as "inheritable" in the metadata will also generate an UpdateSecurityEvent event.

## Transactions

Using the Content Java API Compatibility Layer, a Content Engine 3.5.x release caller can request transactional behavior only for batch processing. All methods of the batch are performed atomically within a new transaction context.

The caller view of a batch operation consists of the following sequence of actions:

1. Initiating the batch using `Session.startBatch`.
2. Accumulating work items into the batch (using standard method calls).
3. Executing the batch using `Session.executeBatch`.

It is the `executeBatch` method that is performed atomically. Callers declare their intentions about transactional behavior via a boolean parameter on the `startBatch` call.

In summary, the Content Java API Compatibility Layer calls are executed without propagating transaction context. The single exception is the `executeBatch` method associated with a `startBatch(asTransaction == true)` call, which is executed as its own complete transaction.

## Limitations

The individual work items of a Content Java API Compatibility Layer batch are performed on the client side (in a JVM other than that used by the Content Engine server). This imposes the following limitations:

- **Getting the InitialContext**

The first step in JNDI lookups is instantiating the `javax.naming.InitialContext` class. However, because the Content Java API Compatibility Layer has no a priori knowledge about the JNDI environment, you might need to use the **WcmApiConfig.properties** file to provide an explicit set of JNDI environment properties. See ["TxJndiProperties" on page 594](#).

If the `InitialContext` class cannot be instantiated, the Content Java API Compatibility Layer will throw an exception.

- **Finding the Transaction**

The Content Java API Compatibility Layer will also use a JNDI lookup to locate the `javax.transaction.UserTransaction` object. Because different application servers have differing standards for the key to be used in the lookup, you might need to specify the `InitialContext` lookup key value in the **WcmApiConfig.properties** file. See ["TxJndiKey" on page 594](#).

If a `UserTransaction` object cannot be located and instantiated, the Content Java API Compatibility Layer will throw an exception.

- **Transport**

Only the FileNet P8 4.0 EJB transport supports transaction propagation; the FileNet P8 Web Services transport does not. The Content Java API Compatibility Layer will call the Content Engine API method, `Connection.setParameter`, passing `ConfigurationParameter.CONNECTION_PARTICIPATES_IN_TRANSACTION` set to true.

If the Connection does not support transaction propagation, the Content Java API Compatibility Layer will catch that exception and rethrow it wrapped in an exception.

- **Applets and Applications**

The J2EE standard does not require that application servers support transactions in applet or application containers, although they can.

In cases where transactions are not supported, the Content Java API Compatibility Layer will catch the exception from JNDI (or other method calls) and rethrow it wrapped in an exception.

- **Active Transaction**

Because it does not interact directly with the J2EE transaction manager, the Content Java API Compatibility Layer cannot suspend a currently active transaction.

In cases where a transaction is currently active, the Content Java API Compatibility Layer will catch the exception from `UserTransaction.begin` and rethrow it wrapped in an exception.

### Transaction Entries in the `WcmApiConfig.properties` File

The 3.5.x Content Java API is configured via the `WcmApiConfig.properties` file. Although the `Session` interface has methods for setting or accessing some specific configuration values, those related to transaction processing can be set only in the `WcmApiConfig.properties` file. The transaction-related entries are as follows:

- **TxTimeout**

Allows you to specify an explicit value for the transaction timeout (in seconds). If this is not specified, the Content Java API Compatibility Layer uses a default of 60 seconds.

- **TxJndiKey**

Allows you to specify an explicit key for the call to `InitialContext.lookup(key)` for locating a `UserTransaction` reference. If this is not specified, the following ordered list is used until a lookup succeeds or the list is exhausted:

```
"java:comp/UserTransaction"
```

```
"javax.transaction.UserTransaction"
```

```
"UserTransaction"
```

- **TxJndiProperties**

Allows you to specify explicit JNDI property values. This setting is required only in those circumstances where it is not possible for you to configure JNDI properties using the standard methods.

The string value is reparsed as a `java.util.Properties` collection. Use newline escaping when specifying this value (as described in the Javadocs for `java.util.Properties.load`). For example, `"TxJndiProperties = p1=v1\n p2=v2\n p3=v3"`.

This `Properties` collection is passed unmodified as the environment when the JNDI `InitialContext` is instantiated. Only the `Context.PROVIDER_URL` property (`"java.naming.provider.url"`) gets special treatment:

- If the `Context.PROVIDER_URL` property is not present (including the case where the `TxJndiProperties` value is not specified), the property is added with a value of the URL being used by the Content Java API Compatibility Layer to connect to the FileNet P8 4.x Content Engine.
- If the `Context.PROVIDER_URL` property is present, and has the literal value `"!"` (exclamation mark), this indicates that the Content Java API Compatibility Layer should not override the default provider URL supplied by the application server. (The `Context.PROVIDER_URL` property containing `"!"` is removed from the `Properties` collection.)
- If the `Context.PROVIDER_URL` property is present and has any value other than `"!"` (exclamation mark), it remains unchanged.

## Events and Subscriptions

### Event Actions

Content Engine 3.5.x supported scripts and COM objects for event actions. In the 4.0 release, event actions must be Java-based. Any `EventAction` objects that your application uses must be modified to use a Java class instead of a script or a COM object. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. (For information on how to create a Java-based event action, see the [IBM FileNet P8 help topics under Content Engine Administration > Events and subscriptions > How to... > Work with event actions](#).)

In addition, changes in the following `EventAction` properties might require modifications in your application:

- `ProgId` - In 3.5.x, this property contained a string (such as `VBScript`) that specified the type of code to be executed when a Content Engine event occurred. In 4.0, this property specifies the fully qualified name of the Java class that executes when the event occurs.
- `ScriptText` - In 3.5.x, this property contained the script text of event procedures that executed when an action-related object's events were raised. In 4.0, the text in this property is ignored and should be left empty.

### EnableOn\* Subscription Properties Removed

The `EnableOn*` properties of the `Subscription` class and subclasses were deprecated in a previous release. These properties have been removed in the 4.0 release, and attempting to access them will generate exceptions. If your existing code uses the properties, convert your code to instead use the `SubscribedEvents` property.

## Priority Property Deprecated

In previous releases, you could specify a value used to order the firing of subscriptions in response to an event via the Priority property on a Subscription object. However, the server did not use this property, and therefore the property is no longer implemented in the 4.0 Content Engine release, and its associated Property constant (PRIORITY) has been deprecated in the Content Java API Compatibility Layer.

## Workflow Subscription Filter Expressions

In 3.5.x, a Process Engine filter expression could be supplied with a workflow subscription using the Expression property. This type of filter expression was evaluated by the Process Engine after the workflow event was launched. Because the filter expression evaluation occurred well after the original transaction was complete, this capability was inefficient and often inaccurate. In the 3.5.x documentation, IBM recommended that you use the FilterExpression property instead, as the FilterExpression property is evaluated within the transaction that initiates the event (before an event is queued).

In 4.0, the Expression property is no longer supported and its Property constant (EXPRESSION) has been deprecated. During the upgrade to Content Engine 4.0, if the Expression property contains a value, then the value will be moved to the FilterExpression property (but only if the FilterExpression property was not populated with a value). If both properties were populated with a value, then the upgrade tool will not move the contents of the Expression property but will produce a report of these cases.

**NOTE** The Expression property supported certain SQL syntax that is not supported in the FilterExpression property. Specifically, with the exception of the LIKE operator, all other operators requiring a database query (such as the EXISTS operator) are not supported in the FilterExpression property. For a list of unsupported operators, see [FilterExpression Property](#) in the *4.0 Content Engine Java and .NET Developer's Guide*.

If your application uses the Expression property, you must modify your code to use the FilterExpression property instead. In addition, if any 3.5.x Expression values included a SQL operator that is not supported in the FilterExpression property, you must resolve these instances by either programmatically setting the value of the FilterExpression property to a string that includes only supported operators or by using Enterprise Manager to correct the value.

## IsolatedRegion Property Deprecated

In the 4.0 release, the IsolatedRegion property on the ClassWorkflowSubscription and InstanceWorkflowSubscription classes has been renamed to IsolatedRegionNumber. The old property (IsolatedRegion) is deprecated, but existing code that accesses this property will continue to work.

## Query

### Content Based Retrieval

Content Engine 3.5.x integrated the Verity VDK for CBR. For Content Engine 4.0, Autonomy K2 (the former Verity product) is integrated for CBR. The following changed search behaviors are a result of integrating Autonomy K2 as the CBR engine:

- When an individual property is identified for CBR, only the specified property is searched.
- You can now do a simultaneous search for text in properties that is different for text in content.
- Search no longer supports identification of a content element.

## Query Results

If your application uses the Search or StoredSearch interfaces to perform a query that retrieves all information from a Content Engine 4.0 table (using SELECT \*), a much larger result set will be returned than in previous Content Engine releases. This is because previously, many system properties were marked as not selectable. In Content Engine 4.0, all properties are selectable. Some of the newly selectable properties can be quite large; in particular, any properties of type ListOfObject (such as Permissions or PropertyDefinitions). This extra data can cause an increase in response times for queries. For any application that traverses all properties in a result set, this can cause a further increase in response time, and could change behavior if new, unexpected data types are encountered. Queries that specify an explicit list of columns in their SELECT clause will not be affected.

## Using MAXRECORDS and TOP N

In 3.5.x, you could specify both MAXRECORDS as a query option and TOP N as part of the query. In 4.0, doing this will result in an invalid SQL statement because the Content Java API compatibility layer converts MAXRECORDS into a TOP N.

## Query-related ObjectStore Properties

The following query configuration-related properties on an ObjectStore are no longer supported:

- EnumBatchSize
- DefaultQueryBatchSize
- DefaultQueryRowLimit
- MaxQueryRowLimit

In the 4.0 release, the functionality supplied by these properties has been replaced with properties on the new ServerCacheConfiguration class:

- QueryPageMaxSize
- QueryPageDefaultSize
- NonPagedQueryMaxSize

Any existing code that accessed the ObjectStore properties will no longer function. The new properties are not accessible from the Content Java API.

## Auditing

In previous releases, attempting to modify properties for which the user did not have appropriate access rights was recorded as an audit failure by the Content Engine. However, in the 4.0 release the Content Java API Compatibility Layer will sometimes throw an exception on the client, and

therefore will not record the failed attempt in the audit log. (Only operations executed on the Content Engine server can be audited; failed or successful client operations are not audited by the server.)

## Document Lifecycle and Classification Actions

Content Engine 3.5.x supported Microsoft ActiveScript technology for document lifecycle actions and document classification actions. In the 4.0 release, lifecycle and classification actions must be Java-based. Furthermore, the Java class must be located on the server; it cannot be on a client and streamed to the server for execution as for CE 3.5.x. Although any 4.0 Content Engine actions must be rewritten to use Java, there are no code changes needed to existing 3.5.x Content Java API applications. For information on how to create a document classification action, see IBM FileNet P8 help topic [Content Engine Administration > Automatic document classification > Concepts > Understanding automatic document classification > Developing a custom classifier](#).

## Feature AddOns

Content Engine 3.5.x supported Microsoft ActiveScript technology for pre-install/post-install scripts for Feature AddOns. In Content Engine 4.0, only JavaScript scripting is supported. If your application creates FeatureAddOn objects and uses any pre-import or post-import scripts written in VBScript, you must:

- Rewrite those scripts in JavaScript.
- In your existing code to create a FeatureAddOn instance, specify `FeatureAddOn.SCRIPT_TYPE_JAVA` in the `scriptType` parameter in the call to `Domain.createFeatureAddOn`. For example:

```
FeatureAddOn oPubAddOn = oDom.createFeatureAddOn("newAddOn1",
FeatureAddOn.TYPE_RECOMMENDED, "XMLManifest 00 11", "preImp 00", "postImp 01",
FeatureAddOn.SCRIPT_TYPE_JAVA, sPredArr);
```

## ServerEnvironment Class

The 3.5.x Content Java API includes the `ServerEnvironment` utility class, which is used to gather information for FileNet Customer Support for configuration problems with the Application server or Content Engine server. Although this class remains in the Content Java API Compatibility Layer, differences in Content Engine platform technologies from 3.5.x to 4.0 have made this class unsupported in Content Engine 4.0 and in the compatibility layer. Methods called on this class will return exceptions.

## MimeType

In previous releases, when a Content Java API application did not supply the `MimeType` for a document's or annotation's `ContentTransfer` element, the 3.x Content Engine used file extension mappings in the Windows registry along with Microsoft APIs to examine the `ContentTransfer` element's content. Content Engine 4.0 identifies the `MimeType` based on a fixed listing of file extensions. Because of this, the resulting (identified) `MimeType` can be different from that identified by Content Engine 3.5.x. To ensure that a `ContentTransfer` element's `MimeType` is identified properly, explicitly specify the `MimeType` in the passed-in instance of the `TransportInputStream` when calling **setContent** on a Document or Annotation object.

## VersionSeries Security

In previous releases, if a user attempted to retrieve a VersionSeries object for which he did not have access rights to the current version, the VersionSeries object would be returned. In the 4.0 release, if the user does not have access rights to the current version, a null VersionSeries object will be returned.

## PropertyDescription DateTime Values

In 3.x releases, if you specified a date earlier than 1 Jan 1753 (which was the constraint specified by the PropertyMinimumDateTime property on the PropertyDescriptionDateTime class) for a custom property, the Content Engine server silently changed the date to 1 Jan 1753. In the 4.0 release, the PropertyDescriptionDateTime class has been updated to reflect the database-specific limits on date/time values. System and custom property values exceeding the database constraints will generate an exception. If your 3.5.x application relied on the above-mentioned PropertyMinimumDateTime behavior, you should modify your application to either pre-validate the date/time values, or to catch the BadPropertyValueException exception.

## Java SecurityManager

Use of java.lang.SecurityManager, which is a class that allows applications to implement a security policy, is not supported for the Content Java API Compatibility Layer (or for the new 4.0 Content Engine Java API). Access to Java system properties and other resources and capabilities is controlled by the Java SecurityManager, and the FileNet Java APIs might periodically need to read or update those resources.

## Exception When Attempting to Modify Root Folder Properties

In previous releases, attempting to move a root folder generated a ContentEngineUnsupportedOperationException. In the 4.0 release, an InsufficientPermissionException is thrown. (Note that the InsufficientPermissionException is also thrown in 4.0 if you attempt to change a root folder's name or delete a root folder.)

## Retrieving an ObjectStore's ClassDescription

In previous releases, if you retrieved the ClassDescriptions property for an ObjectStore, the ClassDescription for that ObjectStore would be included in the collection. In 4.0, this is no longer the case. However, as in 3.5.x, you can continue to retrieve the ClassDescription for an ObjectStore by retrieving the ClassDescription property on the ObjectStore object.

## Publishing

### Copy publishing

In previous releases, copy publishing (that is, publishing without doing a transform, which renders the content to a new format) was a synchronous process; in the 4.0 release, it is an asynchronous process. Applications that synchronously wait for the operation to complete will require modification. Also, less security will be required to complete the operation (similar to the security required for rendition publishing).



## PublishRequests

In the 4.0 Content Engine release, the PublishRequest class is a subclass of the new Subscribable class. Because of this, PublishRequest objects no longer inherit properties from the CustomObject class. In addition, some PublishRequest properties have been renamed or removed in the 4.0 release.

In previous releases, certain properties were inherited from the CustomObject class; in the 4.0 release, they are no longer on the PublishRequest class. Although it is unlikely that any PublishRequest code used these properties, you must modify any existing code that attempts to access or manipulate these properties as these attempts will result in an exception being thrown. The following table identifies these properties:

Properties No Longer Available on the PublishRequest Class		
AccessMask	LockTimeout	OIID
ActiveMarkings	LockToken	Owner
Annotations	Name	SecurityParent
Containers	ObjectStore	SecurityPolicy
LockOwner	ObjectType	WorkflowSubscriptions

In addition, the TargetDocument property is no longer supported. In previous releases, this property was used as a placeholder for saving the resultant publication object while it was being created during the copy publishing process (which was a synchronous process). You could retrieve the resultant target document when the publish operation completed. For example:

```
PublishRequest pr = document.publish(publishTemplate, pubOptions);
Document target = (Document) pr.getPropertyValue(Property.TARGET_DOCUMENT);
```

However, in the 4.0 release the publishing process does not require this placeholder as the copy publish process is handled asynchronously. Therefore, any existing code that attempts to access this property must be modified.

The following table identifies PublishRequest properties that have been renamed in this release. You must modify any existing code to use these new names. Note that the Content Java API Compatibility Layer does not include Property constants for these new names. Also, although



existing Property constants for the old names have not been deprecated, attempting to use them will throw an exception.

PublishRequest Property	New Name for Property in 4.0.0 Release
ErrorCategory (of type Integer)	ErrorCode (of type String) (see NOTES)
Options	PublishRequestType
PublishingPluginServer	DequeueHost
RetryNumber	RetryCount
Status	PublishingStatus (see NOTES)
StyleTemplate	PublishStyleTemplate

## NOTES

- In the 4.0 release, a publishing error is reported through the ErrorCode and ErrorDescription properties. ErrorCode contains the main error message string that describes the error, while ErrorDescription contains (if any) the cause of the error.
- If your existing code uses the Property.STATUS constant to access the Status property, your code will continue to work. However, if your existing code retrieves the Status property by name, then your code will need to be updated to use the new name, PublishingStatus.

## GCD Changes

In previous releases, the GCD (Global Configuration Data) was stored in the local file system. For the 4.0 release, the GCD is generally stored in a separate database instance, leveraging the RDBMS used by Content Engine.

During the upgrade operation from FileNet P8 3.5.x, the DateCreated and Creator properties for the following objects are replaced with the date of the upgrade and the name of the user performing the upgrade, respectively:

- ObjectStores
- MarkingSets
- FixedContentDevices

The DateCreated and Creator properties are not changed by the upgrade operation for objects that reside within an object store (such as Documents, Folders, and CustomObjects).

## Process Java API

### Process Router

The Process router used in previous releases was configured using the Process Task Manager, and the configuration information was stored in a local XML file. The Process Engine API

referenced the router via the RMI registry. For the 4.0 release, the Process router has been replaced by Connection Points. The connection points are implemented in the Content Engine API classes PEConnectionPoint and IsolatedRegion, configured using Enterprise Manager, and stored in the GCD. When logging on to a Process Engine session, existing Process Engine applications will need to supply a connection point name (instead of a router name) and the Content Engine URI referencing the connection point.

Note that the Process Engine Web Service requires only the connection point name.

## Directory Services and Authentication

In FileNet P8 3.5.x, LDAP directory lookups were done using a system configured account. This is no longer the case. The Process Engine now uses JAAS and the Content Engine Java API to perform authentication and directory service lookups.

## Content Engine API Files

Because the Process Engine uses the Content Engine API for authentication and directory service access, applications being upgraded for FileNet P8 4.0.0 will need to have the Content Engine API JAR files in their classpath. IBM recommends that you use the Web Services transport to connect your custom Process Engine applications to the Content Engine. However, for applications that use the J2EE framework with facilities specific to an application server (such as JMS), the EJB transport should be employed.

If your applications use the Web Services transport, follow the instructions below to configure the applications to run in a 4.0.0 environment. If your applications use the EJB transport, refer to the transport information in the Getting Started section of the Content Engine Java and .NET Developer's Guide, or to "Configuring the Component Integrator" in the Configuring the Process Development Environment section of the Process Engine Developer's Guide.

### To configure the applications for the 4.0.0 environment (Web Services transport only)

---

1. Copy the following directories from the CE\_API directory on the Process Engine server. These are located under the fnsw directory (for example, `c:\fnsw\CE_API`):

```
CE_API
  \config
  \lib
  \wsi
```

The other directories are not needed.

This directory is referenced in the following instructions as `<CE_API>`.

2. Add the following jar files to the classpath of the Process Engine applications:

```
<CE_API>\wsi\lib\wasp.jar
<CE_API>\lib\Jace.jar
```

Note that wasp.jar will need to precede the (existing) reference to pe3pt.jar.

3. Add the following JVM system properties to the command line of the Process Engine applications:

-Djava.security.auth.login.config=<CE\_API>\config\jaas.conf.WSI

-Dwasp.location=<CE\_API>\wsi

4. Specify the Content Engine URI.

The Content Engine server URI needs to be provided to the Process Engine API for the lookup of the connection points. Process Engine applications can use the `VWSession` methods to specify this or, for minimal code change, do one of the following:

- Put the `WcmApiConfig.properties` file in a directory that is in the classpath of the Process Engine application. This `WcmApiConfig.properties` file should, at a minimum, contain the following line:

```
RemoteServerUrl = cemp:http://CEServerName:CEServerPort/wsi/FNCEWS40DIME/
```

- Add the JVM system property, `filenet.pe.bootstrap.ceuri`, and specify the CE URI as follows:

```
-Dfilenet.pe.bootstrap.ceuri=http://CEServerName:CEServerPort/wsi/  
FNCEWS40DIME/
```

## Web Application Toolkit

Developers should be aware of the following information related to the backward compatibility of custom applications built with the Web Application Toolkit.

### Upgrading Custom Applications to Web Application Toolkit Version 4.0

When using JAAS container-managed authentication, Workplace does not have access to full user credentials and therefore, does not generate user tokens since it has nothing to encrypt as the basis for the token. As a result, in situations where Workplace might expose an action that invokes some piece of UI hosted by, for example, Records Manager (RM), and container-managed authentication is in use, the user will be presented with a second sign-in page when launching the RM-based action. To avoid the second sign-in page, the site must install and configure an SSO solution which offers JAAS support, such as Netegrity SiteMinder.

**NOTE** In FileNet P8 4.0.0, an Application Integration client cannot participate in a JAAS-integrated SSO deployment. The client does not support integration to leverage SSO authenticated user credentials in this type of deployment. Although no workaround is currently available for this issue, it will be addressed in a future release of FileNet P8.

### Upgrading Custom Applications to Web Application Toolkit Version 3.x and Later

Two settings, `windowIdCompatibility` and `internalTokensEnabled`, are provided in Web Application Toolkit 3.x and later to assist with backward compatibility. By default, if these settings are *not defined*, the assumed behavior will be for maximum backward compatibility. Therefore, the main reason for a custom application to define these settings would be to set them such that features introduced in 3.x and later can be leveraged (assuming that the 3.x behavior does not negatively impact the existing application).

When `windowIdCompatibility` is set to `false`, the `windowId` features of Toolkit version 3.0 are activated. These include support for `windowId` expiration and `CREATE_INLINE` window ID support. When `internalTokensEnabled` is set to `true`, 3.0 sign-in behavior is activated. In 2.x, SSL-based sign-in behavior used a mechanism where the SSL server would proxy directly into the browser/user's session on the non-secured server to set credentials. In 3.x and later, an internal token is requested from the non-secured server, eliminating the need to proxy into the user session, which also eliminates the need for origin session ID, IP, or port information by the SSL server. The internal token is then passed from the SSL server to the non-SSL server via a client-side redirect.

For applications that use a controller directly extended from `WcmController`, these settings are defined via the servlet descriptor, **WEB-INF/web.xml**. This would include existing applications developed against the 2.x Toolkit. In 3.x and later, the Workplace application is configured to use the 3.x `windowId` behavior and the token-based sign-in behavior. Therefore, if you copy **p8controller.xml** from the Workplace application to your custom application environment, you will have to change the default `windowIdCompatibility` and `internalTokensEnabled` settings if you want to maintain backward compatibility in your application.

The `ConfigurableController` controller class introduced in Web Application Toolkit 3.x extends and provides a full, configurable implementation of the abstract `WcmController`. For new or existing applications that wish to take advantage of `ConfigurableController`, the compatibility settings are set in the `ConfigurableController` configuration file, **WEB-INF/p8controller.xml**.

## Upgrading Custom Applications to Access Workflows on a P8 System 4.0

The following procedure shows you how to upgrade 3.5.x Toolkit-based applications to access workflows stored and processed on a 4.0 Content Engine and Process Engine. The procedure does not require modification of custom code, and it applies to Workplace as well as custom Toolkit-based applications.

**NOTE** The 3.5.x Process applets, for example, Administrator and Designer, are not supported in the 4.0 P8 environment.

The procedure assumes that:

- A custom application includes a **Site Preferences** page with a "Process Router - Host:port/name" option.
- The P8 system 4.0 is installed and configured.
- You know the name of the Process Engine Connection Point.
- You have access to the directory structure of the P8 4.0 Application Engine.

### To update the 3.5.x Toolkit-based application

---

1. Stop the 3.5.x application (custom or Workplace).
2. From the 4.0 Application Engine, copy the following JAR files in `<app>/WEB-INF/lib` to the corresponding directory in your 3.5.x Application Engine.
  - jace.jar
  - javaapi.jar
  - pe.jar
  - peResources.jar
3. In `<app>/WEB-INF/WcmApiConfig.properties` of the 3.5.x application, change the remote server settings to point to the 4.0 Content Engine.
4. Start the Toolkit-based application, and sign in as a user who is a member of the **Application Engine Administrators** group.
5. Click **Admin**, then click **Site Preferences**.
6. In the **General** preference view, scroll to the **Tasks** category.
7. In the "Process Router - Host:port/name" option, change the "vmrouter" setting to the Process Engine Connection Point name.
8. Apply your change.
9. Restart the Application Engine.
10. To verify that the upgrade works, go to the **Tasks** primary view. You should get a list of workflows from the P8 4.0 system.

## Obsolete

The APIs listed below, which were previously deprecated, are obsolete in the FileNet P8 4.0.0 release.

- WcmImageAnchor class
- WcmFieldUtil class
- WcmControllerBehaviorInterface interface

## FormProcessor Class

- doMapParameters(WorkflowPolicy workflowPolicy, FormData formData), use doMapParametersList method instead.

## WcmAuthoringDataProvider Class

- getRealmGroups(java.lang.String realmId, boolean refresh)
- getRealmUsers(java.lang.String realmId, boolean refresh)

## WcmController Class

- CREATE\_WINDOW\_ID
- FIRST\_WINDOWID\_KEY
- HOME\_KEY
- POPUP\_KEY
- PROPOGATE\_WINDOW\_ID
- WINDOW\_SIGNED\_IN\_KEY
- windowIdParams
- configurePage(javax.servlet.ServletContext applicationValue, javax.servlet.http.HttpServletRequest request, boolean windowIdRequired)
- configurePage(javax.servlet.ServletContext applicationValue, javax.servlet.http.HttpServletRequest request, int windowIdMode)
- configurePage(javax.servlet.ServletContext applicationValue, javax.servlet.http.HttpServletRequest request, int windowIdMode, java.lang.String[] windowIdParams)
- getHomeURLKey(java.lang.String windowId)
- getWindowIdFromReferer(java.util.Map popups, java.lang.String referer)
- isNewWindowIdRequired(javax.servlet.http.HttpServletRequest request)
- setAutoFixWindowId(boolean value)

### WcmCredentialsServlet Class

- getCredentialsRequestURL(WcmDataStore ds, java.lang.String remoteHost, boolean wasBugFix, java.lang.String originScheme, java.lang.String originIP, java.lang.String originPort, java.lang.String originSessionId, java.lang.String encodedSessionId, java.lang.String userId, java.lang.String password, java.util.Map extraParameters)

### WcmDateUtil Class

- convertDateToW3CDate(java.util.Date date, boolean bDateOnly)

### WcmEProcessDataProvider Class

- doEvaluateExpression(java.lang.String subscriptionName, java.lang.String expression, java.lang.String objectStoreId, java.lang.String objectId, java.lang.String verSerId, int objectType)
- findGroups(java.lang.String searchPattern, int searchType, int sortType, int maxBufferSize)
- findUsers(java.lang.String searchPattern, int searchType, int sortType, int maxBufferSize)

### WcmException Class

- loadResource(java.io.InputStream in)
- loadResource(java.util.Map m)

### WcmGetContentServlet Class

- handleExternalDocument(com.filenet.wcm.api.Document doc, javax.servlet.http.HttpServletRequest response, WcmDataStore ds)

### WcmIcons Class

- load(java.lang.String sFilePath)

### WcmModule Class

- getClassPropertyKeys()
- getModulePropertyKeys()

### WcmSearchDataProvider Class

- checkForChoiceLists(org.w3c.dom.Document searchDefinition)
- checkForUserListProperties(org.w3c.dom.Document searchDefinition)
- getAllObjectStores(boolean refresh)
- getChoiceList(java.lang.String objectStoreName, java.lang.String searchID)
- getSearchRequest(java.util.Set selectProperties, java.util.Map propertyItems, java.util.Map verityItems, java.util.Set excludedClasses, int maxRecords)

### WcmServerCredentials Class

- getSessionToken(WcmDataStore ds, java.lang.String appld, java.lang.String user, java.lang.String password, java.lang.String domain)
- getUserDomain()

### WcmSignInPolicy Class

- createServerCredentials(WcmDataStore ds, java.lang.String appld, java.lang.String user, java.lang.String password, java.lang.String domain)
- getEncodedSessionId()
- getLoginRoutingInfo()

### WcmSignInPolicyInterface Class

- getLoginRoutingInfo()

### WcmSSLInfo Class

- getSslHost()
- getSslHostDecoded()

### WcmString Class

- loadResource(java.io.InputStream in)
- loadResource(java.util.Map m)

### WcmStringEditor Class

- encode(java.lang.String str)

### WcmStringResourceLoader Class

- load(java.lang.String path, java.lang.String name, java.lang.Class resourceClass)

### WcmStringResources Class

- addLegacyPRB(boolean exception, java.io.InputStream in)
- addLegacyPRB(boolean exception, java.util.Map m)

### WcmXSLUtil Class

- setXSLParameter(java.lang.String key, org.w3c.dom.Node node)
- setXSLParameter(java.lang.String key, java.lang.String value)
- setXSLParameter(java.lang.String key, org.w3c.dom.NodeList nodes)



## ***Workplace Application Integration Toolkit***

In FileNet P8 4.0.0, an Application Integration client cannot participate in a JAAS-integrated SSO deployment. The client does not support integration to leverage SSO authenticated user credentials in this type of deployment. Although no workaround is currently available for this issue, it will be addressed in a future release of FileNet P8.

The following Workplace Application Integration Toolkit members are now obsolete:

### **Obsolete**

The APIs listed below, which were previously deprecated, are obsolete in the FileNet P8 4.0.0 release.

- ICheckinCmd interface, use ICheckinCmd2 interface instead.
- ISelectObjectCmd interface, use ISelectObjectCmd2 interface instead.
- ISessionLogin interface, use IStdSessionLogin interface instead.
- ISessionLogin::Initialize method; IStdSessionLogin::InitializeUnifiedLogin and IStdSessionLogin::InitializeViaSessionName methods should be used instead.

## Task 18: Upgrade Server-Side Scripts and COM Objects

As with Content Engine 3.5.x, you can develop and plug in the following server-side components for the 4.x version of Content Engine: event action handlers, lifecycle action handlers, and document classifiers. Whereas these components had to be implemented as COM objects or scripts for Content Engine 3.5.x, they must be implemented as Java classes in Content Engine 4.x. Therefore, for Content Engine 4.x, you must convert any event action handlers, lifecycle action handlers, and document classifiers that you developed for Content Engine 3.5.x.

This section gives you an idea of the effort involved in converting 3.5.x script-based event action handlers to 4.x Java-implemented handlers. For details on developing and deploying event action handlers, lifecycle action handlers, and document classifiers, see the *Java and .NET Developer's Guide* and the *Java API Reference*.

In the following listings, compare the 3.5.x script-based event action handlers to the corresponding Java-implemented handlers required for Content Engine 4.x.

### ***File Document Handler***

This handler files a document to a specified folder, determined by the event fired on the document.

**Content Engine 3.5.x JScript Version**

```
function OnEvent (Event, Subscription)
{
  var doc = Event.SourceObject;
  if ( Event.IsOfClass("CreationEvent") )
  {
    FileDocInFolder("/Docs", doc);
  }
  else if (Event.IsOfClass("ChangeClassEvent") )
  {
    FileDocInFolder("/Archives", doc);
  }
}

function FileDocInFolder(otherFolderName, doc)
{
  var os = doc.ObjectStore;
  var rootFld = os.RootFolder;
  var fldSet = new Enumerator(rootFld.SubFolders);
  var subFld;
  for ( ; !fldSet.atEnd(); fldSet.moveNext() )
  {
    subFld = fldSet.item();
    if (subFld.Name == otherFolderName)
    {
      subFld.File(doc, 0, doc.DocumentTitle);
    }
  }
}
```

**Content Engine 4.x Java Version**

```
import com.filenet.api.constants.*;
import com.filenet.api.constants.DefineSecurityParentage;
import com.filenet.api.core.*;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class FileDocumentAction implements EventActionHandler
{
    public void onEvent(ObjectChangeEvent event, Id subscriptionId)
        throws EngineRuntimeException
    {
        Document doc = (Document)event.get_SourceObject();
        try
        {
            if (event.getClassName().equalsIgnoreCase("CreationEvent"))
                FileDocInFolder("/docs", doc);
            else if (event.getClassName().equalsIgnoreCase("ChangeClassEvent"))
                FileDocInFolder("/Archives", doc);
        }
        catch (Exception e)
        {
            throw new EngineRuntimeException(ExceptionCode.E_FAILED);
        }
    }

    public void fileDocInFolder(String folderName, Document doc)
    {
        try
        {
            Folder folder = (Folder)doc.getObjectStore().getObject("Folder", folderName);
            ReferentialContainmentRelationship rel = folder.file (doc,
                AutoUniqueName.AUTO_UNIQUE, doc.get_Name(),
                DefineSecurityParentage.DO_NOT_DEFINE_SECURITY_PARENTAGE);
            rel.save(RefreshMode.NO_REFRESH);
        }
        catch (Exception e)
        {
            e.printStackTrace();
        }
    }
}
```

## Log Event Handler

This handler records events to a log file when documents are created.

### Content Engine 3.5.x VBScript Version

```
Public Sub OnEvent (EventObject, Subscription)
  Dim doc, message)
  Set doc = EventObject.SourceObject)
  WriteToLogFile (doc.Name & " was created on: " & Date))
End Sub

Public Sub WriteToLogFile (message))
  Dim fso, ts, logFile)
  Set fso = CreateObject("Scripting.FileSystemObject"))
  Set logFile = fso.CreateTextFile("C:\log.txt"))
  logFile = nothing)
  Set ts = fso.OpenTextFile("C:\log.txt", 8, True))
  ts.Write (message))
  ts.WriteBlankLines (2))
  ts.Close)
  Set fso = Nothing)
  Set ts = Nothing)
End Sub
```

**Content Engine 4.x Java Version**

```
import java.io.File;
import java.io.FileWriter;
import java.io.IOException;

import com.filenet.api.core.Document;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ErrorRecord;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class LogEventAction implements EventActionHandler
{
    public void onEvent(ObjectChangeEvent event, Id subscriptionId) throws EngineRuntimeException
    {
        try
        {
            Document doc = (Document)event.get_SourceObject();
            WriteToLogFile(doc.get_Name() + " was created on: "
                + new java.util.Date() + "\r\n");
        }
        catch (Exception e) {
            ErrorRecord er[] = {new ErrorRecord (e)};
            throw new EngineRuntimeException(e, ExceptionCode.EVENT_HANDLER_THREW, er);
        }
    }

    public void writeToLogFile(String message)
    {
        try
        {
            File outputFile = new File("C:\\log.txt");
            FileWriter out = new FileWriter(outputFile, true);
            out.write(message);
            out.close();
        }
        catch (IOException e) {
            ErrorRecord er[] = {new ErrorRecord (e)};
            throw new EngineRuntimeException(e, ExceptionCode.EVENT_HANDLER_THREW,er);
        }
    }
}
```

## Send eMail Handler

This handler sends an email when a new document has been created.

### Content Engine 3.5.x VBScript Version

```
Public Sub OnEvent (EventObject, Subscription)
  Dim myMail, MessageBody
  Set myMail = CreateObject("CDONTS.NewMail")
  myMail.From = "user1@company.com"
  myMail.To = "sysAdmin@company.com"
  myMail.Subject = "Event Notification--New Document created"
  MessageBody = "A document titled "" & Source.DocumentTitle & "" was Created at " & time & " on " &
date & "."
  MessageBody = MessageBody + vbCrLf + Subscription.UserString
  myMail.Body = MessageBody
  myMail.Send
  Set myMail = Nothing
End Sub
```

**Content Engine 4.x Java Version**

```
import java.util.Date;
import java.util.Properties;

import javax.mail.Message;
import javax.mail.MessagingException;
import javax.mail.Session;
import javax.mail.Transport;
import javax.mail.internet.InternetAddress;
import javax.mail.internet.MimeMessage;

import com.filenet.api.core.*;
import com.filenet.api.engine.EventActionHandler;
import com.filenet.api.events.ObjectChangeEvent;
import com.filenet.api.exception.EngineRuntimeException;
import com.filenet.api.exception.ExceptionCode;
import com.filenet.api.util.Id;

public class EMailAction implements EventActionHandler
{
    public void onEvent(ObjectChangeEvent event, Id subscriptionId) throws EngineRuntimeException
    {
        Document doc = (Document)event.get_SourceObject();
        try
        {
            Properties props = new Properties();
            props.put("mail.smtp.host", "smtp.company.net");
            props.put("mail.smtp.port", "25");
            Session session = Session.getInstance(props);
            try {
                Message msg = new MimeMessage(session);
                msg.setFrom(new InternetAddress("user1@company.com" ));
                InternetAddress[] address = {new InternetAddress("sysAdmin@company.com" )};
                msg.setRecipients(Message.RecipientType.TO, address);
                msg.setSubject("Test E-Mail through Java");
                msg.setSentDate(new Date());
                msg.setText("Document " + doc.get_Name() + " created with ID " + doc.get_Id());
                Transport.send(msg);
            }
            catch (MessagingException mex) {
                mex.printStackTrace();
            }
        }
        catch (Exception e)
        {
            throw new EngineRuntimeException(ExceptionCode.E_FAILED);
        }
    }
}
```



## Task 19: Upgrade ISRA Servlet

.As part of upgrading FileNet Image Services Resource Adapter (ISRA) software to run with FileNet P8, you must do the following:

- Ensure that Application Engine has been upgraded.
- Ensure that existing FileNet ISRA software is installed and configured.

For information on installing, configuring, and deploying FileNet ISRA, refer to the ISRA documentation on the FileNet ISRA Installation package.

**TIP** Use the Sample Application shipped with FileNet ISRA to confirm that the ISRA installation was successful.

**WARNING** In an ISRA upgrade, take care to use the same library name (JNDI connection factory name) that has been previously set in the ISRA install. Changing this variable can cause conflicts when accessing documents.

- Do the following, as documented later in this task topic:
  - Upgrade the Application Engine ISRA Servlet, taking the following into account:
    - The servlet must be deployed on the same application server as FileNet ISRA.
    - The servlet does not need to be collocated with the Application Engine.
  - Configure Workplace site preferences for ISRA support.

### ISRA SSL Support

The following table details the supported for SSL configurations.

SSL Configuration	SSL Support
ISRA Servlet and AE Collocated. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet and AE Collocated. AE Configured for SSL logon redirect to a local host.	Supported
ISRA Servlet and AE Collocated. AE and ISRA Servlet running under SSL.	Not Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a non-local host.	Supported
ISRA Servlet remote from AE. AE configured for SSL logon redirect to a local host.	Supported
ISRA Servlet remote from AE. AE running under SSL, ISRA Servlet not running under SSL.	Supported

SSL Configuration	SSL Support
ISRA Servlet remote from AE. AE and ISRA Servlet running under SSL.	Not Supported

### To upgrade the Application Engine ISRA Servlet

The FileNet P8 Application Engine installation CDs contain the ISRA servlet installation programs for the supported P8 AE platforms.

1. Log on to the application server.

(UNIX) Log on as a user with write access to the **/bin** directory and read, write, execute access to the directory where you plan to install ISRA Servlet.

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

2. Stop the application server if it is running.
3. Access the ISRA installation package, and start the Application Engine ISRA Servlet Setup wizard:

(UNIX) - Execute **<Platform>filenet\_ae\_israservlet\_setup.bin**.

(Windows) - Execute **WINfilenet\_ae\_israservlet\_setup.exe**

4. Complete the Setup screens as follows:

In this screen...	Perform this action...
License Agreement	Review and accept the license agreement for FileNet P8 software, then click <b>Next</b> .
Directory Name	<p>For the Directory Name field, enter or browse to the location where you want to install the ISRA Servlet, or accept the default location:</p> <ul style="list-style-type: none"> <li>• UNIX - /opt</li> <li>• Windows - C:\Program Files\</li> </ul> <p>Click <b>Next</b></p> <p><b>UNIX</b></p> <pre>&lt;AE_israservlet_install_path&gt;/FileNet/ ApplicationEngineISRAServlet&gt;</pre> <p><b>Windows</b></p> <pre>&lt;AE_israservlet_install_path&gt;\FileNet\ApplicationEngineI SRAServlet&gt;</pre>

In this screen...	Perform this action...
Create war File	Select this check box if you use a WebSphere application server.  If the check box is selected, the installation program will create the <b>ae_isra.war</b> file and a script that can also generate the <b>ae_isra.war</b> file.
User Token Security	Make a note of the user token crypto key path. Click <b>Next</b> .
Ready to Install	Verify your selections, and click <b>Next</b> to install the ISRA Servlet.
Completing the Setup	Click <b>Finish</b> to complete the ISRA Servlet installation Wizard.

5. Check the file **filenet\_ApplicationEngineISRAServlet\_install\_log.txt**, located in the **<AE\_israservlet\_install\_path>\FileNet** directory, to see if any errors occurred during the installation.
6. Install unlimited strength jar files.  
  
Perform this step only if your site is generating or accepting unlimited strength user tokens. Your system must be configured as follows:
  - The Application Engine ISRA Servlet is deployed on a different application server from the Application Engine server.
  - The **Create unlimited strength key** option was selected in the Application Engine User Token Security step of the Application Engine installation.

**NOTE** Failure to perform the step will cause an EncryptionException when you log in to the IS Server.
7. (WebSphere 5.x only) To allow users to save annotations through ISRA, copy the **<AE\_israservlet\_install\_path>\FileNet\jar\ISRA.jar** file to the **<AE\_install\_path>\FileNet\AE\Workplace\WEB-INF\lib** directory.
8. Start the application server.
9. Deploying the Application Engine ISRA Servlet is similar to deploying Workplace.
  - WebSphere  
Deploy **<AE\_israservlet\_install\_path>\FileNet\ApplicationEngineISRAServlet\ae\_isra.war** in the same way you deployed the **app\_engine.war** file for Workplace.
  - WebLogic  
Deploy **<AE\_israservlet\_install\_path>\FileNet\ApplicationEngineISRAServlet** in the same way you deployed **<AE\_install\_path>\FileNet\AE\Workplace**.
10. Stop and restart the application server.
11. Verify the Application Engine ISRA Servlet installation. Do the following to use an available diagnostic tool to verify that the ISRA Servlet is installed and deployed correctly.

- a. Launch your browser.
- b. Enter the URL for the Application Engine ISRA Servlet. For example,

```
http://<ApplicationEngineISRAServlet_servername>:<port>/
ApplicationEngineISRAServlet/ISRA
```

**NOTE** ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

If the ISRA Servlet is installed and deployed correctly, a Congratulations message displays. For example:

```
Congratulations! ISRA Interface Servlet is configured at this URL.
WcmApiConfigFile = D:\ISRAInterface\jsp\WEB-INF\WcmApiConfig.properties
WcmApiConfig file exists
```

```
CryptoKeyFile/UserToken = C:\Program
Files\FileNet\Authentication\UTCryptoKeyFile.properties
CryptoKeyFile/UserToken exists
```

```
FileNet ISRA classes are in the classpath
com.filenet.is.ra.cci.FN_IS_CciConnectionSpec
```

### **To configure the Workplace site preferences for ISRA support**

---

Application Engine Setup installs a pre-configured external service called Image Service, which includes the parameterized values necessary to access FileNet Image Service libraries from Workplace. Enable the service by setting the *Image Service* value in Site Preferences to **Show** (the default is Hide), as described in the following procedure.

1. Sign in to Workplace as a user having the Application Engine Administrators access role.
2. Launch Site Preferences as follows:
  - a. Select Admin.
  - b. Select Site Preferences.
3. Select External Services from the left options list.
4. Select Modify for the Image Service (under External Reference Services).  
 The External Reference Service Settings site preference page displays.
5. Under General Information, locate *Show on Select File page* and change the value to Show.
6. Click **Accept**.
7. Click **Apply**.

### To set the ISRA Interface Servlet URL

---

1. Select **Bootstrap**.
2. Under Preferences Settings, set the value of ISRA Interface Servlet URL. For example:

`http://<servername>:<port>/ApplicationEngineISRAServlet/ISRA`

**NOTE** ApplicationEngineISRAServlet is the default context root. If you specified a different name for the context root when deploying the Application Engine ISRA Servlet, change the URL to match your configuration.

3. Click **Apply**.
4. Click **Exit** to exit the Site Preferences.

### To log on to Image Services via LDAP

---

To log on to the Image Services library using your LDAP account, configure ISRA and Image Services for LDAP authentication. If the LDAP account with which you accessed Workplace is not valid for the Image Services library, or if LDAP authentication is not configured, you will be prompted to log on to the Image Services library.

For information on configuring LDAP authentication for ISRA, refer to the *ISRA Installation and Deployment Guide*. For information on configuring LDAP authentication for Image Services, refer to the *Image Services System Tools Reference Manual*.

### To access IS library documents

---

For information about accessing IS library documents, see FileNet P8 Help topic [User Help > Actions, preferences and tools > Actions > Documents > Add a document \(Workplace\)](#).

## Task 20: Install Service Packs for Add-On Components

Install any service packs, fix packs and/or interim fixes required for the add-on components. To determine whether additional service packs, fix packs or interim fixes are needed, contact your service representative.

## Remove Software

This section includes:

- “Remove the IBM FileNet P8 Documentation” on page 624
- “Remove Content Search Engine” on page 626
- “Remove Content Engine” on page 628
- “Remove Process Engine (Windows)” on page 631
- “Remove Process Engine (UNIX)” on page 633
- “Remove Application Engine (WebSphere)” on page 634
- “Remove Application Engine (WebLogic)” on page 636
- “Remove Application Engine (JBoss)” on page 637
- “Remove the Application Engine ISRA Servlet” on page 639

For instructions on removing the Rendition Engine software, see the IBM FileNet P8 guide [FileNet P8 Documentation > FileNet P8 System Installation > Rendition Engine Installation and Upgrade](#).

## Remove the IBM FileNet P8 Documentation

The following topic explains how to remove the documentation for the FileNet P8 Platform and its expansion products.

**NOTE** Because of the number of application server configuration possibilities, these examples should be used only for reference. Your specific installation directories and application names may vary.

### To remove the FileNet P8 documentation

---

**NOTE** In some Windows installations where NTFS is used for the file system (Not FAT or FAT32), there is a known issue with deleting files (and folders) that are longer than 256 characters. For example, if you use a default WebSphere installation location, you may encounter an error where the FileNet P8 Platform documentation files cannot be properly deleted due to the number of characters in the file/folder names. See the Microsoft Knowledge Base article <http://support.microsoft.com/?kbid=320081> for additional information about deleting files (and folders) in this environment.

### To remove the FileNet P8 documentation from a WebSphere server

1. Log on to the WebSphere FileNet P8 documentation server.
  - UNIX - Log on as a user with write access to where the FileNet P8 Platform documentation files are installed.
  - Windows - Log on with a user account that has local Administrative (or Account Operators and Server Operators) rights.
2. Verify that the WebSphere server is running.
3. Log on to the WebSphere administrative console (for example, <http://localhost:9060/ibm/console>).
4. Expand Applications > Enterprise Applications.
5. Select the FileNet P8 documentation site (for example, *ecm\_help.war*).
6. Click **Stop**.
7. Click **Uninstall** and follow the remaining screen prompts.
8. Click **Save** and follow the remaining screen prompts.
9. Delete the entire FileNet P8 documentation folder (for example, *ecm\_help.war*) structure from the installation location.
10. Delete all folders and files including any **temp** folder(s) for the FileNet P8 documentation (for example, *ecm\_help*).

**WARNING** Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

### To remove the FileNet P8 documentation from a WebLogic server

1. Log on to the WebLogic application server.



- UNIX - Log on as a user with write access to where the **ecm\_help** files are located.
  - Windows - Log on with a user account that has local Administrative rights (or Account Operators and Server Operators).
2. Verify that the WebLogic server is running.
  3. Execute and log on to the WebLogic Server Administration Console (for example, <http://<machinename>:7001/console>).
  4. Expand Deployments > Web Application Modules.
  5. Right-click **ecm\_help** and click **Delete**.
  6. Click **Yes**, then click **Continue**.
  7. Delete all folders and files including any **temp** folder(s) for the FileNet P8 documentation.

**WARNING** Do not remove any other FileNet P8 application (for example, Workplace) files that are installed on the web application server.

## Remove Content Search Engine

To completely remove Content Search Engine (Autonomy K2) and collections from your IBM FileNet P8 platform installation, you must disable full-text indexing and remove the Autonomy K2 installations from all servers associated with the Content Engine.

**CAUTION** If you remove or disable Autonomy K2 before you disable CBR and full-text indexing in Enterprise Manager, your system will be rendered unusable and require considerable reconstruction.

### NOTES

- This procedure presumes you have a running installation of Content Search Engine and that you have existing collections.
- If you intend to remove Content Engine, skip to step 4.
- The paths listed in this procedure assume you have used the suggested install path. If you have installed to another location, substitute the path as appropriate.

### To remove Autonomy K2 and CBR

---

1. Launch Enterprise Manager.
2. Disable CBR for any classes that have been enabled for CBR.
  - a. Right-click the class you want to configure and click **Properties**.
  - b. Click the **General** tab.
  - c. Clear the **CBR Enabled** checkbox and click **OK**.
  - d. A dialog will ask if you wish to propagate this change to the subclasses of this class, click **Yes**.
  - e. Repeat this procedure to disable CBR for all classes.
3. Run an index job and re-index any of the following that were previously enabled:
  - Document
  - Annotation
  - Custom Object
  - Folder

This will disable all full-text indexing and content-based retrieval settings and will delete any associated collections. Once the indexing job is complete, proceed to step 4.

4. Log onto the K2 Master Administration Server host machine as an administrator.
5. Shutdown the Tomcat service:

Windows:

- a. From a command window, access **C:\Program Files\FileNet\ContentEngine\Verity\apps-server\bin**

- b. Enter the following command:

```
shutdown.bat
```

UNIX:

- a. Access **Opt/Verity/appser/bin**  
b. Enter the following command:

```
./shutdown.sh
```

6. (Windows only) Remove the K2 Dashboard service:

- a. From a command window, access **C:\Program Files\Filenet\ContentEngine\Verity\apps-erver\bin**  
b. Enter the following command:

```
service remove k2
```

7. Remove the Autonomy K2 installation:

Windows:

- a. From a command window, access **C:\Program Files\Filenet\ContentEngine\Verity**  
b. Enter the following command:

```
k2\_nti40\bin\vconfig -cfg "C:\Program  
Files\filenet\contentengine\verity\config.vcnf" -dir "C:\Program  
Files\filenet\contentengine\verity" -verbose -log log.txt -uninstall
```

UNIX:

- a. Access Opt/Verity/  
b. Enter the following command:

```
k2/<platform>/bin/vconfig -cfg "/opt/verity/config.vcnf" -dir "/opt/verity" -  
verbose -log log.txt -uninstall
```

Substitute one of the following for *<platform>*:

- *\_ssol26* (Solaris 8.0, 9.0 or 10.0)
- *\_hpux* (HP-UX 11i with -AA compiler flag)
- *\_rs6k43* (AIX 5.2 and 5.3)
- *\_ilnx21* (Red Hat Advanced Server 3.0 and 4.0, SUSE 8 and 9)

The Autonomy K2 Administration Server service and Tomcat will be uninstalled at the completion of the vconfig command.

8. Delete the install directory.

## Remove Content Engine

You can uninstall an entire Content Engine installation or selected Content Engine components.

**NOTE** Uninstalling Content Engine does *not* undeploy it. You must use the application server console or commands to remove the Content Engine EAR file from the application server.

Use one of the following procedures to uninstall part or all of Content Engine.

### To remove part or all of a Content Engine installation using the Windows Control Panel

---

1. Choose **Control Panel > Add/Remove Programs**.
2. Highlight Content Engine 4.0.0 in the list of currently installed programs and click **Remove** to launch Content Engine Setup.
3. Complete the Content Engine Setup screens as follows:

In this screen...	Perform this action...
Welcome	Click <b>Next</b> to proceed with removing the Content Engine installation.
Select Components	Select the check boxes of the components to be uninstalled, and then click <b>Next</b> .
Summary	Verify the list of components to be removed, and then click <b>Uninstall</b> .
Summary Information	Click <b>Finish</b> to exit Content Engine Setup.

4. If you want to completely remove all traces of the Content Engine installation, delete the **C:\Program Files \FileNet** directory.

### To remove an entire Content Engine installation (UNIX)

---

1. Navigate to the directory **\_uninst2**, created by the Content Engine installer.
2. To uninstall Content Engine interactively, run the following command:

```
uninstaller.bin
```

### To remove Content Engine silently

---

1. Record a response file, as follows:
  - (UNIX) `./uninstaller.bin -options-record <path_to_response_file>`
  - (Windows) `uninstaller.exe -options-record <path_to_response_file>`

**CAUTION** Do not specify a response file path within the Content Engine installation directory; if you do this, the file will be deleted.

2. To uninstall Content Engine silently, run one of the following commands:

- (UNIX) `uninstaller.bin -options <path_to_response_file> -silent`
- (Windows) `uninstaller.exe -options <path_to_response_file> -silent`

**To remove selected components of a Content Engine installation**

---

By default, the Content Engine uninstaller removes all of Content Engine. If there is an individual Content Engine component you want to retain, specify it by either checking the components on the Select Components dialog box or as a command-line option. For example, in a UNIX environment, the command line is as follows:

```
uninstaller.bin -P <compID>.activeForUninstall=false
```

**NOTE** You cannot uninstall Application Server LDAP Provider, even if it appears on the Uninstaller dialog.

The table below shows the IDs of the components you can retain during an uninstall:

Component	Component ID
Content Engine Server	beanFeatureServer
Application Server LDAP Provider	beanFeature LDAP
.NET Clients (Windows only)	beanFeatureDotNetClient
Enterprise Manager (Windows only)	beanFeatureEM
COM Compatibility Client API (Windows only)	beanFeatureCCL
Java Clients	beanFeatureJavaClient
Object Store Upgrade Tool (Windows only)	beanFeatureClient

If you want to retain multiple Content Engine components during an uninstall, specify each as a command-line option. Use white space to separate successive options, as follows:

```
-P <compID-1>.activeForUninstall=false -P <compID-2>.activeForUninstall=false
```

Alternatively, you can specify the component(s) to be retained by listing them in an options file, one component per line. For example, an options file with the following two lines specifies retention of Content Engine Server and Object Store 3.5. to 4.0 Upgrade Tool:

```
-P beanFeatureServer.activeForUninstall=false
-P beanFeatureClient.activeForUninstall=false
```

To uninstall Content Engine, except for the two components specified in the options files, run the uninstaller command, as in the following Windows example where the options file is **options.txt**:

```
uninstaller.exe -options options.txt
```

The components in the table above have parent-child relationships, which are indicated in the Select Components screen of Content Engine Setup (see [“Select Components” on page 164](#) in

["Install and Deploy Content Engine" on page 160](#)). These relationships are shown again in the following table:

Parent Component	Child Component
Content Engine Server	Application Server LDAP Provider
.NET Clients	Administration Tools
.NET Clients	COM Compatibility Client API
Java Clients	Object Store 3.5 to 4.0 Upgrade Tool

If you want to remove a child component and retain the parent component, you must specify the parent and the child component IDs, *in that order*, in the command line:

```
-P <parentID>.activeforUninstall=false -P <childID>.activeForUninstall=true
```

or in the options file:

```
-P <parentID>.activeforUninstall=false  
-P <childID>.activeForUninstall=true
```

For example, the following options file specifies retention of the .NET Clients component and removal of the Object Store 3.5 to 4.0 Upgrade Tool:

```
-P beanFeatureDotNetClient.activeForUninstall=false  
-P beanbeanFeatureCCL.activeForUninstall=true
```

### **To remove data associated with Content Engine**

---

After uninstalling Content Engine, you can remove its associated data, as follows:

1. Use the application server console or command lines to undeploy Content Engine.
2. Use your database tools to drop any databases/tablespaces for object stores and GCD.
3. Use your LDAP tools to delete users and groups you created in [Task 2 "Specify IBM FileNet P8 Accounts" on page 60](#).
4. Use your operating system commands to delete any directories, users and groups used for installing and administering Content Engine; and delete file-storage-area directories containing content (for example, documents) and index-area directories (K2 collections).

## Remove Process Engine (Windows)

This task includes Process Engine removal instructions for Windows platforms.

**CAUTION** You must remove the software in the order listed below. If you remove Image Services (Step 8) before you remove Process Engine (Step 3) the Process Engine software will be left in a state that will not allow removal with this procedure.

### To remove the Process Engine software

1. Stop all of the following components that are running. For procedures and further details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Shut-down and Startup](#).

Component	Location
Process Simulator	Process Simulator
Process Analyzer	Process Analyzer
Component Manager	Application Engine
Application Engine	Application Engine
Content Engine	Content Engine
Process Service	Process Engine
Process Task Manager	Process Engine or Application Engine

2. Navigate to Control Panel > Add/Remove Programs.
3. Click **Remove** for the Process Engine application.
4. Click **Next** at the Welcome screen for Process Engine uninstallation.
5. Click **Uninstall** to confirm you want to remove Process Engine for Windows 4.0 installation.
6. Click **Next** to stop FileNet BPM software components.
7. Click **Finish** after you've read the summary information.
8. Click **Remove** for the FileNet Image Services 4.0.50 software.
9. Click **Yes** to confirm you want to remove the Image Services software.
10. Press **Enter** to continue with the uninstall.
11. Close the Add/Remove snap-in.
12. Close the Control Panel.

**CAUTION** (MS SQL Server only) If you plan to reinstall Process Engine and will configure Process Engine to use a different MS SQL database, you must remove the database that was configured

for the Process Engine installation. In addition, you must remove the following FileNet user IDs from the SQL Server Security folder before you reinstall Process Engine software:

- f\_sw
- f\_maint



## Remove Process Engine (UNIX)

This task includes Process Engine removal instructions for UNIX platforms.

### To remove the Process Engine software

1. Stop all of the following components that are running. For procedures and further details, see the IBM FileNet P8 help topic [FileNet P8 Administration > Enterprise-wide Administration > Shut-down and Startup](#).

Component	Location
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. On AIX, execute the following:

```
slibclean
```

4. Log on to the Process Engine server as the root user.
5. At a command prompt, run the following command:

```
/fnsw/_pws_uninst2/pws_uninstall.bin
```

6. Follow the prompts on the screen to remove the Process Engine software.
7. At a command prompt, run the following script:

```
/fnsw/etc/uninstall
```

The uninstall script shuts down the Process Engine software, removes the software from the **/fnsw** directory, and removes the directory structure under **/fnsw**. The uninstall script removes the Process Engine entries in the **/etc/services**, **/etc/system**, **/etc/inittab** and **/etc/devlink.tab** files. On an AIX-based Process Engine, the script also removes smit entries.

## Remove Application Engine (WebSphere)

This task includes Application Engine removal instructions for WebSphere on UNIX and Windows platforms.

### To remove the Application Engine software

---

1. Log on to the application server.
  - UNIX - Log on as a user with write access to the **/bin** directory and the directory where Application Engine is installed.
  - Windows - Log on as a user with Administrative rights.
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. The default location is:
  - UNIX - **/opt/FileNet/AE/\_uninst**
  - Windows - **C:\Program Files\FileNet\AE\\_uninst**
5. Run the uninstall program:
  - UNIX - **filenet\_app\_engine\_setup\_uninstall.bin**
  - Windows - **filenet\_app\_engine\_setup\_uninstall.exe**

**NOTE** For Windows, you can also use **Add/Remove Programs** from the Control Panel to remove the IBM FileNet Application Engine.

On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

6. Delete the Workplace folder:  
**<WAS\_HOME>/temp/<node\_name>/<application\_server\_name>/Workplace**
7. Delete the **<AE\_install\_path>** directory.
8. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the **vpd.properties** file. If it exists, delete it.

**WARNING** In the following step, do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

9. (UNIX) Remove the P8TASKMAN\_HOME system environment variable.

If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN\_HOME system environment variable to complete the uninstallation.

## Remove Application Engine (WebLogic)

This task includes Application Engine removal instructions for WebLogic on UNIX and Windows platforms.

### To remove the Application Engine software

---

1. Log on to the application server.
  - UNIX - Log on as a user with write access to the **/bin** directory and the directory where Application Engine is installed.
  - Windows - Log on as a user with Administrative rights.
2. Undeploy the Workplace application.
  - a. Stop the Workplace Web Application Module.
  - b. Delete the Workplace Web Application Module.
3. Navigate to the **/\_uninst** folder under the Application Engine installation location. The default location is:
  - UNIX - **/opt/FileNet/AE/\_uninst**
  - Windows - **C:\Program Files\FileNet\AE\\_uninst**
4. Run the uninstall program:
  - UNIX - **filenet\_app\_engine\_setup\_uninstall.bin**
  - Windows - **filenet\_app\_engine\_setup\_uninstall.exe**

**NOTE** For Windows, you can also use **Add/Remove Programs** from the Control Panel to remove Application Engine.

On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

5. Delete the **<AE\_install\_path>** directory.
6. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the **vpd.properties** file. If it exists, delete it.
7. (UNIX) Remove the P8TASKMAN\_HOME system environment variable.

If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN\_HOME system environment variable to complete the uninstallation.

**WARNING** Do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

## Remove Application Engine (JBoss)

This task includes Application Engine removal instructions for JBoss on Linux and Windows platforms.

### To remove the Application Engine software

---

1. Log on to the application server:
  - UNIX - logon as a user with write access to the **/bin** directory.
  - Windows - Log on as a member of the local Administrators group or a user with equivalent permissions.
2. Shut down JBoss.
3. Navigate to the **/\_uninst** folder under the Application Engine installation location. The default location is:
  - UNIX - **/opt/FileNet/AE/\_uninst**
  - Windows - **C:\Program Files\FileNet\AE\\_uninst**
4. Run the uninstall program:
  - UNIX - **filenet\_app\_engine\_setup\_uninstall.bin**
  - Windows - **filenet\_app\_engine\_setup\_uninstall.exe**

**NOTE** For Windows, you can also use **Add/Remove Programs** from the Control Panel to remove Application Engine.

On Windows, when you click **Next** on the Application Engine Uninstaller Welcome screen, the second screen asks you to wait while Windows Task Manager service shuts down before continuing the uninstall. This can take a few moments. Wait for the shutdown to complete, then complete the uninstall wizard from the next screen.

5. (If Application Engine is the only IBM FileNet P8 application installed on the server) Search for the **vpd.properties** file. If it exists, delete it.
6. (UNIX) Remove the P8TASKMAN\_HOME system environment variable.

If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN\_HOME system environment variable to complete the uninstallation.

**WARNING** Do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

7. Remove temporary files and directories.
  - a. Remove the **Workplace.war** folder from **<JBoss\_HOME>\server\default\deploy\**.
  - b. Remove the Temp working directory for Workplace from **<JBoss\_HOME>\server\default\work\MainEngine\localhost\**
8. Delete the **<AE\_install\_path>** directory.

9. (LINUX) Remove the P8TASKMAN\_HOME system environment variable.

If Application Engine is the *only* IBM FileNet P8 application running on the server you must remove the P8TASKMAN\_HOME system environment variable to complete the uninstallation.

**WARNING** Do *not* remove the system environment variable if any other IBM FileNet P8 application is installed on the server.

## Remove the Application Engine ISRA Servlet

This task includes Application Engine ISRA Servlet removal instructions for Windows and UNIX environments.

**NOTE** Since the installed names for the ISRA Servlet are configurable on the supported application servers, the information below may not be the same as your environment. Make the appropriate name changes as required for your environment.

### To remove the Application Engine Servlet software

---

1. Log on to the application server.

**UNIX** - Log on as a user with write access to the /bin directory and the directory where ISRA Servlet is installed.

**Windows** - Log on as a user with Administrative rights.

2. Undeploy the ApplicationEngineISRAServlet application. This step is similar to that required to undeploy the Workplace application.

**WebSphere:**

- a. Stop the ApplicationEngineISRAServlet process in the Admin console
- b. Uninstall the ApplicationEngineISRAServlet application from Enterprise Applications

**WebLogic:**

- a. Stop the ApplicationEngineISRAServlet Web Application Module.
- b. Undeploy or delete the ApplicationEngineISRAServlet Web Application Module.

3. Navigate to the **/\_uninstISRAServlet** directory under the ISRA Servlet installation location. The default location is:

**UNIX** /opt/

**Windows** C:\Program Files\

4. Run the uninstall program:

**UNIX** filenet\_ae\_israservlet\_setup\_uninstall.bin

**Windows** filenet\_ae\_israservlet\_setup\_uninstall.exe

**NOTE** For Windows, you can also use Add/Remove Programs from the Control Panel to remove the FileNet Application Engine ISRA Servlet.

5. Navigate to the **\FileNet** directory. If there is no other FileNet software installed under this directory, delete the **\FileNet** directory. If there is some other FileNet software installed under this directory, delete only the **\ApplicationEngineISRAServlet** subdirectory.
6. WebSphere only. Delete the following temporary working folders for the Application Engine ISRA Servlet:

### WebSphere 5.x

```
<WAS_Home>\WebSphere\AppServer\installedApps\<servername>\ApplicationEngineISRAServlet.ear\
```

### WebSphere 6.x

```
<WAS_Home>\WebSphere\AppServer\profiles\default\installedApps\<servername>\ApplicationEngineISRAServlet.ear\
```



## *Appendixes*

This appendixes section contains the following major topics:

- [“Encrypt Passwords for Silent Installations and Upgrades” on page 642](#)
- [“IBM FileNet P8 Port Numbers” on page 643](#)
- [“Process Engine SQL Scripts” on page 647](#)
- [“IBM FileNet P8 Database Character Sets” on page 653](#)
- [“New Content Engine Classes and Properties” on page 656](#)

# Encrypt Passwords for Silent Installations and Upgrades

You can use resource-input files to install and upgrade IBM FileNet P8 Platform components silently, as well using interactive Setup programs, as noted throughout this guide.

Several passwords are required to silently install and upgrade IBM FileNet P8 components. To accommodate your security requirements, you can encrypt these passwords, as follows, before you enter them into the silent installation and upgrade resource files.

## To encrypt a password

---

1. The encryption tool is located on the installation media for Process Engine or Content Engine software, in the Tools directory. Copy the following files from the Tools directory to a local drive:
  - `fnencryptutils` - an executable .jar file
  - `RunEncryptApp` - a batch file
2. Run one of the following executable files to invoke the application:
  - (Windows) **RunEncryptApp.bat**
  - (UNIX) **RunEncryptApp.sh**

**NOTE** Before running the file, be sure Java is installed and its location is in your PATH environment variable.

3. Enter the appropriate values for the user name and password.
4. Reenter the password to confirm it.
5. Click **Generate**.
6. An encrypted password will be generated and displayed in the encrypted password field.
7. Copy and paste this password into the appropriate silent installation or upgrade resource file.

## IBM FileNet P8 Port Numbers

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

Content Engine Ports	
LDAP / SSL	389 / 636
WebSphere (EJB / WSI)	2809 / 9080
WebLogic	7001
JBoss (EJB / WSI / IIOP)	1099 / 8080 / 3528
MSSQL Server / Oracle / DB2	1433 / 1521 / 50000
Process Engine (RMI)	32771

Process Engine Ports	
SMTP (E-mail Notification)	25
MSSQL Server / Oracle / DB2	1433 / 1521 / 50000
Oracle Services for Microsoft Transaction Server (Windows)	2030
Process Engine (RMI)	32771
TMS	32768 / TCP
COR	32769 / TCP
NCH	32770 / UDP
fn_snmpd	161 / UDP
fn_trapd	35225 / UDP
SNMPD SMUX (AIX only)	199 / TCP
Rules Listener	32774 (TCP/IP)
Component Manager (Event Port)	32773
Process Engine Communication Port (IOR port)	32776
Process Engine Broker Port	32777
BPM Web Services Reliable messaging client port	32767 / TCP 32767 / UDP

**NOTE** If the port number assigned to Component Manager conflicts with the port number required by another application or service that runs on the Process Engine or the Application Engine

server, Process Task Manager will not start and the necessary **vwtaskman.xml** will not be automatically created. If this happens, make a copy of the sample **vwtaskman.xml.sample** file located on the Process Engine or Application Engine. On Process Engines, the file is located in the **/fnsw/bin** directory. On Application Engines, the file is located in **Drive:\Program Files\FileNet\AE\Router** on Windows and in **/opt/FileNet/AE/Router** on UNIX. Open **vwtaskman.xml.sample** with a text editor, change the port element value to an available port number, and save the file to **vwtaskman.xml** in the same directory.

<b>Application Engine Ports</b>	
WebSphere / WebSphere SSL	9080 / 443
WebLogic / WebLogic SSL	7001 / 7002
JBoss / JBoss SSL	8080 / 8443
Process Engine (RMI)	32771
Component Manager (Event Port)	32773
Web Services Reliable Messaging Client Port	32777

<b>Autonomy K2 Search Engine Ports</b>	
K2 Administrative Server	9950
K2 Dashboard/Business Console	9990
K2 Index Server	9960
K2 Broker Server	9900
K2 Server	9920
K2 Ticket Server	9910

<b>Rendition Engine Ports</b>	
Liquent / Print Server on Local Port (Liquent XML Printer)	<Drive>:\Program Files \Liquent\filenet\liquent.xml
Liquent / Print Server on Local Port (Liquent PDF Printer)	<Drive>:\Program Files \Liquent\filenet\ESPS.PS
Liquent notify port	2868

Rendition Engine Ports	
Liquent event port	2869
Liquent admin port	2870
Liquent file transfer port	2871
Microsoft SQL Server	1433
Oracle	1521

Process Analyzer Ports	
Registry Port	32771
Database Port	1433

Process Simulator Ports	
Process Analyzer	32771
Registry	32771
Database Port	1433
Return	0 (anonymous)

System Manager Ports	
Listener (first)	32775
Listener (subsequent)	OS defined
Database Port	1433
Return	0 (anonymous)

**NOTE** If necessary, you can limit the subsequent port numbers to a specific range using the settings in the **PchConfig.properties** file.

Content Federation Services for Image Services	
tms	32768 / TCP
cor	32769 / TCP
nch	32770 / UDP
fn_snmpd	161 / UDP
snmp trap	162 / UDP
fn_trapd	35225 / UDP
Native default SNMP port (HP only)	8000 / UDP
IBM FileNet specific SNMP port (HP and Solaris only)	8001 / UDP
Migration notify	anonymous

**NOTE** On UNIX platforms, Image Services port assignments are made in the `/etc/services` file. For more information, see the *IBM FileNet P8 Content Federation Services for Image Services Guidelines*. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs” on page 21](#).

## Process Engine SQL Scripts

Process Engine Setup runs several SQL scripts for Oracle and SQL Server databases. However, the scripts can be run the following ways:

- Run the scripts manually before you start Process Engine Setup
- Let the setup software prompt you for the Oracle SYS or SQL server administrator (sa) password and then run the scripts automatically
- Run the scripts silently using operating system authentication

Use operating system authentication only in a trusted environment or when Process Engine is configured with a local database.

If the scripts are executed from the Process Engine Setup program, the ability to connect to the database is validated during setup. If the connection fails, the user can correct the errors and proceed with the installation.

If the scripts are executed manually, the database connections are not validated until the end of the installation.

If the scripts are executed manually for a SQL Server database, the SQL Server client software does not have to be installed on the Process Engine server.

If the scripts are copied to the database server and executed manually by the DBA, there is no need to provide access to the database / sysadmin password for the person who executes Process Engine Setup.

### SQL Scripts for SQL Server

Script Name	Action
<b>CreatePEinstallSP_1.sql</b>	Creates PE_createDbUsers stored procedure in the Process Engine database.
<b>CreatePEinstallSP_2.sql</b>	Creates a stored procedure that is called during Process Engine installs.
<b>CreatePEinstallSP_3.sql</b>	Creates fn_error stored procedure.

## SQL Scripts for Oracle

Script Name	Action
<b>pe_install_scripts.sql</b>	A wrapper script that executes the following scripts: <b>pe_filenet_site.sql</b> , <b>pe_create_stored_procedures.sql</b> , and <b>pe_grant_sp_permissions.sql</b> .  The <b>pe_install_scripts.sql</b> script generates an Oracle output spool file. After the script completes, Process Engine Setup scans the output file for errors.
<b>pe_upgrade_scripts.sql</b>	A wrapper script that executes the following scripts: <b>pe_create_stored_procedures.sql</b> and <b>pe_grant_sp_permissions.sql</b> .  The <b>pe_upgrade_scripts.sql</b> script generates an Oracle output spool file. After the script completes, Process Engine Setup scans the output file for errors.
<b>pe_filenet_site.sql</b>	Creates the database users for IBM FileNet P8. See <a href="#">“Changes made by pe_filenet_site.sql and pe_oracle_users_defaults.sql”</a> on page 649 for details.
<b>pe_create_stored_procedures.sql</b>	Creates several stored procedures. See <a href="#">“Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql”</a> on page 650 for details.
<b>pe_grant_sp_permissions.sql</b>	Executes grants and creates synonyms for stored procedures. See <a href="#">“Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql”</a> on page 650 for details.
<b>pe_oracle_users_defaults.sql</b>	Sets user tablespace defaults and privileges for IBM FileNet P8 users. See <a href="#">“Changes made by pe_filenet_site.sql and pe_oracle_users_defaults.sql”</a> on page 649 for details.

All the scripts are located in the root of the Process Engine installation directory. Process Engine Setup cannot complete until all these scripts have run successfully. If an error message indicates that any of the scripts did not run, you must resolve the errors before you proceed. See the *IBM FileNet P8 Platform Troubleshooting Guide* for more information. To download this guide from the IBM support page, see [“Access IBM FileNet Documentation, Compatibility Matrices, and Fix Packs”](#) on page 21.



### Changes made by `pe_filenet_site.sql` and `pe_ora_users_defaults.sql`

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

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

\* Changed automatically during Process Engine software installation from *filenet* to an internal secret password.

\*\* Manually changed after installation with the `set_f_maint_pw` tool, as described in each “Install Process Engine <Platform>” topic of this guide.

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

### ***Changes made by pe\_create\_stored\_procedures.sql and pe\_grant\_sp\_permissions.sql***

The information in columns 1, 2 and 3 is created by pe\_create\_stored\_procedures.sql. The information in columns 4 and 5 is created by **pe\_grant\_sp\_permissions.sql**.

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

The scripts run with the following options:

*<run-time user>*

This option value is one of the following:

f\_sw - if you use the default operating system and database users when you run Process Engine Setup.

alias for f\_sw - if you define aliases for the operating system and database users when you run Process Engine Setup.

*<maintenance user>*

This option value is one of the following:

f\_maint - if you use the default operating system and database users when you run Process Engine Setup.

alias for f\_maint - if you define aliases for the operating system and database users when you run Process Engine Setup.

<data tablespace> (Oracle only)

The data tablespace name that will be entered in Process Engine Setup.

<index tablespace> (Oracle only)

The index tablespace name that will be entered in Process Engine Setup. This is an optional tablespace. If it doesn't exist, enter the data tablespace value.

<temp tablespace> (Oracle only)

The temp tablespace name that will be entered in Process Engine Setup.

<PE database name> (SQL Server only)

The database name that will be entered in Process Engine Setup.

<DSN> (SQL Server only)

The ODBC data source name that will be entered in Process Engine Setup.

### To run SQL scripts manually (SQL Server)

---

1. Copy **CreatePEinstallSP\_1.sql**, **CreatePEinstallSP\_2.sql**, and **CreatePEinstallSP\_3.sql** from the Process Engine CD to the local disk on the SQL Server database server.
2. Run **CreatePEinstallSP\_3.sql** first, then **CreatePEinstallSP\_2.sql**, then **CreatePEinstallSP\_1.sql**.

```
osql -E -D <DSN> -d <PE database name> -i CreatePEinstallSP_3.sql -n -o output3.log
```

```
osql -E -D <DSN> -d master -i CreatePEinstallSP_2.sql -n -o output2.log
```

```
osql -E -D <DSN> -d <PE database name> -i CreatePEinstallSP_1.sql -n -o output1.log
```

Instead of -E the following can be used:

```
-U sa -P <sa password>
```

Successful execution will record nothing in the **output3.log** and **output2.log** files. On a SQL Server 2000 configuration, the **output1.log** file may contain:

```
Cannot add rows to sysdepends for the current stored procedure because it depends on the missing object 'sys.sp_validname'. The stored procedure will still be created.
```

This error can be ignored.

3. Create a new text file containing the following:

```
PE_createDbUsers '<run-time user>', '<maintenance user>', '<PE database name>'
```

For example:

```
PE_createDbUsers 'f_sw', 'f_maint', 'simDB'
```

Then execute this procedure

```
osql -E -D <DSN> -d simDB -i input.txt
```

Instead of -E the following can be used:

```
-U sa -P <sa password>
```

### To run SQL scripts manually (Oracle)

---

1. From the Process Engine CD, copy the scripts to the destination computer.
2. Start SQL Plus. For example:

```
sqlplus "sys/<password> as sysdba"
```

3. At the SQL prompt, enter:

```
@pe_install_scripts.sql <run-time user> <maintenance user> <data tablespace> <index  
tablespace> <temp tablespace>
```

For example:

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

# IBM FileNet P8 Database Character Sets

## IBM FileNet P8 Database components character sets

The table below lists the character sets supported for the IBM FileNet P8 database components.

Database type	Character sets supported in the regular character set (via varchar and char data types)	Character sets supported in the national character set (via nvarchar, nchar, and graphics data types)	Character set used by Content Engine	Character set used by Process Engine
<b>Oracle 9i</b>	All standard Oracle character sets  AL32UTF8 Unicode character sets	Not used by FileNet P8	Regular character set  AL32UTF8 recommended	Regular character set
<b>Oracle 10g</b>	All standard Oracle character sets  AL32UTF8 Unicode character sets	Not used by FileNet P8	Regular character set  AL32UTF8 recommended	Regular character set
<b>MS SQL Server</b>	Only non-Unicode code pages are supported here	UCS-2 (Generically called Unicode on SQL Server)	Unicode/national character set only	Regular character set only
<b>DB2</b>	UTF-8 and all non-Unicode character sets	UCS-2	National character set only	Regular character set only

### Additional Oracle character set information

The Oracle database character set you choose depends on whether the database is shared or dedicated.

- The recommended character set for a database dedicated to Content Engine or shared with Process Engine is AL32UTF8 (Unicode 3.1 UTF-8 universal character set).

- If the database is dedicated to Process Engine, choose from among the supported character sets listed in the following table.

Description	Character set
Unicode 3.1 UTF-8 universal character set	AL32UTF8
Western European	(UNIX) WE8ISO8859P1 or WE8ISO8859P15 (Windows) WE8MSWIN1252
Eastern European	EE8ISO8859P2
South European	SE8ISO8859P3
Northern & Northeastern Europe	NEE8ISO8859P4
Latin/Cyrillic	CL8ISO8859P5
Latin/Arabic	AR8ISO8859P6
Latin/Greek	FL8ISO8859P7
Latin/Hebrew	IWISO8859P8
Western European & Turkish	WEISO8859P9
North European	NE8ISO8859P10
ASCII 7-bit American	US7ASCII

Be aware of the special NLS\_LANG settings for character sets and locale on Oracle Client machines. Consider the following:

- You must ensure that the NLS\_LANG Oracle environment variable on an Oracle Client machine matches the character set/locale of the client operating system.
- Under Windows, by default the Oracle Client installer sets the NLS\_LANG value in the Windows registry to match the locale of the Oracle Client machine's operating system. For Process Engine you do not need to override the registry value with the user environment variable. The default NLS\_LANG value is adequate for either a Unicode (for example, AL32UTF8) or non-Unicode (for example, WE8MSWIN1252) database character set.
- Under UNIX, the Oracle Client installer does not automatically set the NLS\_LANG (as it does under Windows). For this reason, after you have installed Process Engine, you must manually set the locale and character set value on each UNIX Oracle Client machine from which users will access Process Engine.
- If the database uses a non-Unicode character set (e.g., WE8MSWIN1252), you can make the NLS\_LANG value on the Oracle Client machine match the character set of the database by choosing the appropriate database character set during database creation. For example, when the client operating system locale is American ANSI 1252, you can designate the WE8MSWIN1252 character set when you create the database. The Windows Oracle Client installer will set NLS\_LANG to be AMERICAN\_AMERICA.WE8MSWIN1252 because the client

operating system locale is American ANSI 1252. Under UNIX, you would manually set this value.

- You set the NLS\_LANG value manually on Oracle Client machines as follows:
  - (UNIX) Add NLS\_LANG to the shell environment login files for each user who will be logging on to the machine to run IBM FileNet P8 software.
  - (Windows) Set or modify the value of the NLS\_LANG key using System Properties in the Control Panel for each user who will be logging on to the machine to run IBM FileNet P8 software.

### ***Additional DB2 character set information***

Content Engine requires a unicode (UTF-8) database. Process Engine will support the UTF-8 code set, code page 1208. Process Engine will also support all other single-byte character sets, for example:

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

## New Content Engine Classes and Properties

This appendix lists the new classes and properties that have been added to Content Engine in the 4.0.0 release. If you have previously defined any classes or properties in a 3.5.x object store with names that conflict with any of the class or property names listed here, it will not be possible to upgrade the object store to 4.0.0. The upgrade tool will check for naming conflicts and require you to change the names. To complete the upgrade, you must change the name of your user-defined class or property.

### Content Engine Classes

The following table displays the new class names in Content Engine 4.0. To ensure successful upgrade, resolve any conflicts between your 3.x object store class names and the new names.

Class Symbolic Name	Class ID
AddOnInstallationRecord	Add On Installation Record
AsyncProcessingConfiguration	Async Processing Configuration
CenteraSiteSettings	Centera Site Settings
CodeModule	Code Module
ComponentRelationship	Component Relationship
ContentCacheArea	Content Cache Area
ContentCacheConfiguration	Content Cache Configuration
ContentConfiguration	Content Configuration
DatabaseStorageArea	Database Storage Area
DirectoryConfiguration	Directory Configuration
DirectoryConfigurationAD	Directory Configuration AD
DirectoryConfigurationADAM	Directory Configuration ADAM
DirectoryConfigurationIBM	Directory Configuration IBM
DirectoryConfigurationNovell	Directory Configuration Novell
DirectoryConfigurationSunOne	Directory Configuration SunOne
Domain	Domain
EntireNetwork	EntireNetwork
FileStorageArea	File Storage Area
FixedStorageArea	Fixed Storage Area
Group	Group



Class Symbolic Name	Class ID
ImageServicesClassDescription	Image Services Class Description
ImageServicesImportAgentConfiguration	Image Services Import Agent Configuration
ImageServicesPropertyDescription	Image Services Property Description
ImageServicesSiteSettings	Image Services Site Settings
IndexArea	Index Area
IndexJob	Index Job
IndexJobClassItem	Index Job Class Item
IndexJobCollectionItem	Index Job Collection Item
IndexJobItem	Index Job Item
IndexJobSingleItem	Index Job Single Item
IsolatedRegion	Isolated Region
PEConnectionPoint	PE Connection Point
PublishCompleteEvent	Publish Complete Event
PublishRequest	Publish Request
PublishRequestEvent	Publish Request Event
PublishingConfiguration	Publishing Configuration
Realm	Realm
RenditionEngineConnection	Rendition Engine Connection
SecurityPrincipal	Security Principal
ServerCacheConfiguration	Server Cache Configuration
ServerInstance	Server Instance
Site	Site
SiteSettings	Site Settings
StorageArea	Storage Area
StoragePolicy	Storage Policy
SubsystemConfiguration	Subsystem Configuration
TakeFederatedOwnershipEvent	Take Federated Ownership Event
TraceLoggingConfiguration	Trace Logging Configuration

Class Symbolic Name	Class ID
UpgradeAddOn	Upgrade Add On
User	User
VerityCollection	Verity Collection
VerityDomainConfiguration	Verity Domain Configuration
VerityIndexArea	Verity Index Area
VerityServerConfiguration	Verity Server Configuration
VirtualServer	Virtual Server

### **Content Engine Properties**

The following table displays the new property names in Content Engine 4.0. To ensure successful upgrade, resolve any conflicts between your 3.x object store property names and the new names.

Property Symbolic Name	Property Name
AbandonedContentCleanupInterval	Abandoned Content Cleanup Interval
AbandonedDBContentCleanupInterval	Abandoned Database Content Cleanup Interval
ActiveDirectorySiteDNS	Active Directory Site DNS
AddOn	Add On
AddOnInstallationRecords	Add On Installation Records
AddOns	Add Ons
AddOnName	Add On Name
AllRealms	All Realms
AllowFederatedDeletes	Allow Federated Deletes
AllowsContentToBeCached	Allows Content To Be Cached
AllowsDelete	Allows Delete
APITraceFlags	API Trace Flags
AppenderNames	Appender Names
ApplyDefinition	Apply Definition
AsynchronousProcessingTraceFlags	Asynchronous Processing Trace Flags
AuditedDeletePrefix	Audited Delete Prefix

Property Symbolic Name	Property Name
BaseClassIDs	Base Class IDs
BatchDelay	Batch Delay
BlobReadAheadSize	Blob Read Ahead Size
BlobWriteCollisionsAvoidanceFlag	Blob Write Collisions Avoidance Flag
BrokerPort	Broker Port
FPPoolBufferSize	Buffer Size
CacheStatus	Cache Status
CanAcceptForwardedRequests	Can Accept Forwarded Requests
CanForwardRequests	Can Forward Requests
CBRLocale	CBR Locale
CBRTraceFlags	CBR Trace Flags
CFSDaemonTraceFlags	CFS Daemon Trace Flags
ChildComponent	Child Component
ChildDocuments	Child Documents
ChildRelationships	Child Relationships
ChildVersionSeries	Child Version Series
ClassDefinition	Class Definition
ClosureDate	Closure Date
FPPoolClusterNonAvailTime	Cluster Nonavailable Time
CodeModule	Code Module
CodeModuleCacheEntryTTL	Code Module Cache Entry TTL
CodeModuleCacheMaxFileSpace	Code Module Cache Maximum File Space
CodeModuleCacheMaxMemory	Code Module Cache Maximum Memory
CodeModuleTraceFlags	Code Module Trace Flags
CollectionName	CollectionName
ComponentCascadeDelete	Component Cascade Delete
ComponentPreventDelete	Component Prevent Delete
ComponentRelationshipType	Component Relationship Type

Property Symbolic Name	Property Name
ComponentSortOrder	Component Sort Order
CompoundDocumentState	Compound Document State
ConcurrentReaders	Concurrent Readers
ConnectionState	Connection State
ConnectionTimeout	Connection Timeout
ContentCacheArea	Content Cache Area
ContentCacheAreas	Content Cache Areas
ContentCacheTraceFlags	Content Cache Trace Flags
ContentElementCount	Content Element Count
ContentElementKBytes	Content Element Kbytes
ContentElementsCreated	Content Elements Created
ContentElementsDeleted	Content Elements Deleted
ContentQueueMaxWorkerThreads	Content Queue Max Worker Threads
ContentStorageTraceFlags	Content Storage Trace Flags
ContentTempDirectoryRoot	Content Temp Directory Root
CSMCache	CSM Cache
CurrentUser	Current User
DatabaseContentUploadBufferSize	Database Content Upload Buffer Size
DatabaseServerPort	Database Server Port
DatabaseTraceFlags	Database Trace Flags
DatabaseType	Database Type
FPPoolDefaultCollisionAvoidance	Default Collision Avoidance
DefaultISDocumentClass	Default IS Document Class
DefaultRetentionDays	Default Retention Days
DefaultRetentionPassThrough	Default Retention Pass Through
DefaultSite	Default Site
DeleteMethod	Delete Method
DeviceAddress	Device Address

Property Symbolic Name	Property Name
DeviceRootDirectory	Device Root Directory
DirectoryConfigurations	Directory Configurations
DirectoryServerHost	Directory Server Host
DirectoryServerPassword	Directory Server Password
DirectoryServerPort	Directory Server Port
DirectoryServerProviderClass	Directory Server Provider Class
DirectoryServerType	Directory Server Type
DirectoryServerUserName	Directory Server User Name
DirectoryStructure	Directory Structure
DispatcherEnabled	Dispatcher Enabled
DispatcherWaitInterval	Dispatcher Wait Interval
DistinguishedName	Distinguished Name
DNSName	DNS Name
DocumentsPerBatch	Documents Per Batch
Domain	Domain
DomainID	Domain ID
EJBForwardingEndpoint	EJB Forwarding Endpoint
EJBTraceFlags	EJB Trace Flags
EmailAddress	Email Address
EmbeddedDataThreshold	Embedded Data Threshold
FPPoolEnableMulticlusterFailover	Enable Multicluster Failover
EncryptionAlgorithm	Encryption Algorithm
EngineTraceFlags	Engine Trace Flags
ErrorCode	Error Code
ErrorDescription	Error Description
EventsTraceFlags	Events Trace Flags
ExpiredBatchSelectionSize	Expired Batch Selection Size
FixedContentDevice	Fixed Content Device

Property Symbolic Name	Property Name
FixedContentDevices	Fixed Content Devices
FixedContentProviderTraceFlags	Fixed Content Provider Trace Flags
FolderCacheMaxAgeDelta	Folder Cache Maximum Age Delta
FolderCacheMaxEntries	Folder Cache Maximum Entries
FolderCacheReapFrequency	Folder Cache Reap Frequency
FromVersions	From Versions
FullTextRowDefault	Full Text Row Default
FullTextRowMax	Full Text Row Max
GCDCacheTTL	GCD Cache TTL
GCDTraceFlags	GCD Trace Flags
GroupBaseDN	Group Base DN
GroupDisplayNameAttribute	Group Display Name Attribute
GroupMembershipSearchFilter	Group Membership Search Filter
GroupNameAttribute	Group Name Attribute
GroupSearchFilter	Group Search Filter
Groups	Groups
ImageServicesClassDescriptions	Image Services Class Descriptions
ImageServicesClassName	Image Services Class Name
ImageServicesDataType	Image Services Data Type
ImageServicesPropertyDescriptions	Image Services Property Descriptions
ImageServicesPropertyName	Image Services Property Name
ImplementationClass	Implementation Class
ImportAgentRetryCount	Import Agent Retry Count
InboundFileNameCacheMaxEntries	Inbound File Name Cache Maximum Entries
IndexArea	Index Area
IndexAreas	Index Areas
IndexItems	Index Items
IndexationID	Indexation ID

Property Symbolic Name	Property Name
InitiatingUser	Initiating User
InlineContentRetrievalLimit	Inline Content Retrieval Limit
InputDocument	Input Document
InstallationDate	Installation Date
InstallationReport	Installation Report
Installer	Installer
IsCBREnabled	Is CBR Enabled
ISDomainName	IS Domain Name
IsFederatedOnly	Is Federated Only
ISOrganization	IS Organization
ISPassword	IS Password
IsSSLEnabled	Is SSL Enabled
ISTempDir	IS Temporary Directory
ISUserName	IS User Name
IsolatedRegion	Isolated Region
IsolatedRegions	Isolated Regions
JNDIDataSource	JNDI DataSource
JNDIXADataSource	JNDI XA DataSource
JobAbortRequested	Job AbortRequested
JobStatus	Job Status
LabelBindValue	Label Bind Value
LeaseDuration	Lease Duration
LocalDomain	Local Domain
MarkingSetCacheEntryTTL	Marking Set Cache Entry TTL
MarkingSetCacheMaxEntries	Marking Set Cache Maximum Entries
MarkingSets	Marking Sets
MarkingUseGranted	Marking Use Granted
MaxBatchSize	Max Batch Size

Property Symbolic Name	Property Name
MaxCollections	Max Collections
MaxInMemoryElementState	Max In Memory Element State
MaxInMemoryQueueItems	Max In Memory Queue Items
MaxObjectsPerCollection	Max Objects Per Collection
MaxReaderSemaphoreWaitTime	Max Reader Semaphore Wait Time
FPPoolMaxConnections	Maximum Connections
MaximumContentElements	Maximum Content Elements
MaxResolutionBatchSize	Maximum Resolution Batch Size
MaximumRetentionDays	Maximum Retention Days
MaximumSizeKBytes	Maximum Size Kbytes
MaximumTimeToLive	Maximum Time To Live
MemberOfGroups	Member Of Groups
MetadataTraceFlags	Metadata Trace Flags
MinimumRetentionDays	Minimum Retention Days
FPPoolMulticlusterDeleteStrategy	Multicluster Delete Strategy
FPPoolMulticlusterExistsStrategy	Multicluster Exists Strategy
FPPoolMulticlusterQueryStrategy	Multicluster Query Strategy
FPPoolMulticlusterReadStrategy	Multicluster Read Strategy
FPPoolMulticlusterWriteStrategy	Multicluster Write Strategy
MyRealm	My Realm
NeverDeleteClipsOrContent	Never Delete Clips Or Content
NoWorkDelay	No Work Delay
NumImportAgents	Num Import Agents
ObjectSecurityCacheEntryTTL	Object Security Cache Entry TTL
ObjectSecurityCacheMaxEntries	Object Security Cache Maximum Entries
ObjectStores	Object Stores
FPOpenStrategy	Open Strategy
OptimizationInterval	Optimization Interval



Property Symbolic Name	Property Name
OriginalObject	Original Object
OriginalOrdinal	Original Ordinal
OutputFolder	Output Folder
OutputLocation	Output Location
ParentComponent	Parent Component
ParentDocuments	Parent Documents
ParentRelationships	Parent Relationships
PartialResolutionChunkSize	PartialResolutionChunkSize
PEConnectionPoints	PE Connection Points
PersistenceType	Persistence Type
PoolAddress	Pool Address
FPPoolPrefetchSize	Prefetch Size
PreloadOnCreate	Preload On Create
Prerequisites	Prerequisites
PrivilegedSettability	Privileged Settability
PruneAmount	Prune Amount
PruneThresholdContentElements	Prune Threshold Content Elements
PruneThresholdSizeKBytes	Prune Threshold Size KBytes
PublicKey	Public Key
PublicationDocument	Publication Document
PublishRequestType	Publish Request Type
PublishStyleTemplate	Publish Style Template
PublishTemplate	Publish Template
PublishTraceFlags	Publish Trace Flags
PublishingStatus	Publishing Status
QueueItemDatabaseTimeout	Queue Item Database Timeout
QueueItemMaxDispatchers	Queue Item Max Dispatchers
QueueItemRetryCount	Queue Item Retry Count

Property Symbolic Name	Property Name
ReferencingActions	Referencing Actions
RegionPassword	Region Password
RenameFileRetryAttempts	Rename File Retry Attempts
RenditionEngineConnection	Rendition Engine Connection
RenditionEngineConnections	Rendition Engine Connections
RenditionEnginePassword	Rendition Engine Password
ResourceStatus	Resource Status
ResourceString	Resource String
RetentionPeriod	Retention Period
RetrievalRetryAttempts	Retrieval Retry Attempts
FPPoolRetryCount	Retry Count
FPPoolRetrySleep	Retry Sleep
ReturnNameAsDN	Return Name As DN
RollFwdBatchRetryAttempts	Roll Forward Batch Retry Attempts
RootDirectoryPath	Root Directory Path
SearchCrossForestGroupMembership	Search Cross Forest Group Membership
SearchServersToAttach	Search Servers To Attach
SearchTraceFlags	Search Trace Flags
SecurityDescCacheMaxEntries	Security Descriptor Cache Maximum Entries
SecurityTraceFlags	Security Trace Flags
ServerInstances	Server Instances
ShortName	Short Name
SingleItem	Single Item
Site	Site
SiteSettings	Site Settings
Sites	Sites
SnapLockPassword	SnapLock Password
SnapLockUserName	SnapLock User Name

Property Symbolic Name	Property Name
SSITraceFlags	SSI Trace Flags
StagingAreaPath	Staging Area Path
Status	Status
StatusDescription	Status Description
StorageArea	Storage Area
StorageAreas	Storage Areas
SubjectCacheEntryTTL	Subject Cache Entry TTL
SubjectCacheMaxEntries	Subject Cache Maximum Entries
SubmittedCount	Submitted Count
SubsystemConfigurations	Subsystem Configurations
TemplateType	Template Type
TempDBContentLifetime	Temporary Database Content Lifetime
TempDirectoryPath	Temp Directory Path
TempFileLifetime	Temporary File Lifetime
ThreadCount	Thread Count
TimeAllSubmitted	Time All Submitted
TimeLastProcessed	Time Last Processed
FPPoolTimeout	Timeout
TimeOutSeconds	Timeout Seconds
ToVersions	To Versions
TraceLoggingEnabled	Trace Logging Enabled
URIValue	URI Value
UserBaseDN	User Base DN
UserDisplayNameAttribute	User Display Name Attribute
UserDomain	User Domain
UserGroup	User Group
UserName	User Name
UserNameAttribute	User Name Attribute

Property Symbolic Name	Property Name
UserPassword	User Password
UserSearchFilter	User Search Filter
UserTokenCacheEntryTTL	User Token Cache Entry TTL
UserTokenCacheMaxEntries	User Token Cache Maximum Entries
Users	Users
VerityAdminServerHostname	Verity Admin Server Hostname
VerityAdminServerPort	Verity Admin Server Port
VerityBrokerName	Verity Broker Name
VerityBrokers	Verity Brokers
VerityCollections	Verity Collections
VerityDomainConfiguration	Verity Domain Configuration
VerityIndexServers	Verity Index Servers
VerityMasterAdminServerHostname	Verity Master Admin Server Hostname
VerityMasterAdminServerPort	Verity Master Admin Server Port
VeritySearchServers	Verity Search Servers
VersionBindType	Version Bind Type
VirtualServer	Virtual Server
VirtualServers	Virtual Servers
WSITraceFlags	WSI Trace Flags

# Index

## A

- access roles, Application Engine Administrators [406](#)
- accounts, create
  - Application Engine [79](#)
  - Content Engine [62](#)
  - Content Search Engine [60](#)
  - Process Engine [71, 74](#)
- Application Engine
  - administration [37](#)
  - administrator access roles [406](#)
  - before you upgrade [551](#)
  - deploy (JBoss) [394](#)
  - deploy (WebLogic) [373, 391](#)
  - deploy (WebSphere) [356, 385](#)
  - install [344](#)
  - multiple instances [443](#)
  - remove (JBoss) [637](#)
  - remove (WebLogic) [636](#)
  - remove (WebSphere) [634](#)
  - remove Workplace [552](#)
  - upgrade [551](#)
- Application Integration
  - install [429](#)
  - Toolkit custom applications [609](#)
  - uninstall [432](#)
  - upgrade [560](#)
  - verify installation [431](#)
  - verify upgrade [561](#)
- applications, custom, upgrade [563](#)
- authentication
  - domains [29](#)
  - IBM Tivoli Directory Server [59](#)
  - Novell eDirectory [58](#)
  - realms [29](#)
  - Sun Java System Directory Server [56](#)
  - Windows Active Directory [54](#)
  - Windows ADAM [55](#)
- automatic ANSI to OEM conversion [98](#)
- Autonomy K2, see Content Search Engine

## B

- banner image, change [405](#)
- bootstrap preferences
  - bootstrap.properties file location [404](#)
  - configure [403](#)
  - SSL configuration [404](#)
- bootstrap.properties file
  - location [404](#)

- share between Application Engines [417](#)

## C

- Centera shared libraries
  - install [178](#)
- character sets [655](#)
- cluster [24](#)
  - bootstrap.properties file [403](#)
  - install generic server [344](#)
- Component Manager [425](#)
- configurations, sample [37](#)
  - baseline [38](#)
  - baseline with options [39](#)
  - demo [41](#)
  - developer [39, 40](#)
- connection point
  - configure for Application Engine [410](#)
  - convert router torouter
    - see connection point
  - create for Process Engine [409](#)
  - migrate process router to [543](#)
- connection pools
  - configure for GCD [170](#)
  - verify for GCD [180](#)
- container-managed authentication
  - and SSO for Application Engine [349, 356](#)
  - impact on custom application upgrade [604](#)
  - WebLogic patch requirements [125](#)
- Content Engine
  - COM API custom applications [563](#)
  - configure API for Process Engine [279, 298, 523, 535](#)
  - configure content-based retrieval [397](#)
  - connection pools, verify [180](#)
  - install [160](#)
  - Java API custom applications [587](#)
  - remove [628](#)
  - Setup program [161](#)
  - silent install [176](#)
  - specify for Application Engine [333, 348](#)
  - uninstall [628](#)
  - upgrade [487](#)
  - Upgrader Tool [500](#)
  - verify deployment [181](#)
  - verify installation [236, 263](#)
  - web service [412](#)
  - WebService custom applications [578](#)
  - XML upgrade status file [501](#)
- Content Engine COM API, custom applications [563](#)

Content Engine Web Service, custom applications [578](#)

Content Java API, custom applications [587](#)

Content Search Engine

- CBR Enabled flag [401](#)
- CBR indexing [401](#)
- CBR locale [400](#)
- class properties [401](#)
- collections [153](#)
- collections directory [397](#)
- configure content-based retrieval (CBR) [153](#)
- create accounts [60](#)
- FileNet Styleset [155](#)
- FileNet\_FileSystem\_PushAPI [155](#)
- install silently [147](#)
- K2 Administrative Users [157](#)
- K2 Broker Server [400](#)
- K2 Dashboard documentation [147](#)
- K2 Dashboard shortcut [149](#)
- K2 Index Server [154](#)
- K2 Search Server [155](#)
- K2 Server [155](#)
- K2 Ticket Server [155](#)
- launch K2 Dashboard [154](#)
- multi-server K2 configuration [447](#)
- set Autonomy K2 Administration Security [156](#)
- Verity Domain Config. tab (VDC) [399](#)

content-based retrieval, configure [153](#)

current release, verify before upgrade [486](#)

custom applications, upgrade [563](#)

## D

data sources

- ODBC for Process Engine [98](#)
- verify for GCD [180](#)

database

- dedicated (DB2) [35](#)
- dedicated (Oracle) [34](#)
- dedicated (SQL Server) [32](#)
- engine requirement [37](#)
- GCD (DB2) [96](#)
- GCD (Oracle) [87](#)
- GCD (SQL Server) [82](#)
- install considerations [32](#)
- local [32](#)
- local vs. remote [32](#)
- remote [32](#)
- shared (DB2) [35](#)
- shared (Oracle) [34](#)

- shared (SQL Server) [32](#)
- storage area [252](#)
- supported character sets [653](#)
- supported types [32](#)
- user (DB2) [35](#)
- user (Oracle alias) [34](#)
- user (SQL Login) [33](#)

DB2

- client install [103](#)
- database collocation restriction [37](#)
- dedicated database [35](#)
- install for Content Engine [92](#)
- install for Process Engine [92](#)
- JDBC drivers, WebLogic [115](#)
- shared database [35](#)
- tablespace, create object store [259](#)

default time mask, Solaris Process Engine [288](#)

deploy from a test environment [24](#)

Deployment Guide [24](#)

Desktop 3.1, collocation issues [429](#)

directory server, see authentication

DNS forwarder, enable [54](#)

document

- check in (verify new object store) [263](#)
- check out (verify new object store) [263](#)
- create (verify new object store) [263](#)

documentation

- Application Engine configuration [350](#)
- configure site preference settings (AE bootstrap) [405](#)
- enable search (JBoss) [141](#)
- enable search (WebLogic) [135](#)
- enable search (WebSphere) [130](#)
- install for FileNet P8 Platform (JBoss) [141](#)
- install for FileNet P8 Platform (WebLogic) [135](#)
- install for FileNet P8 Platform (WebSphere) [130](#)
- server upgrade [478](#)
- server URL [270](#), [291](#), [310](#), [328](#), [522](#), [533](#)
- server verification [485](#)

domain

- configure more than one [29](#), [60](#)
- controller (Windows) [44](#)
- DNS (Domain Name Service) [25](#)
- issues for container-managed authentication [376](#)
- Sun Java System Directory Server [31](#)
- users group [261](#)
- Verity (K2) [60](#), [400](#)
- WebLogic [110](#), [124](#), [161](#), [376](#)

- Windows 30
- Windows NT 30
- domain, FileNet P8
  - configure permissions 191
  - create 187
- domain, Windows, configure multiple 29

## E

- EJB transport 412
- Enterprise Manager
  - enable SSL 414
  - file storage area 264
  - install 183
  - install additional 427
  - run from developer workstations 40
  - upgrade 559
- extensible authentication ports 643

## F

- farms 24
- file storage area
  - defined 252
  - display status 264
  - specify 260
- file store, see file storage area
- File Tracker
  - install 433
  - uninstall 434
  - upgrade 560
- FileNet P8 documentation
  - configure for Application Engine help 350
  - configure for Content Engine help 166
  - configure for Process Engine help 270
  - remove 624
  - upgrade 478
- FileNet P8 Platform documentation
  - install (JBoss) 141
  - install (WebLogic) 135
  - install (WebSphere) 130
- FileNet Styleset, Content Search Engine 155
- firewall 28
- Fix Packs and Test Fixes
  - install for optional components 457
  - minium level required 462
- fixed storage area 252

## H

- high availability environments 24

## I

- IIO protocol 412
- Image Services Integration, ISRA 557
- install
  - silent, encrypting passwords for 642
- installation
  - checklist 42
  - configuration and startup tasks 396
  - optional tasks 422
  - planning 23
  - prerequisite tasks 53
  - sample configurations 37
- IP address
  - AIX host file 25
  - Application Engine SSL host 416
  - Content Engine server 164, 188
  - Content Search Engine server 147
  - database server 46, 170, 236, 241
  - domain controller 44
  - Process Engine 28
  - Solaris host file 25
  - static for each server 28
  - UNIX /etc/hosts file 25
  - Windows Active Directory configuration 161
  - Windows host file 25
- iPlanet, see Sun Java System Directory Server
- IS errorlog message redirection (Windows) 283
- isolated region, create 408
- ISRA
  - install servlet 452
  - remove servlet 639
  - SSL configurations 451
  - supported version 557
  - upgrade servlet 618

## J

- JBoss
  - configure for Application Engine 126
  - configure for Content Engine 118
  - deploy Application Engine 394
  - remove Application Engine 637
- JDBC drivers 106, 170
  - WebLogic, install for DB2 115
  - WebLogic, install for Oracle 116
  - WebLogic, install for SQL Server 115

## K

- K2, see Content Search Engine
- Kerberos 29, 412

**L**

LANG settings, Solaris Process Engine install [288](#)  
 localized environments, install into [24](#)

**M**

Microsoft Office, integration install [429](#)  
 multi-server configuration [443](#)  
   Application Engine [443](#)  
   Content Engine [437](#)  
   Content Search Engine [447](#)  
   Enterprise Manager [427](#)  
   high availability and clusters [24](#)

**N**

network security considerations [28](#)  
 NLS, see character sets  
 non-English environments, install into [24](#)

**O**

object store  
   create [258](#)  
   DB2 database [96](#)  
   name requirements [258](#)  
   prepare for upgrade [504](#)  
   upgrade [506](#)  
   verify new [263](#)  
 ODBC data source for Process Engine [267](#)  
 ODBC datasource [530](#)  
 operating system considerations for install [25](#)  
 Oracle  
   client install [100](#)  
   database collocation [37](#)  
   dedicated database [34](#)  
   install for Content Engine [85](#)  
   install for Process Engine [85](#)  
   JDBC drivers, WebLogic [116](#)  
   shared database [34](#)  
   tablespace, create object store [259](#)  
   user name, create object store [259](#)  
   version number [524](#)

**P**

passwords, encrypting for silent installations and upgrades [642](#)  
 patches, see Fix Packs and Test Fixes  
 Performance Tuning Guide [24](#), [462](#)  
 plan installation [23](#)  
 ports [643](#)

Process Analyzer, upgrade procedure [558](#)  
 Process Engine [524](#)  
   configure Content Engine API [279](#), [298](#), [523](#), [535](#)  
   create users and groups [286](#), [305](#), [323](#)  
   enable port 32771 [301](#)  
   enable port 32771, upgrade [525](#)  
   Java API custom applications [601](#)  
   passwords, reset [300](#), [317](#), [335](#)  
   remove [631](#)  
   remove (UNIX) [633](#)  
   remove (Windows) [631](#)  
   SQL scripts [531](#)  
   start or restart software [539](#)  
   static IP address [28](#)  
   time zone settings [321](#)  
   uninstall [631](#)  
   update security [542](#)

Process Java API, custom applications [601](#)  
 process router  
   see connection point

Process Simulator, upgrade procedure [558](#)

Process Task Manager [340](#)

publishing components  
   install and configure [423](#)  
   upgrade [562](#)  
 publishing templates [423](#)

**R**

realm  
   alternatives to Realm.Users and Realm.Groups [579](#)  
   Application Engine container-managed authentication [376](#)  
   authentication [29](#)  
   configure more than one [29](#), [421](#)  
   Content Engine AllRealms property [658](#)  
   Content Engine MyRealm property [664](#)  
   Content Engine Realm class [657](#)  
   IBM Tivoli Directory Server [32](#)  
   Microsoft ADAM [30](#)  
   Novell eDirectory [31](#)  
   object store administrator [261](#)  
   Realm and Domain interface methods [588](#)  
   retrieve via custom applications [577](#)  
   Sun Java System Directory Server [31](#)  
   WebLogic domain for Content Engine [111](#)  
   Windows Active Directory domain [30](#)  
 redirect IS errorlog messages [283](#)



reenable Oracle password complexity verification [283](#), [526](#), [536](#)

Release Notes, FileNet P8 [24](#)

remove

Application Engine (JBoss) [637](#)

Application Engine (WebLogic) [636](#)

Application Engine (WebSphere) [634](#)

Content Engine [628](#)

FileNet P8 documentation [624](#)

fnsr and oracle users from ORA\_DBA [284](#)

ISRA servlet [639](#)

Process Engine (UNIX) [633](#)

Process Engine (Windows) [631](#)

Rendition Engine [423](#), [562](#)

Rendition Engine Installation and Upgrade Guide [25](#), [462](#)

requirements for hardware and software [24](#), [478](#)

restore custom modifications [301](#), [318](#), [336](#), [524](#)

## S

search documentation

enable (JBoss) [141](#)

enable (WebLogic) [135](#)

enable (WebSphere) [130](#)

security

authentication [55](#), [56](#), [59](#)

info preference, SSL [404](#)

install considerations [28](#)

specify accounts [60](#)

SSL, set up for Application Engine [416](#)

SSL, set up for Content Engine [412](#)

update for PE [542](#)

service packs and interim fixes

install and upgrade [20](#)

silent install

Content Engine [176](#)

Content Search Engine [147](#)

encrypting passwords for [642](#)

Process Engine (AIX) [316](#)

Process Engine (HP-UX) [334](#)

Process Engine (Solaris) [299](#)

Process Engine (Windows) [280](#)

Single Sign-On, see SSO

Site preference settings

documentation URL [405](#)

name and location [405](#)

SQL scripts on Process Engine [531](#)

SQL Server

client install [98](#)

collocation [37](#)

dedicated database [32](#)

install for Content Engine [81](#)

install for Process Engine [81](#)

JDBC drivers, WebLogic [115](#)

shared database [32](#)

SSL

additional procedure for AE on WebSphere [419](#)

Java applets [419](#)

security info preference [404](#)

set up for Application Engine [416](#)

set up for Content Engine [412](#)

SSO

Application Engine (WebSphere), container-managed authentication [356](#)

edit configuration files, Application Engine (WebLogic) [373](#)

edit configuration files, Application Engine (WebSphere) [359](#)

Kerberos [29](#)

planning considerations [29](#)

requirements [29](#)

Statement Cache Size [191](#)

subfolder, create, (verify new object store) [263](#)

Sun Java System Directory Server [56](#)

Sun ONE, see Sun Java System Directory Server

symmetric encryption, unlimited strength .jars [425](#)

System Manager [456](#)

## T

T3 protocol [412](#)

time mask, Solaris Process Engine [288](#)

troubleshooting

installation problems [24](#)

Troubleshooting Guide [24](#)

tuning product performance [24](#), [462](#)

## U

uninstall, see remove

UNIX

remove Process Engine [633](#)

unlimited strength encryption (symmetric encryption) [425](#)

upgrade

Application Engine [551](#)

Application Engine, additional tasks [556](#)

Application Integration [560](#)

checklists [470](#)

Content Engine [487](#)

- custom applications [563](#)
- documentation [478](#)
- File Tracker [560](#)
- FileNet Enterprise Manager [559](#)
- ISRA Servlet [617](#)
- plan [461](#)
- Process Engine (UNIX) [517](#)
- Process Engine (Windows) [527](#)
- publishing components [562](#)
- release levels required to [486](#)
- silent, Content Engine [498](#)
- silent, encrypting passwords for [642](#)
- silent, Enterprise Manager [559](#)
- Upgrader Tool, Content Engine [500](#)
- User Token
  - configure [404](#)
  - Crypto key file path (AE) [353](#)
- users and groups, create
  - AIX Process Engine [305](#)
  - HP-UX Process Engine [323](#)
  - Solaris Process Engine [286](#)
- UTCryptoKeyfile.properties [353](#)
- UTF-8 encoding, Application Engine (WebSphere) [361](#), [364](#), [368](#)

## V

- verify /etc/services settings [283](#)
- verify current release before upgrading [486](#)
- verify TCP/IP settings [283](#)
- Verity K2, see Content Search Engine
- version verification tools [486](#)
- virtual host configuration, WebSphere Application Engine install [386](#)
- VistaPubDir shared folder for publishing [423](#)
- vwlog [520](#), [530](#)
- vwtool [518](#), [527](#)

## W

- Web Application Toolkit, custom applications [604](#)
- web farm
  - bootstrap.properties file [403](#)
  - install generic server [344](#)
- web services
  - Content Engine [412](#)
  - Process Engine [412](#)
- Web Services 3.1, collocation issues [429](#)
- web.xml
  - edit for Application Engine SSO (WebLogic) [373](#)

- update for Content Search Engine [149](#)
- upgrade Application Engine [551](#)
- WebLogic
  - configure for Application Engine [124](#), [126](#)
  - configure for Content Engine [110](#)
  - deploy Application Engine [373](#), [391](#)
  - remove Application Engine [636](#)
- WebSphere
  - configure for Application Engine [123](#)
  - configure for Content Engine [106](#)
  - custom setup [347](#)
  - deploy Application Engine [356](#), [385](#)
  - remove Application Engine [634](#)
  - SSL setup, additional procedure [419](#)
  - typical setup [347](#)
  - UTF-8 encoding [361](#)
  - UTF-8 encoding (Application Engine) [364](#), [368](#)
  - virtual host configuration (Application Engine) [386](#)
- Windows
  - remove Process Engine [631](#)
- Workplace
  - Application Engine install option [347](#)
  - Application Integration Toolkit custom applications [609](#)
  - Application Integration, collocation issues [429](#)
  - Application Integration, install [429](#)
  - Application Integration, upgrade [560](#)
  - application name [52](#)
  - avoid logon failures in multi-forest domains [114](#)
  - configuration [37](#)
  - configure SSL access to [416](#)
  - define workflow features for users [421](#)
  - deployment in farm or cluster [351](#)
  - deployment types [349](#)
  - design publishing templates for users [420](#)
  - design searches for users [420](#)
  - directory, back up for upgrade [551](#)
  - enable access to tasks and work items [410](#)
  - enable for access to IS documents [451](#), [617](#)
  - File Tracker install [433](#)
  - File Tracker upgrade [560](#)
  - install source code [347](#)
  - remove [552](#), [634](#), [636](#)
  - set bootstrap properties [403](#)
  - set documentation server URL [405](#)
  - set SSO proxy parameters for [373](#)
  - specify Application Engine administrators [406](#)

**X**

XA and non-XA transactions [259](#)

XML upgrade status file [501](#)



---

## Notices

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

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

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

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

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

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106-0032, Japan

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

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

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

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

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

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

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

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

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

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

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

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

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://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 (<sup>®</sup> or <sup>™</sup>), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

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

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

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.









Program Number: 5724-R76, 5724-R81

Printed in USA

GC31-5488-04

