


High Availability Technical Notice



High Availability Technical Notice

Note

Before using this information and the product it supports, read the information in "Notices" on page 183.

Table of contents

Revision log	12
About this document	13
Intended audience	13
Typographical conventions	14
Acronyms	15
About IBM FileNet documentation	16
Copy Web documents into the help system and make them searchable	16
Links to additional information	16
Gather reference documentation	17
Installation reference documents	17
Other available documentation	18
Autonomy K2 software documentation for configuring the FileNet P8 Content Search Engine	20
Access IBM FileNet documentation, compatibility matrices, and fix packs	20
Feedback	20
IBM FileNet P8 high availability	21
Terms and concepts	21
What is high availability?	21
Disaster recovery and high availability	21
Server consolidation and high availability	21
Industry terminology	22
The FileNet P8 architecture in a highly available environment	23
Choosing between a farm and a cluster	24
Server farms	24
Server clusters	25
Geographically dispersed clusters	26
Cluster configurations	26
Asymmetric clusters	27
Symmetric clusters	27
N + 1 clusters	27
N to M clusters	28
High availability products	30
Examples of high availability products	30
Load balancing/farming	30
Server clusters	30
Symantec offers its Veritas Cluster Server product for all of these platforms	30
High availability products tested with FileNet P8	31
Load balancing/farming	31
WebLogic clusters (farms)	32

WebSphere clusters (farms)	32
JBoss clusters	33
Server clusters	33
Microsoft cluster server	34
Veritas cluster server	34
Database services	34
Oracle RAC	35
Plan and prepare for high availability	36
Plan the installation	36
Plan your installation	37
Content Engine	37
Content Engine procedure overview	37
Content Engine cache refresh delays	38
Content Search Engine	38
General information about CSE installation	38
Content Search Engine procedure overview	38
Process Engine	39
Process Engine procedure overview for UNIX	39
Process Engine procedure overview for Windows	41
Application Engine/Workplace XT	41
Application Engine/Workplace XT procedure overview	42
Rendition Engine	42
Rendition Engine procedure overview	42
CFS-ICI	43
CFS-ICI procedure overview	43
eForms	44
eForms procedure overview	44
Process Analyzer	44
General information about PA installation	44
PA Configuration Files	44
Process Analyzer procedure overview	44
Other highly available components	45
Records Manager	45
Procedure Overview for Records Manager	46
IBM FileNet Connector for SharePoint Web Parts, version 2.2	46
IBM FileNet Connector for SharePoint Document Libraries, version 2.2	46
Definition of roles	47
Installation administrator (IA)	47
IT administrator (ITA)	47
Security administrator (SA)	47
Database administrator (DBA)	47
Application server administrator (ASA)	48
Use the installation worksheet	48
Perform required preparation tasks	49
Installation administrator tasks	50

Content Engine	50
Content Engine installation location	50
Content Search Engine	50
Process Engine	50
Application Engine/Workplace XT	50
Rendition Engine	50
CFS-ICI	50
eForms	51
Process Analyzer	51
<i>IBM FileNet Connector for SharePoint Web Parts</i>	51
IBM FileNet Connector for SharePoint Document Libraries	51
Application server administrator tasks	52
JBoss clustering of Content Engine and Application Engine	52
Application Engine/Workplace XT	52
Configure the application server	52
Configure load-balancer or proxy server	52
Plan for Workplace/Workplace XT shared settings	53
Plan for Component Manager	54
Content Engine	54
Install JDBC drivers for your selected database	54
Install and configure the application server for high availability.	54
Configure the application server for Content Engine.	55
(WebSphere only) Configure Light weight Third Party Authentication (LTPA)	55
Configure load-balancer or proxy server	55
CFS-ICI	56
Pre-installation tasks	56
IT administrator tasks	57
Content Search Engine	57
(UNIX only) Update the k2_os_user profile	57
Create required shared directories	57
Application Engine/Workplace XT	58
Create a local or shared directory for the shared configuration files	58
Process Engine	58
Configure hardware load balancer	58
Configure the Process Engine nodes hosts file	58
Configure your operating system for Process Engine Cluster (UNIX)	60
Install and configure Veritas Cluster Server (UNIX)	61
Install and configure Veritas Cluster Server (Windows)	62
Install and configure Microsoft Cluster Server	62
Process Analyzer	63
Create a shared Temp directory	63
Configure Microsoft Cluster Server (MSCS) PA Resource Group	63
Security administrator tasks	64
Rendition Engine	64
Assigning the <i>FNRE_Admin</i> user	64
Records Manager	64
(UNIX) Verify write permission on the /bin directory	64
Database administrator tasks	65
Content Engine	65

Install and configure the database	65
Process Engine	65
Oracle RAC configuration	65
Process Analyzer	65
PA and Oracle RAC	65
Rendition Engine	66
Configure remote databases	66
Install and configure for high availability	67
Overview	68
Install a highly available Content Engine	70
Installation and deployment tasks	70
Creating object stores in FileNet P8 4.5	73
Enable application server connections after database failover	74
Deployment verification	76
File store high availability configuration	78
Install a highly available Process Engine	79
Install Process Engine in a server farm (UNIX, Windows)	79
Install Process Engine in a Veritas cluster server (UNIX)	81
Installation tasks	81
Process Engine installation verification	83
Install Process Engine in a Veritas cluster server (Windows)	85
Installation tasks	85
Process Engine installation verification	88
Install Process Engine in a Microsoft cluster server	89
Deployment tasks	89
Deployment verification	91
Process Engine database reconnection	92
Examples of errors returned on database disconnect	92
Known issues	93
Install a highly available Application Engine/Workplace XT	94
Install and configure Application Engine/Workplace XT	94
Installation tasks	94
Application Engine/Workplace XT deployment verification	97
Enabling the Process Engine Component Integrator	98
Install a highly available Rendition Engine	99
Install and configure Rendition Engine	99
Installation tasks	99
Rendition Engine installation verification	100
Install a highly available Content Search Engine	101
Install and configure Content Search Engine	101
Installation tasks	101

Install a highly available CFS-ICI	104
Install and configure CFS	104
Installation tasks	104
Deployment verification	107
Install a highly available eForms for P8	108
Install and configure eForms	108
Installation and deployment tasks	108
eForms deployment verification	109
Install a highly available Process Analyzer	110
Install and configure Process Analyzer	110
Installation tasks	110
Install a highly available Records Manager	112
Install and configure Records Manager	112
Pre-installation tasks	112
Installation tasks	113
Deployment verification	116
Install a highly available IBM FileNet Connector for SharePoint Web Parts 2.2	117
Install and configure IBM FileNet Connector for SharePoint Web Parts 2.2	117
Installation tasks	117
Installation verification	120
Install a highly available IBM FileNet Connector for SharePoint Document Libraries 2.2	
122	
Install and configure IBM FileNet Connector for SharePoint Document Libraries 2.2	122
Installation tasks	122
Installation verification	126
Plan and prepare for high availability upgrades	130
Plan the upgrade	130
Plan your upgrade	131
Content Engine	131
Content Engine upgrade procedure overview	131
The 3.5.x/4.5 Content Engine farm/cluster after upgrade	131
Object store availability in a mixed 3.5.x/4.5 environment	132
Process Engine	132
Process Engine upgrade procedure overview	132
Application Engine/Workplace XT	133
Application Engine/Workplace XT upgrade procedure overview	133
Records Manager	133
Rendition Engine	133
Rendition Engine upgrade procedure overview	133
IBM FileNet Connector for SharePoint Document Libraries	133

IBM FileNet Connector for SharePoint Web Parts	134
Perform required preparation tasks	135
Installation administrator tasks	136
Content Engine	136
Process Engine	136
Application Engine/Workplace XT	136
Rendition Engine	136
Application server administrator tasks	137
Content Engine	137
Configure the application server for Content Engine (upgrades from 3.5.x)	137
Application Engine/Workplace XT	137
IT administrator tasks	138
Process Engine	138
Configure the Process Engine nodes hosts file	138
Security administrator tasks	140
Content Engine	140
Specify IBM FileNet P8 Accounts for upgrades	140
Records Manager	140
(UNIX) Verify write permission on the /bin directory	140
Database administrator tasks	141
Content Engine	141
Create the GCD data source for Content Engine (upgrades from 3.5.x)	141
Upgrade and configure for high availability	142
Upgrade a highly available Content Engine	143
Content Engine upgrade	143
Upgrade a highly available Process Engine	146
Migrating Process Engine 3.5.x clusters to Process Engine 4.5 farms	146
Process Engine pre-upgrade tasks	150
Upgrading server farms (UNIX, Windows)	151
Upgrade tasks	151
Upgrading Process Engine in a Veritas cluster server (UNIX)	153
Upgrade tasks	153
Upgrade Process Engine in a Veritas or Microsoft cluster server (Windows)	156
Upgrade tasks	156
Upgrade a highly available Application Engine/Workplace XT	158
Upgrade tasks	158
Upgrade a highly available Rendition Engine	160
Rendition Engine upgrade	160
Uninstall previous version of Rendition Engine	160
Upgrade Rendition Engine	160

Rendition Engine upgrade verification	160
Upgrade a highly available CFS-ICI	161
Upgrade a highly available Records Manager	162
Upgrade Tasks	162
Appendixes	164
Appendix A – Setting up a secure native mode domain installation	165
Configure the domain controller	165
Configure node 1 and node 2 servers	166
Appendix B – Connecting to a highly available Content Engine	168
Administrative applications	168
User applications	168
Content Engine in an application server cluster using EJB transport	169
Connecting using Content Engine Web Service Transport (CEWS)	170
Appendix C – Using load balancers with IBM FileNet P8	171
Load balancer support for IBM FileNet P8	171
Configure the Time-To-Live value for load balancers that use dynamic DNS servers	174
Appendix D – Recompiling the FileNet Process Engine Veritas agent	176
Appendix E – Managing Workplace/Workplace XT settings in a highly available environment	178
Appendix F - Installation and upgrade worksheet	181
Notices	183
Trademarks	185
U.S. Patents Disclosure	186

Revision log

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

	Revision
05/09	<p>Added requirements for TCP monitoring in “Configure hardware load balancer” on page 58.</p> <p>Added “Enabling the Process Engine Component Integrator” on page 98.</p> <p>Added information about FileNet Enterprise Manager in “Appendix B – Connecting to a highly available Content Engine” on page 168.</p> <p>Updated list of Content Engine load balancer requirements in “Appendix C – Using load balancers with IBM FileNet P8” on page 171.</p>
12/08	<p>Update with information and procedures for the following FileNet components:</p> <ul style="list-style-type: none"> • “Install a highly available Content Search Engine” on page 101. • “Install a highly available CFS-ICI” on page 104. • “Install a highly available Process Analyzer” on page 110. • “Install a highly available Records Manager” on page 112. • “Install a highly available IBM FileNet Connector for SharePoint Web Parts 2.2” on page 117. • “Install a highly available IBM FileNet Connector for SharePoint Document Libraries 2.2” on page 122. • “Upgrade a highly available CFS-ICI” on page 161. • “Upgrade a highly available Records Manager” on page 162.
11/08	Initial revision of the <i>IBM FileNet P8 High Availability Technical Notice, Version 4.5</i> .

About this document

This document provides high availability installation and upgrade information about IBM FileNet P8 Platform, which includes Application Engine (AE), Content Engine (CE), and Process Engine (PE) as well as a number of add-on products such as Rendition Engine (RE) and eForms.

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

To download these documents from the IBM FileNet support Web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Intended audience

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

The procedures in this document require experience with:

- IBM FileNet P8 system installations.
- Installing and configuring application server clusters or farms.
- Setting up and configuring fail-over configurations on supported high availability software. See [“High availability products” on page 30](#).
- Setting up and configuring hardware or software load balancers.
- Supported operating systems and hardware platforms, volume management software, and database platforms.

Typographical conventions

This document uses the following typographical conventions:

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

Convention	Usage	Example
“Text with quotation marks”	User-interface fields that do not use initial capitalization and document headings referenced within a document	See the “Part number” field for the part number. For more information, see “Typographical conventions” on page 14.
UPPERCASE	Case-sensitive text, where uppercase text is required.	Copy the file by entering the following command: <code>COPY filename</code>

Acronyms

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

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

About IBM FileNet documentation

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

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

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

Copy Web documents into the help system and make them searchable

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

For more information, see:

- “Other available documentation” on page 18
- “Install IBM FileNet P8 Platform documentation” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
- “Update documentation search index” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

Links to additional information

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

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

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

Gather reference documentation

Following are two tables with information about the IBM FileNet P8 documents that are available as part of the FileNet P8 release. To download these documents from the IBM FileNet support Web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Installation reference documents

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

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

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

Other available documentation

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

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

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

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

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

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

Access IBM FileNet documentation, compatibility matrices, and fix packs

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

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

Customer support

For information about contacting customer support:

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

Feedback

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

IBM FileNet P8 high availability

Terms and concepts

What is high availability?

High availability is the ability to provide a service to an end-user with as little perceived downtime as possible. If a system component fails for any reason, the high availability solution ensures that another component takes over for the failed component, and that the newly composed system will maintain the same machine identifications (host names and IP addresses) as the system prior to failure, minimizing the disruption to the user.

This technical notice addresses only the hardware and software issues. Other issues that might impact high availability at a customer's system include: human error, process problems, and external issues.

NOTE High availability does not mean that a service is guaranteed to always be available. Analysts describe a range of high availability targets, from the so-called “five nines” availability, with 99.999% up-time, at the high end, to basic availability at 95%. Five nines availability translates to five minutes or less downtime in a full year of 24 by 7 operations. By contrast, 99% availability allows up to 87 hours of downtime per year, and 95% allows up to 436 hours, or 18 days, of downtime, with the cost of providing high availability increases exponentially as the target moves from 95% to 99% to 99.999%.

Disaster recovery and high availability

Disaster recovery solutions provide for business continuity in the face of natural or man-made disasters that cause the loss of an entire production system, whereas high availability solutions provide business continuity in the face of localized failures, such as a single server failure or hard disk crash.

While the goal of both high availability and disaster recovery solutions is the same—keeping your IBM FileNet P8 system available for continued business operations—the solutions themselves are quite different. For more information, see the *IBM FileNet P8 Disaster Recovery Technical Notice*. To download this document from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Server consolidation and high availability

Server consolidation is the process of combining software from two or more servers on a single, more powerful server to make server maintenance easier and thus lower the total cost of ownership.

One way of doing this is logical partitioning, where one server is configured to act like several separate physical servers. IBM's LPARs product for AIX® is one example.

If you plan to implement logical partitioning in an HA environment you must take into account the distribution of the IBM FileNet P8 components across the servers. If a single physical failure on a consolidated server can take down two or more logical servers running in parallel on that server, then redundant servers for a server farm or server cluster should not co-reside on that consolidated server unless the risk of that failure is deemed to be exceptionally low. Otherwise, a single point of failure can take down an entire farm or cluster, violating a central premise for high availability.

Industry terminology

Vendors and industry analysts use a variety of differing terms to describe high availability concepts. Most vendors refer to both server farms and server clusters as “clusters”, for instance, which can be confusing. This document uses two different terms, server farms or server clusters, to distinguish these two different technologies.

For more information, see:

- [“Server farms” on page 24](#)
- [“Server clusters” on page 25](#)

Table 1: Industry terms mapped to the terminology used in this document.

Vendor or Analyst	Term for server farm	Term for server cluster
BEA	WebLogic cluster	N/A
Gartner Group	Web server cluster	Application server cluster, database cluster
IBM	WebSphere server group of clones; also cluster	Cluster
Microsoft®	NLB cluster; also cluster farm	Server cluster; also cluster server
Symantec	Parallel group	Failover group

NOTE IBM documentation refers to the term “Cluster Failover” in reference to WebSphere Clones (farms). For all intents and purposes this can be viewed as “Session Failover” in the scope of this document.

The FileNet P8 architecture in a highly available environment

An ECM system typically consists of multiple tiers. For example, a system might have a web server tier, an application server tier for business logic, a content repository tier for managing content, and a data storage tier. A chain is only as strong as its weakest link, so a high availability solution requires a high availability configuration for every tier. The different tier types and server types require different technical solutions for high availability, primarily server farms with no sharing or server clusters with shared storage.

In a highly available IBM FileNet P8 environment however, components such as Application Engine and Content Engine are each installed and deployed as all-inclusive applications (presentation and business logic components) on multiple running instances.

Example:

The Application Engine architecture does not separate the images and static content of the Workplace application on separate server instance from the servlets and Java server pages of the Workplace application. Instead all these components are deployed on one server instance.

Furthermore, neither Content Engine nor Application Engine require any applications or components be running on any server that is not intended to run the Content Engine or Application Engine itself.

Example:

When you use a load-balancer or proxy server in a configuration of IBM FileNet P8 components, you are not required to install and run Component Manager on the load-balancer or proxy server.

For more information on the FileNet P8 architecture, see the *IBM FileNet P8 platform System Overview*. To download this document from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Other essential HA building blocks include:

- A reliable and tested backup process.
- Hardened servers with uninterruptible power supplies, RAID disk arrays, redundant network interface cards, auto-restart capabilities, and other high availability features.

Choosing between a farm and a cluster

All IBM FileNet web applications adhere to industry standards for HA farming. To be consistent Process Engine supports that HA model as well, in addition to clustering.

Generally, load-balanced server farms have a scalability advantage, exhibit better server utilization, and recover more quickly from server failures when compared to active-passive server clusters.

Server farms

A server farm is a group of identical servers accessed through hardware or software load balancing technology. All the servers are active, provide the same set of services, and are effectively interchangeable. A load balancer distributes incoming client requests over the servers in the group.

Server farms are best suited to server tiers such as Content Engine and Application Engine web applications that are processing-centric rather than data-centric, as all the servers in the farm are clones of each other. Processing logic for these components does not change often, making it easy to keep all the servers identical..

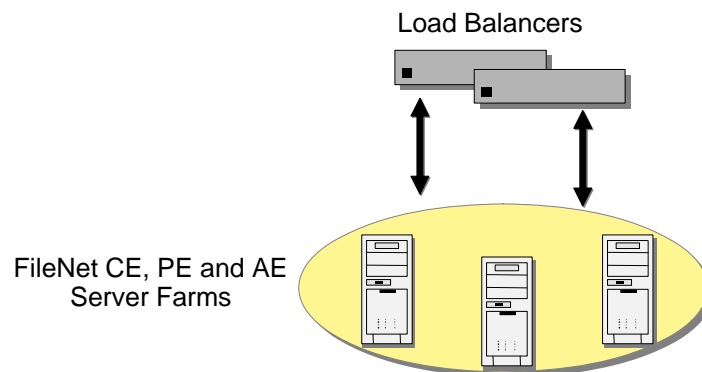


Figure 1: The basic server farm setup.

A load-balanced server farm provides both better availability and better scalability than a single server. When a server fails, the load balancer automatically detects the failure and redirects user requests to another server in the farm, thereby keeping the site available. Administrators can increase system performance and capacity by adding servers to the farm.

With a hardware-based load balancing solution, redundant load balancers are required to avoid a single point of failure. The software-based load balancers are typically designed to avoid a single point of failure by running the network load balancing software on each server in the farm.

For a brief discussion about available load balancing/farming products, see [“Examples of high availability products”](#) on page 30.

Server clusters

Server clusters are based on the concept of shared software configuration data storage, making it suitable for data-centric tiers, such as file servers and data servers, because their data content is constantly changing.

Server hardware and software vendors offer vendor-specific server clustering products as their high availability offering for these kinds of data-centric servers.

- Two or more servers share a highly available disk array for data storage, shown in the figure below. The array incorporates redundant copies of the data, but appears as a single disk resource to the servers, thereby avoiding the need for data replication between servers. The servers may each have their own local disk for static storage of operating system, utilities, and other software.
- A common set of applications run on each server.
- Server clients see the cluster as a single virtual server.
- If one of the servers fails, the other server picks up the workload of the failed server (failover). When the failed server is repaired and ready to run again, the workload is shifted back over from the other server (failback). In some configurations, the repaired server simply becomes the new backup server, and no failback is required.

NOTE IBM FileNet P8 does not support session replication across cluster nodes.

- In most server clusters, only one server is actively serving clients at a time. This is called an active-passive configuration. Some cluster server products also support another mode, called an active-active configuration. In this mode, all the servers in the cluster can be actively sharing part of the workload at the same time. It typically requires an application designed to

partition data sets among the servers to avoid data integrity problems resulting from concurrent updates to the same data from multiple servers.




Figure 2: Basic server cluster using RAID storage.

Server clusters typically communicate through a broadcast or share a central repository to keep track of cluster information and cluster node status. For a brief discussion about available clustering products, see [“High availability products tested with FileNet P8” on page 31](#).

Geographically dispersed clusters

Data replication and geographically-dispersed servers should be reserved for a disaster recovery configuration, where data replication and the difficulties introduced by wide area networking cannot be avoided.

In “geographically-dispersed” active-passive clustering configurations, the active and passive nodes within a single cluster are separated geographically from each other. The passive node might have its own local data copy which tracks the production copy via some form of replication over the connecting wide area network.

For more information, see the *IBM FileNet P8 Disaster Recovery Technical Notice*. To download this document from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Cluster configurations

Cluster software vendors support several distinct types of cluster configurations. Using Symantec's terminology, the notable configurations tested by the FileNet P8 4.5 Platform are:

- Asymmetric clusters
- Symmetric clusters
- N+1 clusters
- N-to-M clusters

Asymmetric clusters

Asymmetric clusters are asymmetric in the sense that the cluster includes both active and passive servers. An asymmetric 1-to-1 cluster has one active server and one passive server. This is the simplest form of server cluster, but also the most expensive, in effect doubling the number of servers required for a given workload. There is no drop-off in performance after a failover in an asymmetric 1-to-1 cluster, assuming the passive server is identical in capability with the active server.




Figure 3: Asymmetric cluster.

Symmetric clusters

Symmetric clusters are termed symmetric because all the servers (also called nodes) in the cluster are active. While there are no passive servers in a symmetric cluster, the servers typically run different services. No two servers provide the same service for the same data set simultaneously, because of the difficulty of maintaining data integrity mentioned earlier.

The net result of a failover in a symmetric cluster is that both services end up running on one node. Note the potential drop in performance if both nodes are not provisioned with enough idle capacity to accommodate running both the Process Engine and database on the same node. For this reason nodes in a symmetric cluster may have to be more powerful than they would be in a non-HA environment, but no extra nodes are required, in contrast to an asymmetric cluster.

NOTE You can only set up symmetric clusters with components that can be collocated. For more information, see the *IBM FileNet P8 Hardware and Software Requirements* document.




Figure 4: Symmetric cluster.

N + 1 clusters

An N+1 cluster has N active nodes and only one (1) passive backup node. The lone passive node acts as the backup for all N active nodes. This is an asymmetric cluster, since it has a mix of

active and passive nodes, but it requires fewer nodes overall than an asymmetric 1-to-1 cluster configuration for each active node. Each node in an N+1 cluster is configured with the software for all the services supported by the cluster. This allows each node to act as the backup for all the active nodes in the cluster. The benefit is that a failed node, once it comes back online, can take on the role of backup node. This avoids the need for a failback from the former backup node to the repaired node, eliminating the disruption in service caused by a failback event.

Example of an IBM FileNet PE 4.5 N+1 cluster with three nodes:

- Process Engine server node (PE server)
- Remote database server node
- Backup server node

All three nodes are configured to run both the Process Engine and the database. If either the Process Engine or the database node fails, a failover to the backup node occurs and the appropriate software is started to take over the role. When the failed node is repaired and brought back up, it becomes the new backup node in the cluster.




Figure 5: N+1 cluster.

N to M clusters

An N to M cluster has N active nodes and M nodes configured as failover nodes. The M failover nodes can be either active or passive. Like N+1 clusters, an N-to-M cluster can be configured to avoid the need for a failback when the failed node is available for use again, in the case where a passive backup node is desired. If there is no passive node in an N-to-M cluster, a failback is presumably preferred to get the cluster back to its highest performance configuration with all nodes active. The advantage of an N-to-M cluster over an N+1 cluster is that it can accommodate multiple failures, since it has M failover targets that can take over for a failed node.

Example of an IBM FileNet P8 N to M cluster with four nodes:

- Process Engine server node (PE Server)
- Content Engine Database server node (CE Database Server)
- Process Engine Database server node (PE Database Server)
- File Store share server node (File Store Server)

If the File Store share server node fails, it could be configured to fail over to the Process Engine Database server. On a second failure, say of the Process Engine, the clustering software on the PE node could be configured to fail over to the Content Engine Database Server node.




Figure 6: N to M cluster.

High availability products

Examples of high availability products

Load balancing/farming

Microsoft Windows Network Load Balancing (NLB) and Cisco are two examples of load balancing applications.

- Microsoft Windows® contains software-based load balancing capability called Network Load Balancing (NLB). As requests come in from remote clients, the load balancing products spread out the requests across the servers to balance the workload evenly.
- Cisco offers hardware-based load balancing switches that automatically spread the incoming client workload across a farm of servers, each providing access to the same content or services.

Server clusters

Server cluster products from the major server vendors include:

- HP's MC/ServiceGuard for HP-UX
- IBM's HACMP™ for AIX
- Microsoft Cluster Server for Windows Server 2003
- Sun Cluster for Solaris.

Symantec offers its Veritas Cluster Server product for all of these platforms

High availability products tested with FileNet P8

IBM has tested the IBM FileNet P8 platform with high availability products from the following vendors:

- [“Load balancing/farming” on page 31.](#)
 - BEA Systems
 - F5 networks
 - IBM
 - JBoss
- [“Server clusters” on page 33.](#)
 - Microsoft. See [“Microsoft cluster server” on page 34](#)
 - Symantec. See [“Veritas cluster server” on page 34](#)
- [“Database services” on page 34.](#)
 - IBM DB2® UDB with Veritas Cluster Server
 - Oracle with Veritas Cluster Server
 - Oracle Real Application Clusters (RAC). See [“Oracle RAC” on page 35.](#)
 - SQL Server with Microsoft Cluster Server (MSCS)

All IBM FileNet web applications adhere to HA industry standards and number of other high availability vendors have products and configurations that should work with the IBM FileNet P8 Platform.

CAUTION If your high availability configuration differs significantly from those described here, you should to do a proof of concept test, or even a pilot project, to confirm that the configuration functions as expected.

Load balancing/farming

A number of hardware and software load-balancing products are available for server farm configurations. For more information, see [“Appendix C – Using load balancers with IBM FileNet P8” on page 171.](#)

Table 2: Tested load-balancing/farming solutions

Vendor	Notes
BEA WebLogic clusters	-
F5 Big IP	-
IBM WebSphere server group of clones	Cluster in later releases.
JBoss cluster	Also called HA partition.

- WebSphere and WebLogic

Both FileNet P8 Application Engine and Content Engine are hosted on Java application servers. Both BEA and IBM's Java application servers have built-in capabilities for providing highly available web services. Each application server product is capable of configuring a collection of server instances that function as a single entity in order to provide an application to an end-user. Both BEA WebLogic and IBM WebSphere calls this collection of server instances a Cluster. Both products function like a server farm.

- JBoss

JBoss Application Server clusters do not use a separate administrative server. In a JBoss cluster all nodes grouped together in a partition are equivalent. Applications must be deployed on each node individually unless the optional JBoss "farm" service is used, and individual nodes are started and stopped independently. For details on the JBoss farm service, see the JBoss documentation.

WebLogic clusters (farms)

Using WebLogic server you can create a network of clustered application servers to provide highly available and scalable applications that are capable of providing service in the event of a machine or hardware component failure.

WebLogic clusters are a network of WebLogic server instances that work together as a server farm to provide a highly available web service. Each instance is a fully functioning application server. Applications can be deployed or undeployed en masse to the entire cluster, thus making the cluster function as a single unit providing the same application across all machines in the cluster.

Web clients connect to the cluster through either a cluster aware Proxy application (an application provided by WebLogic which includes built-in load balancing) or a third party load balancer. In either case the web client connects to one address and is transparently redirected to a machine within the cluster.

WebSphere clusters (farms)

IBM's WebSphere Application Server (WAS) provides high availability through clusters (previously called "server groups of clones"), a collection of servers that work together as a server farm and can be managed as a single server. You can deploy, start, and stop all servers in the cluster as a single entity.

Servers in a WAS cluster are copies of an original WAS server. They function as separate individual servers, yet their configuration is based on the same template configuration from an original server. Clusters can be used to group together cloned servers. WebSphere clusters together with clones are equivalent to WebLogic clusters, both of which are server farms.

Web clients connect to the cluster through a proxy HTTP server. For IBM WebSphere the proxy HTTP server uses the IBM WebSphere HTTP plug-in. The HTTP plug-in is cluster-aware and redirects requests to the server group of cloned application servers.

JBoss clusters

JBoss servers can be grouped together into a cluster to provide high availability, but each server is managed individually. The server on each node is started and stopped individually, and applications must be deployed separately to each node unless the optional JBoss “farm” service is used.

NOTE Since the farm service is not necessary for a JBoss cluster to operate, it will not be discussed further in this document; applications will be assumed to be deployed individually on each node.

Each node in a JBoss cluster is an independent server running the cluster service and configured with the same HA partition name. Servers with the same partition name are configured separately but communicate with each other via multicast.

Web clients connect to the cluster through a proxy HTTP server such as Apache HTTP Server with the Apache Tomcat Connector. The Tomcat Connector, mod_jk, is configured with the location of each JBoss server in the cluster and redirects each request to a machine in the cluster. For more information, see JBoss documentation and [“Appendix B – Connecting to a highly available Content Engine” on page 168](#).

Server clusters

Veritas Cluster Server and Microsoft Cluster Server are two examples of similar server cluster products. Each machine in the cluster is referred to as a node. Each node in the cluster monitors the local services it is running and broadcasts this information on a private network connection. This private network connection allows all nodes in the cluster to know the status of all clustered resources. In the event that a service on one node fails, another node receives this status through the private network connection and in response, can start the service locally to maintain high availability for the service.

In addition to Oracle being supported by out-of-the-box Veritas high availability agents, a certified Veritas custom high availability agent for the Process Engine is provided as part of the IBM FileNet P8 Platform software. See [“Install Process Engine in a Veritas cluster server \(UNIX\)” on page 81](#) for more details on the Veritas high availability agent for the Process Engine.

Table 3: Tested server clusters

Vendor	Notes
Microsoft Cluster Server	Windows
Veritas Cluster Server	The Symantec Veritas product runs on a number of hardware and operating systems platforms, including: <ul style="list-style-type: none"> • AIX • HP-UX • Sun Solaris • Microsoft Windows Server

Microsoft cluster server

Microsoft offers clustering technology within its Windows operating system. Microsoft Cluster Server (MSCS) offers the ability to set up a number of Windows cluster nodes in order to provide highly available services, with functionality similar to other cluster products. The maximum number of nodes that can be clustered depends on the version of the operating system being used. For more information, consult your operating system vendor's documentation.

Veritas cluster server

Symantec provides a software package called Veritas Cluster Server (VCS) that provides a network of servers that are capable of running applications in a high availability environment with clustered storage.

VCS works by monitoring resources and applications associated with a provided service (for example an application server, database, or network storage). When a provided service goes offline on one server in the cluster, it is automatically started on another node in the cluster.

VCS Agents monitor, start, and stop services in a cluster. Agents are a middle layer between the user interface, and the services running in a cluster. Commands are given to the Agents and the Agents are responsible for fulfilling the command and verifying that everything executed without error. When you execute a command in VCS to bring a resource offline this is in effect telling the Agent to go and take the resource offline.

VCS provides a highly configurable framework for creating your own Agents to control services in a cluster. Symantec also provides, for purchase, several Agents for popular products such as Oracle, Microsoft SQL Server, and Microsoft Exchange.

Database services

The Content Engine and Process Engine both use database servers to store content. If the database is inaccessible users will be unable to add and retrieve content on the Content Engine or Process Engine. You should include any database services in a cluster as part of a highly available environment. No additional configuration steps are necessary on the FileNet P8 side to use highly available database instances for any FileNet P8 components (with the exception of Oracle Real Application Clusters as noted below).

Table 4: Tested database services

Vendor	Notes
IBM DB2 UDB with Veritas Cluster Server	-
Oracle with Veritas Cluster Server	-
Oracle RAC	See "Oracle RAC" on page 35 .
SQL Server with Microsoft Cluster Server (MSCS)	-

Oracle RAC

In addition to the traditional database server clusters FileNet has also tested IBM FileNet P8 platform components with Oracle Real Application Clusters (RAC). Oracle RAC allows for the active availability of multiple database servers running in parallel. This is in contrast to the traditional cluster topologies covered above, where only one instance was running at any given time. With Oracle RAC multiple database servers are running at the same time and provide access to the same data. The benefit of this is the ability to continue to provide high availability in conjunction with load balancing functionality. Oracle also provides a feature called Transparent Application Failover (TAF) to automatically reconnect client applications to a database in the event of a connection failure, making the failover of database sessions transparent to the application. The settings controlling load-balancing and TAF are maintained by the Oracle client.

The Oracle client settings required to enable load-balancing and TAF in an Oracle RAC configuration are not specific to any application, including IBM FileNet P8 platform components. See your Oracle documentation for more information on how to enable client settings for load balancing and/or Transparent Application Failover as desired in an Oracle RAC configuration.

NOTE For Oracle clusters the JDBC URL requires a different format from the default. For more information, see [“Configure Content Engine instances.” on page 70](#) and [“\(WebSphere and WebLogic\) Oracle RAC JDBC connection URL.” on page 73](#).

Plan and prepare for high availability

This section, Plan and Prepare for High availability, provides instructions for preparing your system to install and upgrade the core IBM FileNet P8 components in a highly available environment. Installation instructions themselves are contained in the IBM FileNet P8 Platform Installation and Upgrade Guide, and separate add on installation guides.

Plan the installation

This section describes how to prepare your environment for the installation of a highly available 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 independent software vendor (ISV) software.

Plan your installation

This topic contains general information and high level installation steps to help you plan your high availability installation.

Content Engine

For a general overview of Content Engine deployment options, see the “Plan your Content Engine deployment” topic in the *Plan and Prepare Your Environment for IBM FileNet P8* document.

Content Engine provides access and storage facilities for content. When a critical component fails, users will no longer have access to existing content or be allowed to add new content. When content is no longer accessible Application Engine and Process Engine components are also affected. This makes the health and stability of the Content Engine a key factor in creating a highly available FileNet P8 environment.

Content Engine is comprised of a number of components. The following are the most important for a highly available FileNet P8 environment:

- Content Engine enterprise application
- Content Engine Global Configuration Data - stored in the database
- Object store data - stored in the database
- File store data - stored in a highly available network share

Content Engine procedure overview

NOTE Depending on your setup/configuration there might be additional steps to the ones below relevant to your setup/configuration. For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

1. Install and configure the database/tablespace.

This step includes installing and configuring your database(s)/tablespace(s), including setting them up to be highly available if required. Follow the instructions in the *Plan and Prepare Your Environment for IBM FileNet P8* for your database requirements, and follow vendor documentation for database installation.

2. Install and configure the application server.

This step includes installing and configuring your application server, including creating and configuring your highly available application server configuration. Follow the instructions provided by your application server vendor.

3. (Content Engine using WSI) Install load balancer.
4. Install Content Engine to one of your application server nodes.
5. Configure Content Engine instances.
6. Install the latest Process Engine Client files on Content Engine servers.
7. Redeploy Content Engine to your application server configuration.

8. Complete post-deployment steps.
9. Establish the FileNet P8 domain and Global Configuration Data (GCD)
10. Create the data sources for your object store databases.
11. Enable application server connections after database failover.

Enable WebSphere/WebLogic/JBoss connection to your GCD and object store databases after failover.

12. Create an object store and verify the Content Engine Installation.

Content Engine cache refresh delays

Metadata and GCD information are cached and any changes made to these areas might not be reflected among the other Content Engine nodes until the refresh has occurred.

The default refresh delay value is:

- 2 minutes for metadata cache (not configurable)
- 30 seconds for GCD cache

The GCD cache value is configurable using FileNet Enterprise Manager via the properties page of the FileNet P8 domain in the “Server Cache” tab.

NOTE After making changes to the GCD, or to object store metadata, you should wait for the respective cache refresh interval before taking actions involving these changes. For example, after creating a class with FileNet Enterprise Manager you should wait two minutes before creating and manipulating instances of the class. This recommendation only applies to the portions of applications performing administrative operations that update metadata or the GCD, such as those within FileNet Enterprise Manager.

Content Search Engine

The IBM FileNet P8 platform relies on Autonomy's K2 Enterprise technology for its content based retrieval (CBR) functionality, also called Content Search Engine. You can install CSE in a highly available environment using Veritas Cluster Server (VCS) or Microsoft Cluster Server (MSCS) technologies.

General information about CSE installation

- CSE only supports failover clustering.
- The CSE installer needs to be run on each cluster node and supplied with the same configuration information.

Content Search Engine procedure overview

1. **NOTE** Depending on your setup/configuration there might be additional steps to the ones below relevant to your setup/configuration. For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide*. Create a Veritas or MSCS cluster.

2. Create highly available storage locations for the following:

UNIX

- (Optional) CSE binaries and configuration files
- CSE temporary files

Windows

- CSE temporary files

3. Install CSE on each cluster node, in the shared directory.
4. Change CSE configuration files to use cluster hostname.
5. Add the resources that will be under cluster control.
 - K2 Admin Service
 - K2 Dashboard Service (optional)
 - File share for Collections directory
6. Failover to each node to ensure that Content Search engine resources startup properly on each node.
7. Configure CSE per install guide instructions, for your cluster.

Process Engine

Process Engine allows users to create, manage, control, and participate in workflows. When Process Engine is unavailable, Application Engine or Content Engine cannot retrieve Process information. Under these circumstances no process work can be done; all workflows remain in their current state until Application Engine or Content Engine can reconnect to a running Process Engine.

The goal of this configuration is to allow Process Engine Services to continue running and provide service to Application Engine and Content Engine in the event that the active node in the cluster or farm fails. Making the Process Engine highly available ensures that Application Engine and Content Engine can reconnect to the Process Engine and continue to provide users with access to workflows.

Process Engine procedure overview for UNIX

The procedure for making the Process Engine highly available differs by platform as well as configuration (cluster or farm).

Server farm environments

The following high-level steps are required to make the Process Engine highly available in a server farm environment on UNIX® or Windows platforms:

1. Install and configure the first Process Engine in the farm as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
2. Configure the hardware load balancer.

3. Install subsequent Process Engines on the farm servers as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide* with some farm specific exceptions.
4. Start Process Engine on the additional farmed servers.
5. Refresh the Process Engine node on the Process Task Manager.

For detailed instructions, see [“Install Process Engine in a server farm \(UNIX, Windows\)”](#) on page 79.

Cluster environments

NOTE A Process Engine cluster must run in an active-passive type of cluster configuration, with only one cluster node running Process Engine Services at any one time.

The following high-level steps are required to make the Process Engine highly available in a cluster environment on UNIX platforms:

1. Create partition resources in the cluster for the Process Engine.
2. Perform Process Engine pre-installation and configuration tasks on all nodes.
3. Install the Process Engine software on one node as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
4. Configure the Process Engine for the cluster.
5. Move the cluster group to second node (repeat this step for each additional cluster node).
6. Install Process engine software on second node (repeat this step for each additional cluster node) as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
7. Create a cluster resource to control the Process Engine (using one of the following two methods).
 - Install/configure the Veritas Certified FileNet Process Engine Agent.
 - Use a generic cluster package to create/configure a resource to control the Process Engine.

NOTE For an overview of command line shutdown and startup procedures, go to the *FileNet P8 Platform* help and navigate to FileNet P8 Administration > Enterprise-wide Administration > Shutdown and Startup.

The following services are available for monitoring:

- TM_daemon (required)
- vworbbroker (required)
- vwior (required)
- SEC_daemon (optional)
- NCH_daemon (optional)
- MKF_writer (optional)
- MKF_clean (optional)

8. Failover to each node to ensure that Process Engine resources startup properly on each node. For detailed instructions, see [“Install Process Engine in a Veritas cluster server \(UNIX\)” on page 81](#).

Process Engine procedure overview for Windows

The procedure for making the Process Engine highly available differs by platform as well as configuration (cluster or farm).

Server farm environments

See [“Server farm environments” on page 39](#). For detailed instructions, see [“Install Process Engine in a server farm \(UNIX, Windows\)” on page 79](#).

Cluster environments

The following high-level steps are required to make the Process Engine highly available in a cluster environment on Windows platforms:

1. Create partition resources in the cluster for the Process Engine
2. Install Process Engine on one node as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
3. Configure Process Engine for the cluster.
4. Move the cluster group to second node (repeat this step for each additional cluster nodes).
5. Install Process engine software on second node (repeat this step for each cluster node) as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
6. Add resources to the cluster to control Process Engine.
 - IMS Control Service
 - Process Engine Services Manager
 - Registry replication
7. Failover to each node to ensure that Process Engine resources startup properly on each node. For detailed instructions, see [“Install Process Engine in a Veritas cluster server \(UNIX\)” on page 81](#) or [“Install Process Engine in a Microsoft cluster server” on page 89](#).

Application Engine/Workplace XT

Application Engine/Workplace XT provides front-end access to FileNet P8 content using the corresponding Workplace/Workplace XT web application. By making Application Engine/Workplace XT highly available you ensure that users will always have access to a running instance of the application.

Application Engine/Workplace XT comprises a number of components. The following are the most important for a highly available FileNet P8 environment:

- The Workplace/Workplace XT web application
- The shared configuration directory
- Component Manager

Application Engine/Workplace XT procedure overview

The following high-level steps are required to make Application Engine/Workplace XT highly available:

1. Create a shared location for storing shared configuration information.
2. Install Application Engine/Workplace XT on all nodes.
3. Configure the application server cluster for Application Engine/Workplace XT.
4. Deploy Workplace/Workplace XT.
5. Perform any necessary post-deployment steps, and restart Application Engine/Workplace XT.

Rendition Engine

In a highly available RE environment a number of RE servers share the rendition administration and workload. Together they provide redundancy in the event of network or server failure. The Ligent software uses a Ligent Vista domain to track all render servers that have been configured and joined to the domain.

Each RE, controlled by the Ligent Vista domain, polls the RE database for submitted rendition jobs in a round-robin fashion. If a job is available, the Ligent Vista domain assigns it to one of the available RE servers. In this scenario the Ligent Vista domain, distributed across the available RE servers, acts as a “load balancer” and assigns jobs to active RE servers. If a server becomes unavailable, the available RE servers pick up the workload.

NOTES

- You can install FileNet Rendition Engine (RE) in highly available configurations using the standard Rendition Engine installer. No extra software is needed to set up a highly available Rendition Engine.
- FileNet P8 relies on the built in Ligent high availability for RE 4.5.
CAUTION This functionality is new for RE 4.0.1 and later, and the instructions in this chapter do not apply to earlier versions of RE.
- If an RE server becomes unavailable during rendering, the job ticket will fail, and a timeout will occur on the ordering CE server. After the timeout, CE resubmits the job once before signaling an error. The new job is picked up by another RE server.

NOTE After a Rendition Engine server becomes unavailable you should inspect both the RE Domain Manager Job tab and the CE EM Publishing Queue for failed jobs. Jobs are automatically retried one time if not returned within the publishing timeout (default: 600 s/10 min) set for the server. Manually retry any jobs in InError state in the CE EM Publishing Queue. For more information, go to the *FileNet P8 Platform* help and navigate to FileNet P8 Administration > Enterprise-wide Administration > Publishing > Publishing Queue > Retry Publishing requests.

Rendition Engine procedure overview

The following high-level steps are necessary to configure Rendition Engine for high availability:

1. Install Content Engine, Process Engine, and Application Engine.

2. Perform prerequisite tasks.
3. Install and configure an initial Rendition Engine.
4. Install and configure additional Rendition Engines.
5. (Optional) Verify Installation.

CFS-ICI

Content Federation Services for IBM Content Integrator (CFS-ICI) is an optional component of the IBM FileNet P8 Platform.

You can install CFS-ICI in highly available configurations using standard application server clustering software. The software is integrated with Content Engine, Process Engine, and Application Engine and integrates with other FileNet P8 components in clustered configurations.

NOTES

- The following procedure covers installation of the Federator Administrator application only.
- The following procedure covers the specific details required for installing CFS in an application server cluster rather than on a standalone server.
- Typically, application servers provide their own clustering/farming (the terminology differs by vendor) for achieving high availability with a web application. Application servers use server farms to make applications highly available. These are groups of identical servers running the same application server, which in turn runs the same web application. A load balancer will typically direct users to different instances throughout the group of servers. When one server goes down a user is automatically directed to an already running instance.

CFS-ICI procedure overview

The following high-level steps are required to make CFS-ICI highly available:

1. Install Content Engine and Process Engine.
2. Install a highly available Application Engine.
3. Install and configure an application server cluster for CFS.
4. Install and configure IICE in the cluster.

Follow the installation instructions provided in the *Content Integrator Enterprise Edition Support* Web site at <http://www.ibm.com/software/data/content-management/support/content-integrator/>.

5. Configure the CFS-ICI environment.
6. Install and Deploy CFS.
7. Perform Post Installation Steps.

eForms

You can install FileNet eForms for P8 on a server with FileNet P8 Application Engine or Workplace XT installed. Because FileNet P8 eForms is an add-on to the Application Engine or Workplace XT the same high availability solution for Application Engine or Workplace XT is used to provide availability for eForms.

eForms procedure overview

The following high-level steps are required to make eForms for P8 highly available:

1. Install eForms and appropriate patches, as documented in the P8 eForms installation and upgrade guide.
2. If Application Engine or Workplace XT is already deployed then redeploy the web application.
3. Perform post-deployment configuration steps as noted in the eForms documentation.

Process Analyzer

Process Analyzer (PA) provides dynamic, chart-based analyses of workflow activity, based on information collected from Process Engine (PE) or Process Simulator (PS).

You can install PA in a highly available environment using Microsoft Cluster Server (MSCS) technologies.

General information about PA installation

- PA only supports failover clustering.
- The PA installer needs to be run on each cluster node and supplied with the same configuration information.
- Configuration changes to the running PA instance should be made using Process Task Manager on the active node of the cluster.

PA Configuration Files

Configuration information is stored in two files located in `PA_install_path\FileNet\Process Analyzer Engine\jpa:`

- `vwtaskman.xml`
- `analyzer.properties`

NOTE PA does not support storing these files in a single location shared among cluster nodes, so if the PA configuration is changed in Process Task Manager, these changes should be propagated to all other nodes. In general, this should be done by copying the files to the other nodes, although it could also be done by running Process Task Manager on each node and making the same configuration changes.

Process Analyzer procedure overview

The following high-level steps are required to make Process Analyzer highly available:

1. Create a shared Temp directory.
2. Configure Microsoft Cluster Server (MSCS) PA Resource Group

3. Install PA Engine on all cluster nodes.
4. Configure PA Engine.
5. Install PA Client.

Other highly available components

In a highly available PA configuration, all other components of the system should be made highly available:

- SQL Server Database Engine
For details, see [“Database services” on page 34](#).
- SQL Server Analysis Services
For details, see [“Database services” on page 34](#).
- PE Database
For details, see [“Database services” on page 34](#).
NOTE While the PE database should be highly available, PA does not support Oracle RAC.
- Content Engine (may be farmed)
For details, see [“Install a highly available Content Engine” on page 70](#).
- Process Engine (may be farmed)
For details, see [“Install a highly available Process Engine” on page 79](#).

Records Manager

You can install FileNet P8 Records Manager version 4.5 in highly available configurations using standard application server clustering software. Records Manager is tightly integrated with Content Engine, Process Engine, and Application Engine and seamlessly integrates with other IBM FileNet P8 components in clustered configurations.

NOTES

- The following procedure covers installation of the Records Manager Web application only.
- The following procedure covers the specific details required for installing Records Manager in an application server cluster rather than on a standalone server.
- Typically, application servers provide their own clustering/farming (the terminology differs by vendor) for achieving high availability with a web application. Application servers use server farms to make applications highly available. These are groups of identical servers running the same application server, which in turn runs the same web application. A load balancer will typically direct users to different instances throughout the group of servers. When one server goes down a user is automatically directed to an already running instance.
- For instructions on installing RM components such as the RM Sweep Application or Crystal Reports, see the *IBM FileNet Records Manager Installation and Upgrade Guide*, version 4.5.

Procedure Overview for Records Manager

The following high-level steps are necessary to configure Records Manager for high availability:

1. Install Content Engine and Process Engine. Follow the instructions in the *IBM FileNet P8 Installation and Upgrade Guide*. For high availability installation of these components, see [“Install a highly available Content Engine” on page 70](#) and [“Install a highly available Process Engine” on page 79](#).
2. Install a highly available Application Engine/Workplace XT and note the location of the `bootstrap.properties` file (a highly available share is recommended). For information, see [“Install a highly available Application Engine/Workplace XT” on page 94](#).

NOTE Do not deploy Workplace. This step will be completed as part of deploying the Records Manager web application.

3. Install Records Manager and apply the relevant client file installers on each node in the cluster where Application Engine/Workplace XT is installed.
4. Configure the application server cluster for Records Manager.
5. Deploy Workplace/Workplace XT and then Records Manager.
6. Perform any necessary post-deployment steps.
7. Start Component Manager on all nodes in the server farm.

IBM FileNet Connector for SharePoint Web Parts, version 2.2

You can install IBM FileNet Connector for SharePoint Web Parts in a highly available SharePoint environment.

NOTES

- You can configure IBM FileNet Connector for SharePoint Web Parts, version 2.2, to connect to a highly available Content Engine or Process Engine, see [“Install IBM FileNet Connector for SharePoint Web Parts on each farm node.” on page 118](#).
- For information on how to install IBM FileNet Connector for SharePoint Document Libraries, version 2.2, in a highly available environment, see [“Install a highly available IBM FileNet Connector for SharePoint Document Libraries 2.2” on page 122](#).

IBM FileNet Connector for SharePoint Document Libraries, version 2.2

You can install IBM FileNet Connector for SharePoint Document Libraries 2.2 in a highly available SharePoint environment.

The setup program installs two sets of services:

- Connector Services installed on each member of the SharePoint farm. For information, see [“Install the Connector Services on the SharePoint Farm nodes.” on page 123](#).
- Connector Administration Services installed in a Microsoft active/passive cluster. For information, see [“Install the Connector Administration Services remotely on each cluster node.” on page 124](#).

Definition of roles

The tasks in this guide as well as the rows in the “Installation and upgrade worksheet” referenced in [“Use the installation worksheet” on page 48](#) are organized by administrative roles, listed below.

Your organization may have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.

Installation administrator (IA)

- Responsible for coordinating the information described in this worksheet. The information itself will require the input from the other roles. See [“Use the installation worksheet” on page 48](#).
- Runs IBM FileNet setup programs during initial setup.
- Runs the Configuration Manager tool during initial setup.
- Launches Enterprise Manager during initial setup.
- Runs IBM FileNet Upgrade programs during upgrades.

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

For tasks assigned to the IA, see [“Installation administrator tasks” on page 50](#).

IT administrator (ITA)

- Responsible for the networking and operating systems issues required by IBM FileNet P8.
- Responsible for those security configurations not assigned to the SA.
- Responsible for providing the information in the rows in the Installation and Upgrade Worksheet with a value of ITA in the Role column.

For tasks assigned to the ITA, see [“IT administrator tasks” on page 57](#).

Security administrator (SA)

- Responsible for configuring the directory servers required by IBM FileNet P8, including Content Engine, Application Engine, CFS Federation Administration application.
- Creates and maintains directory server user and group accounts
- Decides on configuration parameters required to connect
- Responsible for providing the information in the rows in the Installation and Upgrade Worksheet with a value of SA in the Role column.

For tasks assigned to the Security Administrator, see [“Security administrator tasks” on page 64](#).

Database administrator (DBA)

- Creates, configures, maintains database installations and database/tablespaces.

- Responsible for creating database accounts needed by FileNet P8.
- Might have responsibilities regarding the JDBC datasources
- Responsible for providing the information in the rows in the Installation and Upgrade Worksheet with a value of DBA in the Role column.

For tasks assigned to the Database Administrator, see [“Database administrator tasks” on page 65](#).

Application server administrator (ASA)

- Responsible for providing the application servers required by FileNet P8.
- Responsible for application server administrative accounts.
- Responsible for providing the information in the rows in the Installation and Upgrade Worksheet with a value of ASA in the Role column.

For tasks assigned to the Application Server Administrator, see [“Application server administrator tasks” on page 52](#)

Use the installation worksheet

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

For more information, see [“Appendix F - Installation and upgrade worksheet” on page 181](#).

Here are some things to know about the worksheet:

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

Perform required preparation tasks

ATTENTION In addition to the preparation tasks listed below, follow the plan and prepare instructions in the *IBM FileNet P8 Platform Plan and Prepare Your Environment for IBM FileNet P8* document for the components you are installing. For add-on components, follow the plan and prepare instructions in the corresponding installation guide. Note that some of the tasks listed below might replace the tasks described for a non highly available environment.

To prepare your environment for a highly available installation

Perform the required preparation tasks for the components you plan to install:

- [“Installation administrator tasks” on page 50](#)
- [“Application server administrator tasks” on page 52.](#)
- [“IT administrator tasks” on page 57](#)
- [“Security administrator tasks” on page 64](#)
- [“Database administrator tasks” on page 65](#)

In addition, if you are using the Installation and Upgrade Worksheet for your environment, all administrators should fill out the worksheet with the relevant information as described in [“Appendix F - Installation and upgrade worksheet” on page 181.](#)

Installation administrator tasks

The items below are the required tasks for the Installation administrator.

Content Engine

In addition to the tasks listed below, review and perform the required Content Engine tasks in the “Perform the required preparation tasks” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Content Engine installation location

You can install the Content Engine software on local storage; no Content Engine configuration files require shared storage.

NOTE File store data must still be shared. For more information, see [“File store high availability configuration” on page 78](#).

Content Search Engine

In addition to the tasks listed below, review and perform the required Content Search Engine tasks in the “Perform the required preparation tasks” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Process Engine

In addition to the tasks listed below, review and perform the required Process Engine tasks in the “Perform the required preparation tasks” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Application Engine/Workplace XT

In addition to the tasks listed below, review and perform the required Application Engine or Workplace XT tasks in the “Perform the required preparation tasks” section of the *Plan and Prepare your Environment for IBM FileNet P8* guide.

Rendition Engine

In addition to the tasks listed below, review and perform the required Rendition Engine tasks in the “Perform the required preparation tasks” section of the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

CFS-ICI

In addition to the tasks listed below, review and perform the required CFS tasks in the “Perform required preparation tasks” section of the *IBM FileNet Content Federation Services Installation and Upgrade Guide*.

eForms

In addition to the tasks listed below, review and perform the required eForms tasks in the “Perform the required preparation tasks” section of the *IBM FileNet eForms for P8 Installation and Upgrade Guide*.

Process Analyzer

In addition to the tasks listed below, review and perform the required Process Analyzer tasks in the “Perform the required preparation tasks” section of the *IBM FileNet Process Analyzer Installation and Upgrade Guide*.

IBM FileNet Connector for SharePoint Web Parts

Before you begin the installation, ensure that all prerequisites are completed. For details, see “Installation planning considerations” in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide*.

IBM FileNet Connector for SharePoint Document Libraries

Before you begin the installation, ensure that all prerequisites are completed. For details, see “Installation Planning Considerations” in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide*.

Application server administrator tasks

The items below are the required tasks for the Application server administrator.

JBoss clustering of Content Engine and Application Engine

In a JBoss cluster environment Application Engine and Content Engine must reside in separate JBoss clusters. Unless given a unique name, all JBoss nodes will be part of the cluster "Default."

For each Application Engine node, uniquely name the cluster something other than default.

Example:

Updated the JBoss startup file on both Application Engine cluster nodes from:

```
JBOSSSH=${JBOSSSH:-"$JBOSS_HOME/bin/run.sh -c $JBOSS_CONF -b $JBOSS_HOST"}
```

to:

```
JBOSSSH=${JBOSSSH:-"$JBOSS_HOME/bin/run.sh -c $JBOSS_CONF -b $JBOSS_HOST" } -  
Djboss.partition.name=JBossP8AE
```

Application Engine/Workplace XT

Configure the application server

Application Engine/Workplace XT is an enterprise application running on a Java application server and should be made highly available using either of the following configurations:

Application Server Clusters

- (WebSphere, WebLogic) Managed Application Server Clusters.

An administrative server manages a number of application server instances. Applications and configuration changes are implemented using an administrative server/interface and sent to each cluster node.

In this type of configuration the application server software provides the components to build and deploy the highly available enterprise application.

- (JBoss) Application Server Clusters

JBoss application server clusters do not use a separate administrative server, and configuration changes are implemented on each cluster node separately.

Farm of Independent Application Server Instances.

A number of separate server instances run independently behind a load-balancer or proxy device. No central administration server is used. You must install and deploy on each farm node.

Configure load-balancer or proxy server

- A load balancing or proxy device will typically direct users to different instances throughout the group of highly available servers. When one server goes down a user is automatically directed to an already running instance.

- The load-balancing device can be a hardware or software implementation of a proxy or load-balancer.

See the application server software's hardware and software support requirements to determine the supported load-balancing or proxy device supported for you configuration.

- When installing any components that request the URL of a deployed Workplace/Workplace XT instance in a highly available configuration, you must use the URL that directs users to your load-balancing or proxy device.

ATTENTION Specifying a single Application Engine/Workplace XT server will introduce a single point of failure in the environment.

- If you are using load-balancer or proxy device in your configuration you must use the load-balancer or proxy device to log on to Workplace/Workplace XT for the first time. Doing this ensures that the Base URL setting, which must be set to the load-balancer or proxy URL, is correctly set.

Example:

If the application is deployed to a farm of application servers, and a load-balancer is configured for this farm, then a URL could be:

Application Engine

`http://loadbalancer URL:loadbalancer port/Workplace`

Workplace XT:

`http://loadbalancer URL:loadbalancer port/WorkplaceXT`

Use this URL when completing the installation steps.

Plan for Workplace/Workplace XT shared settings

Workplace/Workplace XT settings can be stored in a single shared configuration directory. This directory can reside on a shared device as needed.

ATTENTION

- If you deploy your Workplace/Workplace XT web application as a Windows service on WebSphere, do not use mapped drives to reference shared configuration folders. Instead, use a UNC share. Note that UNC shares cannot be Windows administrative shares such as `\\host\c$`.
- If the share is mountable the mount must be online when Workplace/Workplace XT is started and accessed.
- If the shared configuration directory is not accessible an error page is returned when a user tries to access the Workplace/Workplace XT sign-in page.
- Aside from the shared configuration directory no other Application Engine/Workplace XT software has to reside on shared storage.

- Overwriting of configuration files on subsequent installations.
 - If you are performing an installation using a shared configuration directory be aware that all installations of Application Engine/Workplace XT will replace any files that exist in the shared configuration directory.
 - If an installation has been completed and settings such as the bootstrap information have already been set then any subsequent installations using the same shared configuration directory will result in the files being overwritten, losing any settings that have been made.

To preserve the settings in the shared configuration directory you should back up the files in the shared configuration directory and restore them when you have completed the installation.

NOTE This is not an issue for configurations where all Application Engine/Workplace XT installations are performed at the same time, but for cases where the installer must be run again (to add a node to the configuration for example).

Plan for Component Manager

If you will use Component Manager you must install and run it on each node in the application server configuration.

NOTE Including Component Manager in the installation also facilitates patching by providing a clean installation base on which to run product updates for the Application Engine/Workplace XT, which will in turn update the Component Manager software.

Content Engine

Install JDBC drivers for your selected database

Install the jdbc driver on all nodes in the CE cluster, add the jdbc driver to the classpath for all servers.

Install and configure the application server for high availability.

Content Engine is an enterprise application running on a Java application server and should be made highly available using either of the following Application Server Cluster configurations:

- (WebSphere, WebLogic) Managed Application Server Clusters.

An administrative server manages a number of application server instances. Applications and configuration changes are implemented using an administrative server/interface and sent to each cluster node.

In this type of configuration the application server software provides the components to build and deploy the highly available enterprise application.

- (JBoss) Application Server Clusters.

JBoss application server clusters do not use a separate administrative server, and configuration changes are implemented on each cluster node separately.

Configure the application server for Content Engine.

To configure the Application Server on each node follow the instructions in the *IBM FileNet Plan and Prepare Your Environment for IBM FileNet P8* with the following modifications:

- WebSphere:

Perform the task “Configure WebSphere for Content Engine.”

NOTE Set the JDBC parameters for the nodes using the administrative console, not on the individual servers.

- WebLogic:

Perform the task “Configure WebLogic for Content Engine.”

- JBoss:

Perform the task "Configure JBoss for Content Engine," except use JBOSS_DIST/all as the base configuration file set instead of JBOSS_DIST/server.

(WebSphere only) Configure Light weight Third Party Authentication (LTPA)

To configure LTPA on the Content Engine cluster, in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, in the task “Configure Application Engine (WebSphere)”, follow the instructions in the “Configure Light weight Third Party Authentication (LTPA)” step for the version of WebSphere you are running, with the following modifications:

- Set timeout to 60,000.
- After saving your configuration settings, synchronize the changes across all nodes.
- If your environment uses Workplace, follow the second part of the “Configure Light weight Third Party Authentication (LTPA)” step to configure the Application Engine server for LTPA.
- If your environment uses Workplace XT, follow the second part of the “Configure Light weight Third Party Authentication (LTPA)” step in the *IBM FileNet Workplace XT Installation Guide* task “Configure IBM FileNet Workplace XT(<Application server>)” to configure the Workplace XT server for LTPA.

Configure load-balancer or proxy server

Load-balancers and proxy-servers are typically used to direct users to different instances throughout a highly available application server configuration. When one server goes down the system automatically redirects the user to another running server instance. For more information, see [“Appendix C – Using load balancers with IBM FileNet P8” on page 171](#).

The following are requirements for configurations using a load-balancer or proxy-server:

- Session affinity is required for Application Engine and Workplace XT only.
- Virtual server name.

If you use a load-balancer or proxy server in this configuration you must use the virtual name when performing installation steps that require a server name for a Content Engine server, with the exception of FileNet Enterprise Manager and other administrative applications. For more information, see [“Appendix B – Connecting to a highly available Content Engine” on page 168](#).

CFS-ICI

Pre-installation tasks

The following procedure assumes that the application server already has a cluster instance configured and is ready for applications to be deployed.

1. Install and configure an application server cluster for CFS-ICI.
2. Verify that the cluster nodes and load balancers are started.

IT administrator tasks

The items below are the required tasks for the Information Technology (IT) administrator.

Content Search Engine

(UNIX only) Update the k2_os_user profile

If you are installing CSE on UNIX, add the following lines to k2_os_user's profile:

```
CATALINA_PID=/verity_install_path/data/host/log/catalina.pid
export CATALINA_PID
```

Create required shared directories

CSE requires three separate locations:

1. A location where the CSE executable files will be stored (e.g., /opt/verity).
This location must be shared by each cluster node and may either be on shared local storage (UNIX and Windows) or mounted from a remote highly-available storage location (UNIX only).
2. A location where temporary files will be placed by CE for indexing by CSE (e.g., /opt/verity/tmp).
This location must be shared by each CSE cluster node and each CE node configured to use CSE. This location may either be on shared storage local to a CE or CSE cluster or mounted by both CE and CSE from a remote highly-available storage location.
3. A location where the CSE collections will be stored.
This location must be on shared storage local to CSE and not mounted from a remote location. It must also be made available as a clustered file share resource that can be mounted remotely by CE.

Create highly available file shares to store the CSE binary files and configuration files, and mount these directories from each CSE cluster node at the following locations:

Table 5: CSE shared directories

File Share	Mount point UNIX	Mount point Windows
Binary Share	/opt/verity	Drive:\Program Files\FileNet\ContentEngine\verity
Temp DirectoryShare	/opt/verity/tmp	Drive:\Program Files\FileNet\ContentEngine\verity\tmp
Collections Share	/opt/verity/collections	Drive:\Program Files\FileNet\ContentEngine\verity\collections

Application Engine/Workplace XT

Create a local or shared directory for the shared configuration files

For more information on the use of the shared configuration directory and how to manage it see [“Appendix E – Managing Workplace/Workplace XT settings in a highly available environment”](#) on page 178.

CAUTION At a minimum the user running the install and the Application Server/Workplace XT processes needs write access to this directory.

NOTE Do not use one of the cluster servers for the file location as this creates a single point of failure. The bootstrap.properties file could, theoretically, be placed on a local Windows share or local NFS export directory from any of the systems in the Application Engine/Workplace XT cluster (i.e. shared out from the default file location from the first installation). However, if the local system holding the file would go down, other Application Engine/Workplace XT instances will be unable to find the bootstrap.properties file and will return error messages.

Process Engine

Configure hardware load balancer

Configure your hardware load balancer to forward and monitor TCP traffic on ports 32776 (IOR) and 32777 (broker) only.

For information about required and default ports for your environment, see the IBM FileNet P8 Ports appendix in the *Plan and Prepare Your Environment for IBM FileNet P8* document.

Configure the Process Engine nodes hosts file

Server Farm

For each Process Engine server, configure the hosts file so that on each node in the farm it:

- Includes the NCH name for every other node in the farm.
- Identifies the loadbalancer virtual as itself.

The hosts file entries should look like this:

	Hosts file entry
<i>PE_farm1</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name loadbalancer_virtual_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name</i>
<i>PE_farm2</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name loadbalancer_virtual_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name</i>
<i>PE_farm3</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name loadbalancer_virtual_name</i>

Example:

loadbalancer_virtual_name (1.2.3.10) is the load balancer virtual name, and *PE_farm1* (1.2.3.4), *PE_farm2* (1.2.3.5), and *PE_farm3* (1.2.3.6) are three PE servers.

The format of *PE_farm1_nch_name* is:

```
PEfarm1-filenet-nch-server
```

The NCH name in the hosts file is the NCH name with -filenet-nch-server appended, and any dashes or underscores removed.

CAUTION (UNIX) Verify that the nsswitch.conf file (AIX netsvc.conf and/or irs.conf) is configured to resolve names via the local hosts file first.

Server Cluster

Configure the hosts file on each cluster node server to follow the following format:

```
PEvirtual_ip PEvirtual_dns_name PEclustered_nch_name local_node_name
```

Server	Hosts file entry
<i>PE_node</i>	1.2.3.4 <i>virtual_dns_server clustered_nch_server local_node_server</i>

Example

```
1.2.3.4 virtual_dns_server.ibm.com clustered_nch_server local_node_server
```

The format of *PEclustered_nch_server* is:

```
PEvirtual_dns_name-filenet-nch-server
```

The NCH name in the hosts file is the NCH name with -filenet-nch-server appended, and any dashes or underscores removed.

Configure your operating system for Process Engine Cluster (UNIX)

Perform the following required pre-installation configuration and verification as documented in the platform specific Process Engine installation tasks found in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- AIX
 - “To set time zone parameters” (perform on all nodes).
 - “To modify /etc/rc.dt and /etc/tunables/nextboot for AIX 5.2 and 5.3” (perform on all nodes).
 - “To create or modify internal Process Engine users and groups” (perform on all nodes).
 - **CAUTION** Make sure you use identical user and group IDs on all nodes.
 - “To create partitions:”
 - FileNet recommends the /fnsd partition/directory be placed on a local (non-clustered) storage. The Process Engine binaries are placed in the /fnsd directory and if they are taken offline there will be no binaries available to shutdown, clean, or troubleshoot a failed node.
 - This is performed from one cluster node, wherever the shared storage resources are active (for example a shared disk resource), as we will use them to create these partitions.
 - After the partitions are created and security is set create Veritas resources for them to bring them under cluster control.
 - “To verify the database connection” (perform on all nodes).
 - “To turn off Oracle password complexity verification” (if necessary this is applicable to the Oracle server).
- HPUX
 - “To perform symbolic links for X11 libraries” (perform on all nodes).
 - “To check and optionally modify the time zone settings” (perform on all nodes).
 - “To configure kernel parameters” (perform on all nodes).
 - “To create or modify internal Process Engine users and groups” (perform on all nodes).
 - **CAUTION** Make sure you use identical user and group IDs on all nodes.
 - “To create partitions:”
 - FileNet recommends the /fnsd partition/directory be placed on a local (non-clustered) storage. The Process Engine binaries are placed in the /fnsd directory and if they are taken offline there will be no binaries available to shutdown, clean, or troubleshoot a failed node.

- This is performed from one cluster node, wherever the shared storage resources are active (for example a shared disk resource), as we will use them to create these partitions.
- After the partitions are created and security is set create Veritas resources for them to bring them under cluster control.
- “To verify the database connection” (perform on all nodes).
- “To turn off Oracle password complexity verification” (if necessary this is applicable to the Oracle server).
- Solaris
 - “To increase the operating system kernel limits” (perform on all nodes).
 - “To create or modify internal Process Engine users and groups” (perform on all nodes).
 - CAUTION** Make sure you use identical user and group IDs on all nodes.
 - “To verify national language character set and time settings” (perform on all nodes).
 - “To create partitions:”
 - FileNet recommends the /fnsw partition/directory be placed on a local (non-clustered) storage. The Process Engine binaries are placed in the /fnsw directory and if they are taken offline there will be no binaries available to shutdown, clean, or troubleshoot a failed node.
 - This is performed from one cluster node, wherever the shared storage resources are active (for example a shared disk resource), as we will use them to create these partitions.
 - After the partitions are created and security is set create Veritas resources for them to bring them under cluster control.
 - “To verify the database connection” (perform on all nodes).
 - “To turn off Oracle password complexity verification” (if necessary this is applicable to the Oracle server).

Install and configure Veritas Cluster Server (UNIX)

1. Install the clustering software prior to installing or configuring any FileNet services for high availability. This includes Veritas Cluster Server, Veritas Volume Manager, data replication software, and backup software.
2. Create a cluster group with the following resources (at a minimum):
 - Clustered storage resources - Any shared storage resource that will be used to host the partitions required for the Process Engine (i.e. the /fnsw/local, fn_sec_db0, and fn_sec_r10 partitions).
 - Clustered IP resource.
3. Verify that the cluster group can failover to all nodes in the cluster, and that the clustered storage can be accessed from all nodes.

Install and configure Veritas Cluster Server (Windows)

1. Install the clustering software prior to installing or configuring any FileNet services for high availability. This includes Veritas Cluster Server, Veritas Volume Manager, data replication software, and backup software.
2. Create a clustered storage drive to store configuration data.

The Process Engine stores configuration data under the `fnsw_loc` directory. Clustering this location ensures that each node uses consistent information and provides the same data to all users.

3. Verify that a cluster group exists with the following resources (at a minimum):
 - Clustered storage resources (can include Veritas Volume Group (VMDg) and MountV, or Mount for a basic disk).
 - A clustered IP resource
 - A Lanman name resource
4. Verify that the cluster group can failover to all nodes in the cluster, and that the clustered storage can be accessed from all nodes.

Install and configure Microsoft Cluster Server

1. Verify that you are using Microsoft Windows Cluster Service-approved hardware. All hardware used in the Cluster configuration must be listed in the Windows Catalogs. For more information, go to the Microsoft Web site at <http://www.microsoft.com> and search for "Windows Catalog and HCL."
2. Verify that your cluster nodes are members of a Windows domain running in Windows Server 2003 mode.
3. Install the clustering software prior to installing or configuring any FileNet P8 services for high availability.
4. Create a clustered storage drive where configuration data will reside.

CAUTION The Process Engine stores configuration data under the `fnsw_loc` directory. Clustering this location ensures that each node uses consistent information and provides the same data to all users.

NOTE FileNet recommends that you keep the `fnsw` directory on a local (non-clustered) storage.

5. Verify that a group of cluster resources already exists with the following resources (at a minimum):
 - Clustered disk resource
 - Clustered IP resource
 - Network Name resource
6. Verify that the cluster group can fail over to all nodes in the cluster, and that the clustered storage can be accessed from all nodes.

Process Analyzer

Create a shared Temp directory

On the PA Engine server, create a shared Temp directory and grant full access rights on the folder to the pa_administrator. Full access rights should be set on the Security and Sharing tabs.

In a highly available environment the UNC path to the temp directory configured in Process Task Manager should use the PA cluster's network name so that the path can be configured identically on all nodes.

The shared temp directory will need to be highly available and part of the group of resources configured for PA. For performance reasons, it should be located on the same cluster as PA.

Configure Microsoft Cluster Server (MSCS) PA Resource Group

1. Add resources to the cluster.

The contents of the PA Resource Group should be:

- IP Address
- Network name
- PA Temp Share
- Process Analyzer Services Manager (VWServicesPA)

To add resources to the cluster

- a. Add a resource for the PA Temp share.
 - i. Set the Resource Type: File Share
 - ii. Check the **Use Network Name for computer name** box.
 - iii. Make the resource dependent on the following resources:
 - Clustered network name
 - Clustered disk resources
- b. Add a resource for the VWServicesPA.
 - i. Set the Resource Type: Generic Service
 - ii. Set the Service name to: VWservicesPA
 - iii. Check the **Use Network Name for computer name** box.
 - iv. Make the resource dependent on the following resource:
 - PA Temp Share
 - Clustered network name
 - Cluster IP address
- c. Bring the cluster group online.

Security administrator tasks

The items below are the required tasks for the Security Administrator.

Rendition Engine

Assigning the *FNRE_Admin* user

For a highly available RE environment the required *FNRE_Admin* user can be either:

- A domain user available to all Rendition Engine servers.
- An identically configured user on each Rendition Engine server.

CAUTION Use the same named account with the same password on all Rendition Engine servers.

Records Manager

(UNIX) Verify write permission on the */bin* directory

Verify that the */bin* directory is write-enabled before running the RM installer. The directory need only be write-enabled during installation, and can be set back to whatever security permission you choose after the installation completes.

Database administrator tasks

The items below are the required tasks for the database administrator.

Content Engine

Install and configure the database.

This step includes installing and configuring your database, including setting it up to be highly available if required. Follow the instructions in the following *IBM FileNet Plan and Prepare Your Environment for IBM FileNet P8* tasks:

- “PrepareMicrosoft SQL Server”
- “Prepare Oracle Server”
- “Prepare DB2 Server for z/OS” or “Prepare DB2 Server for Linux®, UNIX and Windows”
- “Configure Microsoft SQL Server Client for Process Engine”
- “Configure Oracle Client for Process Engine”
- “Configure DB2 Client”

For more Content Engine requirements, see the “Perform the required preparation tasks” in the *IBM FileNet Plan and Prepare Your Environment for IBM FileNet P8*.

Process Engine

Oracle RAC configuration

- Oracle 9i

When using Oracle 9i RAC with FileNet P8 you must configure the database instance for Broadcast on Commit (BOC) by setting the MAX_COMMIT_PROPAGATION_DELAY to a value of 0. You must do this on all databases used by FileNet P8 Content Engine.

- Oracle 10g and 11g

When using Oracle 10g or 11g RAC with FileNet P8 you must perform the following actions for any databases used with FileNet P8 Content Engine.

```
grant select on pending_trans$ to public;  
grant select on dba_2pc_pending to public;  
grant select on dba_pending_transactions to public;  
  
grant execute on dbms_system to user;
```

CAUTION Execute the last statement for each object store user in the database that Content Engine is using to access the Oracle database.

Process Analyzer

PA and Oracle RAC

While the PE database should be highly available, note that PA does not support Oracle RAC.

Rendition Engine

Configure remote databases

Perform all required Prerequisite tasks as outlined in the *FileNet Rendition Engine Installation and Upgrade Guide*, with the following modifications:

SQL Server

Perform task “Install Microsoft SQL Server for Rendition Engine.”

NOTE Install the database as a remote database, available to all Rendition Engine servers.

Oracle

Perform task “Install Oracle Server for Rendition Engine.”

NOTE Install the database as a remote database, available to all Rendition Engine servers.

Install and configure for high availability

This section provides the information needed to install a highly available FileNet P8 4.5 environment.

- [“Install a highly available Content Engine” on page 70.](#)
- [“Install a highly available Process Engine” on page 79.](#)
- [“Install a highly available Application Engine/Workplace XT” on page 94.](#)
- [“Install a highly available Rendition Engine” on page 99.](#)
- [“Install a highly available Content Search Engine” on page 101.](#)
- [“Install a highly available CFS-ICI” on page 104.](#)
- [“Install a highly available Process Analyzer” on page 110.](#)
- [“Install a highly available Records Manager” on page 112.](#)
- [“Install a highly available IBM FileNet Connector for SharePoint Web Parts 2.2” on page 117.](#)
- [“Install a highly available IBM FileNet Connector for SharePoint Document Libraries 2.2” on page 122.](#)

Overview

Depending on your configuration, you might have to install at least a portion of the additional fix pack and interim fix software while you work your way through the standard IBM FileNet P8 platform installation.

For more information about available service packs and interim fixes, see the *FileNet P8 Release Notes* and the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

To install FileNet P8 components in a highly available environment

To install and configure IBM FileNet P8 platform in an HA environment, follow the procedures as outlined in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following HA specific modifications:

- Database Services.

This step typically only involves third-party software and does not require you to install or configure any FileNet software.

You will typically create the database server instance and necessary tablespaces, including tablespaces for Global Configuration Data (GCD) and object stores, before you install any FileNet software making it easier to troubleshoot problems that are not FileNet software related.

- Content Engine Server Farm.

This step involves the installation and configuration of the Content Engine component in a highly available application server environment.

As part of this step you will create the FileNet P8 Domain.

For more information, see [“Install a highly available Content Engine” on page 70](#).

- Process Engine Cluster or Farm.

This step involves installing and configuring a highly available Process Engine.

Prior to setting security information in the Process Task Manager as part of the post installation configuration steps, the following must be true:

- Content Engine must be online.
- The GCD domain must be created.
- Security must have been set.

Installing Process Engine at this stage makes it possible to verify Process Engine functionality after Application Engine comes online.

For more information, see [“Install a highly available Process Engine” on page 79](#).

- Application Engine/Workplace XT Server Farm.

This step involves installing and configuring a highly available Application Engine/Workplace XT.

As Application Engine/Workplace XT requires all other FileNet P8 components online in order for you to log in, browse content, and perform any activities with the Process Engine, it is most often the component installed and/or configured last.

For more information, see [“Install a highly available Application Engine/Workplace XT”](#) on page 94.

Install a highly available Content Engine

This section covers the tasks and requirements for making Content Engine highly available in a FileNet P8 environment.

Installation and deployment tasks

To install Content Engine

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
2. (JBoss) Verify that the Content Engine and Application Engine clusters have different names. For information, see [“JBoss clustering of Content Engine and Application Engine” on page 52](#).
3. Install Content Engine.

In this configuration, install the Content Engine to the Administrative server in order to configure the security and JDBC resources needed by the Content Engine.

When installing Content Engine in an Application Server Cluster, an administrative server manages a number of application server instances. Applications and configuration changes are done through the administrative server/interface.

Follow the instructions in “Install Content Engine” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

4. Verify that all instances in your highly available application server configuration are running.
5. Configure Content Engine instances.

Follow the instructions in the “Configure Content Engine instances” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

- Run the CM on your administrative node.
- Configure JDBC Data Sources tab:
 - Database Name - If you are using a highly available database for the GCD make sure to use the virtual server name when prompted for the GCD database information.
 - Oracle RAC
 - Select Oracle JDBC Driver (RAC support)
- Configure LDAP tab
 - Directory service server host name - Use the fully qualified domain name or DNS alias for the LDAP servers.
- Deploy Application tab
 - Deployment type - Cluster.

- Application server cluster name - (Application Server Cluster configurations only) For Application server cluster name, enter the name of the HA cluster you have configured.

This configures the GCD data sources and security settings at the administrative level so that all cluster nodes have access to this vital configuration information.

6. Install the latest Process Engine Client files on Content Engine servers.

Follow the instructions in the “Install the latest Process Engine Client files on Content Engine servers” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

Install the Process Engine Client files on all servers on which CE is installed.

7. (Application server cluster only) Deploy Content Engine instances to the cluster nodes.

Follow the instructions in the “Deploy Content Engine instances” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

Verify that Deployment type is set to CLUSTER.

NOTE JBoss clusters do not have an administrative server. When installing Content Engine choose a single JBoss server on which to install CE, then copy the necessary files to the rest of the nodes in the cluster. To deploy a Content Engine instance on each additional JBoss server in your cluster, follow the instructions in the “Deploy Content Engine to additional JBoss servers in a cluster” subtopic of the “Deploy Content Engine instances” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

To configure Content Engine post installation

1. Install and launch FileNet Enterprise Manager (EM).

Follow the instructions in “*Install FileNet Enterprise Manager*” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task.

2. Configure the FileNet P8 domain.

- a. Create a FileNet P8 domain.

Follow the instructions in “To create a FileNet P8 domain” in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Establish the FileNet P8 domain and Global Configuration Data (GCD).”

NOTE Do this once per HA environment only.

CAUTION Using a load-balancer for the Enterprise Manager connection is not supported. For more information, see “[Appendix B – Connecting to a highly available Content Engine](#)” on page 168.

- b. Configure directory service authentication.

Follow the instructions in “To configure directory service authentication” in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Establish the FileNet P8 domain and Global Configuration Data (GCD).”

- c. Configure permissions for a FileNet P8 domain.

Follow the instructions in “To configure permissions for a FileNet P8 domain” in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task *Install and Deploy Content Engine* task “Establish the FileNet P8 domain and Global Configuration Data (GCD).”

NOTE Once this is complete every other server instance that starts up in a farm or application server cluster will be able to see the P8 domain.

3. Perform post-deployment configuration steps.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* tasks.

4. (If you have made changes that requires CE to restart) Restart the Content Engine application in the cluster after performing post-deployment configuration.

Creating object stores in FileNet P8 4.5

To create object stores follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* tasks “Create the data sources for an object store,” “Create the initial object store,” and “Verify the Content Engine installation”.

For highly available environments some or all of the following information might apply to your environment.

When creating your profile using the Create New Configuration Profile wizard make sure you do the following:

- Create a profile for a cluster.
- JDBC connection URL:
 - If you are using a highly available database cluster, use the virtual name of the cluster when entering the server/host name.
 - (WebSphere and WebLogic) Oracle RAC JDBC connection URL.

If you are using Oracle RAC for your backend database you must use a specific formation for the JDBC connection URL property in your data source as follows:

WebSphere

```
jdbc:oracle:@(DESCRIPTION=(load_balance=yes)(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=rac node 1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=rac node 2)(PORT=1521)))(CONNECT_DATA=(SERVICE_NAME=DB service name)(failover_mode=(type=select)(method=basic)(retries=retry number)(delay=delay value))))
```

Example:

```
jdbc:oracle:@(DESCRIPTION=(load_balance=yes)(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=testrac1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=testrac2)(PORT=1521)))(CONNECT_DATA=(SERVICE_NAME=testracdb)(failover_mode=(type=select)(method=basic)(retries=32)(delay=4))))
```

WebLogic

Follow the instructions available on the BEA Product Documentation web site,

<http://edocs.bea.com>.

Search for “Configuring JDBC Multi Data sources.”

JBoss

```
jdbc:oracle:@(DESCRIPTION=(load_balance=yes)(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=rac node 1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=rac node 2)(PORT=1521)))(CONNECT_DATA=(SERVICE_NAME=DB service name)(failover_mode=(type=select)(method=basic)(retries=retry number)(delay=delay value))))
```

Example:

```
jdbc:oracle:@(DESCRIPTION=(load_balance=yes)(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=testrac1)(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=testrac2)(PORT=1521)))(CONNECT_DATA=(SERVICE_NAME=testracdb)(failover_mode=(type=select)(method=basic)(retries=32)(delay=4))))
```

Enable application server connections after database failover

Depending on your application server, do one of the following:

WebSphere

- [“To enable WebSphere connection to the GCD and object store databases after failover” on page 74.](#)

WebLogic

- [“To enable WebLogic connection to the GCD and object store databases after failover” on page 75.](#)

JBoss

- [“To enable JBoss connection to the GCD and object store databases after failover” on page 75.](#)

To enable WebSphere connection to the GCD and object store databases after failover

To enable WebSphere to connect to the database after database failover, do the following steps for each of the two data sources (XA and non-XA) you created for your object store (replacing *Data Source Name* with the name of your XA and non-XA data sources):

WebSphere 6.1.x

1. Log on to WebSphere Administrative Console.
2. Navigate to JDBC Providers > JDBC Providers for GCD Datasource > Data sources > *Data Source Name* > WebSphere Application Server data source properties.
3. Set the **Pretest connection properties** check box in WebSphere to enable and accept the default parameter values for pretest connection retry interval and pretest SQL string.
4. (RAC only) Navigate to JDBC Providers for GCD Datasource > Data sources > *Data Source Name* > custom properties, and remove the `serverName` and `portNumber` properties.
5. Save your changes.

WebSphere 7.x

1. Log on to WebSphere Administrative Console.
2. Navigate to Resources > JDBC > Data sources > *Data Source Name* > WebSphere Application Server data source properties.
3. Set the "Validate New Connections" check box. Under this, set number of retries to 100 and retry interval to 3.
4. Set the "Validate existing pooled connections" box and set the retry interval to 0 seconds.
5. Under validation options, enter an appropriate query for your database type.

Example, For DB2, enter

```
SELECT 1 FROM sysibm.sysdummy1
```

6. (RAC only) Navigate to Resources > JDBC > Data sources > Data Source Name > Custom Properties, and remove the `serverName` and `portNumber` properties.
7. Save your changes.

To enable WebLogic connection to the GCD and object store databases after failover

To enable WebLogic to connect to the GCD database after failover, do the following steps for each of the two data sources (XA and non-XA) created during installation of Content Engine (replacing *GCD Data Source Name* with the name of your XA and non-XA data sources):

1. Log on to WebLogic Administration Console.
2. In the Domain Structure tree, expand Services > JDBC, and then select Data Sources.
3. On the Summary of Data Sources page, click *GCD Data Source Name* or Data Source Name.
4. In the Configuration tab, click the **Connection Pool** tab, click the **Advanced** menu, then select the **Test connections on reserve**. By default the option is collapsed beneath a drop down box.
5. Click **Advanced** and select the Test Connections On Reserve check box.
6. Save and activate your changes.

To enable JBoss connection to the GCD and object store databases after failover

Edit the `JBoss_HOME/server/all/deploy/*-ds.xml` files you created for the data sources and uncomment the following lines:

Oracle

```
<valid-connection-checker-class-name>org.jboss.resource.adapter.jdbc.vendor.OracleValidConnectionChecker</valid-connection-checker-class-name>.
```

DB2 and SQL Server

```
<check-valid-connection-sql>SQL statement</check-valid-connection-sql>
```

Where *SQL statement* is specific for DB2 or MS SQL Server. For information, see JBoss documentation for "Connectors on JBoss."

Deployment verification

To verify successful deployment

1. Verify that the Content Engine application is running on the cluster/farm.
2. Verify Content Engine access on each node.

For each server node in your configuration, in a browser, go to:

`http://CE_server:port#/FileNet/Engine`

to load the CE “Startup Context” page and verify that CE is running on the server.

Where:

CE_server is ONE of your servers in your farm.

port# is the port used to access CE on server running CE.

The “Startup Context” page displays certain data regarding CE, such as CE instance name you are connected to, server uptime, and product build number. View the startup context.name of the load-balancer or proxy device to verify the functionality of the load-balanced environment, or the name of the physical server to verify a single server.

3. Verify Content Engine access through the load balancer.

In a browser, go to:

`http://CE_load_balancer:port#/FileNet/Engine`

to load the CE “Startup Context” page and verify that CE is accessed correctly using the load balancer.

Where:

CE_load_balancer is name of the load-balancer or proxy device to verify the functionality of the loadbalanced environment.

port# is the port used to access CE on the load balancer or server running CE.

NOTES

- Depending on your configuration you can connect to the load balanced CE installation, and after confirming that CE is up and running refresh your browser to have the load balancer cycle through all managed nodes for verification.
- For JBoss, this cyclic verification will not work as JBoss only reports the name of the cluster and not the individual servers.

4. Create an Object Store and Verify the Content Engine Installation.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Create Object Stores” and “Verify the Content Engine Installation.”

5. Verify failover.
 - a. In a browser, go to:
`http://CE_server.port#/FileNet/Engine`
 - b. Fail one of the Content Engine nodes.
 - c. Verify that you can reload the Startup Context page, and that you can access your object stores using Enterprise Manager.
 - d. Repeat [Step a](#) and [Step b](#) for all nodes in your configuration.
 - e. Restart all nodes you failed during this test.

File store high availability configuration

Unlike previous versions FileNet P8 Content Engine version 4.0 and higher doesn't use a Windows based file store service to manage file store content. Instead, file store content is managed from the Content Engine application.

To make file store data highly available

To ensure continuity of access to content in file stores, the following conditions must be met:

- All File Stores
 - File systems used for file stores must be shared or mounted through the NFS or CIFS protocol.
 - Each Content Engine instance in a farm must have access to the same shares or mount points.
- Standard File Stores
 - Standard file stores must be installed in file systems made highly available.
- Fixed File Stores
 - The file system front-end must be made highly available as indicated above, and the remote device must be made highly available using its built-in capability (refer to vendor documentation for details).

NOTES

- A fixed file store comprises a file system-based front-end, similar to a standard file store, and a remote storage system such as Centera or NetApp/IBM N-Series SnapLock.
- The method and technology used to make file store data highly available can vary greatly, from file servers built on general-purpose hardware and cluster technology, such as Veritas Cluster Server or IBM HACMP, to specialized devices with built-in high availability such as NAS- and SAN-based offerings from EMC, NetApp or IBM.
- The implementation and use of these technologies vary, but they can all be used to provide highly available access to data so long as they meet the criteria discussed in this section.

Install a highly available Process Engine

This section covers the tasks and requirements for making Process Engine highly available in a FileNet P8 environment.

Install Process Engine in a server farm (UNIX, Windows)

This section lists the steps required to install, configure, and use FileNet Process Engine in a farm environment.

To install and configure Process Engine

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
2. Install and configure the first Process Engine in the farm as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
NOTE Perform all required tasks including installing software updates, installing client files, and so on.
3. Configure the hardware load balancer.
 - a. Verify that your hardware load balancer is set up to allow ALL PROTOCOLS traffic, not just TCP.
 - b. Verify that the load balancer virtual name is in the host file of the local PE server and that it resolves to that PE server IP address.

(UNIX) To enable the vworbbroker process to start, the name resolution must come from the host file first.
 - c. Configure the Process Task Manager for PE farm.
NOTE This configuration is only required for the first PE. These are steps that should be performed in connection with the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure Process Task Manager.”
 - i. Using Process Task Manager (vwtaskman), on the Servers node, add the hardware load balancer virtual name registered in DNS.
 - ii. On the Process Engine node, select the **Advanced** tab and add the following property:
`vworbbroker.endPoint = giop:tcp:loadbalancer_virtual_name:32777`
 - d. Configure the Isolated Region connection pointers on the CE. All connection pointers for the farm should use the load balancer virtual name as the DNS name.
4. Install subsequent Process Engines on the farm servers as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide* with the following farm-specific exceptions:

- Use a unique NCH name for each server in the farm. For information, see [“Configure the Process Engine nodes hosts file” on page 58.](#)

CAUTION As farm nodes act as independent servers they must each have a unique NCH name.

- When running SQL scripts, choose the option “I have already run the pre-install scripts manually.”
- Configure all additional Process Engines to use the same database as the initial Process Engine.

CAUTION Make sure the settings listed in the “Complete *database* database screens” step are identical to what you initially set up during your FileNet PE 4.5 installation.

- You do not need to perform the task “Configure Process Task Manager” as this was done when you configured the first Process Engine. See [“Configure the Process Task Manager for PE farm.” on page 79](#) above.

5. Start Process Engine on the additional farmed servers.
6. Refresh the Process Engine node on the Process Task Manager.

You should now see all the farmed PE servers under the servers node.

Install Process Engine in a Veritas cluster server (UNIX)

This section lists the steps required to install, configure, and use the FileNet custom agents on Veritas Cluster Server; all steps are applicable for Solaris, AIX, and HP-UX.

Symantec uses “Service Groups” of resources to provide highly available services to users. The following instructions specify how to add resources to a pre-existing service group to provide a highly available Process Engine Service.

Symantec provides a framework to create custom agents that can control and monitor particular resources. FileNet has created custom agents using the Veritas framework to allow users to create, monitor, and control a Process Engine resource. The agents have been certified by Symantec for use with their Veritas Cluster Server product on Solaris, AIX, and HP-UX platforms.

CAUTION In a Veritas Cluster Environment the user accounts, by design, must be local and not Network Information Services (NIS and NIS+) users if they are to be used by an agent. This requirement includes the account specified as the Oracle Owner in the Veritas Oracle agent, as well as the fns user account that is used in the FileNet Process Engine agent. For more information, refer to the Veritas Cluster Server and Veritas Agent documentation.

NOTE From time to time it can become necessary to recompile the FileNet Process Engine Veritas Agent to accommodate a change in the operating system (if the agent is to be used on an operating system on which the agent was not compiled) or a change in the startup and shutdown behavior of the agent (for example to increase the level of monitoring). To recompile the FileNet Process Engine Veritas Agent to produce a new agent binary if needed, follow the instructions in [“Appendix D – Recompiling the FileNet Process Engine Veritas agent” on page 176](#).

Installation tasks

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
2. Install Process Engine in the cluster.
 - a. Install the Process Engine on the first node as documented in the platform specific Process Engine installation tasks “Install Process Engine interactively” or “Install Process Engine silently” found in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - b. Once the installation completes, perform the following system specific post-configuration steps on each system:
 - “(Optional: Oracle databases only) To remove fns from the Oracle database administrators group (UNIX).”
 - “To restore any custom modifications for root and fns users (UNIX).”
 - “To verify the /etc/services file settings (UNIX).”
 - “To clean up before starting Process Engine (UNIX).”
 - “To edit the /etc/inittab file.”

- c. After system-specific post-installation tasks have been completed, stop the Process Engine services manually, then switch the cluster to another node.

CAUTION When the cluster group comes online on the next node, remove all files and directories underneath /fns/local. Repeat the installation procedure starting at [Step a on page 81](#) above.

NOTE Repeat, one server at a time, for all Process Engine nodes.

3. Configure Process Engine for the cluster.

NOTE Perform the following on one node only. The changes made here will apply to the whole cluster.

- a. Start the Process Task Manager.
- b. Make sure the Process Service is stopped.
- c. From the General tab, change the IP Address field to the clustered IP resource's address and click **Apply**.
- d. Start the Process Service from the Process Task Manager, and verify that it starts without errors.

4. Enable automatic start of Process Engine services from the task manager.

Failover to each node and:

- a. Start the Process Task Manager.
- b. Check the box to enable automatic startup of Process Engine Processes.

5. Configure the Process Engine using the highly available Content Engine URL. For more information, see ["Connecting using Content Engine Web Service Transport \(CEWS\)" on page 170](#).

For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task "Configure Process Task Manager."

6. Install the FileNet P8 4.5 Process Engine Agent.

NOTE Prior to putting the Process Engine under cluster control, ensure that you can start Process Services from each node on which the cluster group is running. Also ensure that all CE connection points can connect to Process Engine. If there are any issues at this point then they must be resolved prior to creating cluster resources to control the Process Engine.

To install the agent

- a. Untar the FileNet Process Engine Agent located under the Veritas folder on the *IBM FileNet P8 Process Engine* CD or eSD image.
- b. Copy the files to all nodes in the cluster.
- c. On all nodes perform the following steps:
 - i. Create the following directory:

```
/opt/VRTSvc/bin/FN_ProcessEngine
```
 - ii. Copy FN_ProcessEngineAgent, the agent binary, to this directory.

iii. Change the permissions on the directory as follows:

```
chmod -R 755 /opt/VRTSvcs/bin/FN_ProcessEngine
```

iv. Copy the file `FN_Types.cf` to `/etc/VRTSvcs/conf/config`.

d. Follow Veritas documentation to prepare the cluster for a configuration update.

NOTE This involves stopping Veritas processes on all nodes.

e. On one of the nodes, update the cluster configuration for the new FileNet type definitions.

When ready for update, integrate the FileNet type definitions with the cluster by adding the following line to the beginning of the cluster's configuration file, `main.cf`:

```
include "FN_Types.cf"
```

f. Restart cluster services according to Veritas documentation.

NOTE A Process Engine resource should now be listed as one of the available resource types.

Type the command:

```
hatype -list
```

One of the resources displayed should be `FN_ProcessEngine`.

g. Create a Process Engine resource.

Enable the resource and verify that it shows as "offline."

h. Make the Process Engine resource dependent on the clustered IP and the Process Engine partitions on clustered storage:

- local partition resource
- `fn_sec_db0` partition resource
- `fn_sec_rl0` partition resource

i. Bring the cluster online.

Process Engine installation verification

You will use Application Engine (Workplace) to verify your deployment. A user should be able to use the Process Engine components through Application Engine and not lose the session connection or be prompted for login credentials throughout this procedure.

1. Create a region and assign it a password on the Process Engine and a corresponding region id and connection point on the Content Engine.
2. When creating the region ID on the Content Engine specify the DNS name of the clustered IP resource for the Process Engine server.
3. Log on to workplace.
4. Select a connection point in Site preferences to use for Tasks functionality.
5. Initialize the region.

6. Create and launch a workflow.
7. Initiate a failover of Process Engine in the cluster.
8. Access the user's inbox and retrieve workflow information during/after the failover.

NOTE There might be a slight delay, but eventually the connection point automatically connects to the Process Engine after the failover.

9. Complete a step in a workflow.
10. Create and launch another workflow.
11. Access the user's inbox and retrieve workflow information.

Install Process Engine in a Veritas cluster server (Windows)

Symantec uses “Service Groups” of resources to provide highly available services to users. The following instructions specify how to add resources to a pre-existing service group to provide a highly available Process Engine Service.

NOTE The Veritas custom agent for the Process Engine is intended only for the Solaris, AIX, and HP-UX platforms. For Process Engine on Windows, FileNet recommends using Veritas out-of-the-box agents.

Installation tasks

CAUTION During the first part of the installation the installer will reboot the server. This will fail the cluster group to another node. Before you login to the machine after the reboot to complete the Process Engine installation, login to another machine in the cluster, and move the cluster group back to the node where the current installation is being run. Once that is done you can log on and continue the installation.

NOTE Before starting your installation of Process Engine, follow the configuration steps outlined in [“Appendix A – Setting up a secure native mode domain installation” on page 165](#).

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
2. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Install Process Engine.”

NOTE Run the Process Engine installer as a domain user created for installing Process Engine software as detailed in [“Create FileNet Users.” on page 165](#). You can also use the domain “fns” user to perform the Process Engine installation.

- a. Make sure the cluster group is online on the system you are installing on.
- b. Start the installation. Enter the same installation information for all nodes.

NOTE The NCH domain name should be set to `virtual cluster name:FileNet`. The installer will verify that this parameter matches a network name resource in the cluster and the installation will fail if this parameter does not match.

- c. Install FNSW (executables) on the local drive.
- d. Install FNSW_LOC (local files) on the clustered disk drive.
- e. Verify that the Process Engine Services are working.
 - i. Verify the following services have started and modify them using the Windows Service Control Manager as follows:
 - IMSService
 - Set the startup type to manual.

- Process Engine Services Manager Service
 - Set the startup type to manual.
 - ii. Open the Process Task Manager and verify that Process Service started without error.
 - From the Process Task Manager, stop the Process Service.
 - From the General Tab change the **IP Address** field to the clustered IP resource's address and click **Apply**.
 - iii. Stop the Process Service from the Process Task Manager.
 - iv. Stop the IMS Control Service and Process Engine Services Manager services from the Windows Service Control Manager. Kill any other active Process Engine processes.
 - v. Installation of the Process Engine software on the first node is complete.
- f. Move the cluster group to subsequent nodes in your cluster and install the Process Engine using the steps outlined above.
3. After installation is complete, configure the Process Engine using the highly available Content Engine URL from one of the nodes.
- For information, see [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#). For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure Process Task Manager.”
4. Configure Process Engine for the cluster.
- a. On all nodes perform the following steps:
 - i. In the Process Task Manager, check the box to automatically start the Process Services.
 - ii. Modify the registry to set the NCHBroadcast value.
 - 1. Navigate to:


```
HKEY_LOCAL_MACHINE\Software\FileNET\IMS\CurrentVersion
```
 - 2. Add a new DWORD Value:
 - Name: NCHBroadcast
 - Value: 0
5. Add resources to the cluster to control the Process Engine.
- a. Add a resource for the IMSService Registry Replication.
 - i. Set the Resource Type: RegRep.
 - ii. Set MountResName to the mount resource for the clustered disk.
 - iii. Leave the ReplicationDirectory location on the clustered disk as default or set as desired.

- iv. Set Keys to:
 - HKLM\SOFTWARE\FileNet\IMS\CurrentVersion
 - HKLM\SYSTEM\CurrentControlSet\Services\IMSService
- v. Make the resource dependent on the following resource:
 - Clustered disk resource.
- b. Add a resource for the IMSService.
 - Resource Type: GenericService
 - Set the ServiceName to: IMSService
 - Make the resource dependent on the following resources:
 - Lanman name resource
 - Clustered disk resource
 - IMSService Registry Replication
- c. Add a resource for the Process Engine Services Manager Registry Replication.
 - i. Set the Resource Type: RegRep
 - ii. Set MountResName to the mount resource for the clustered disk.
 - iii. Leave the ReplicationDirectory location on the clustered disk as default or set as desired.
 - iv. Set Keys to:
 - HKLM\SYSTEM\CurrentControlSet\Services\VWServicesPE
 - v. Make the resource dependent on the following resource:
 - Clustered disk resource
- d. Add a resource for the Process Engine Services Manager.
 - i. Set the Resource Type: GenericService
 - ii. Set the ServiceName to: VWServicesPE
 - iii. Make the resource dependent on the following resources:
 - IMSService
 - Process Engine Services Manager Registry Replication
 - iv. Set the Offline Timeout value for the resource to a minimum of 60 seconds.
 - v. Set the Clean Timeout value for the resource to a minimum of 60 seconds.
6. Bring the cluster group online.

Process Engine installation verification

You will use Application Engine (Workplace) to verify your deployment. A user should be able to use the Process Engine components through Application Engine and not lose the session connection or be prompted for login credentials throughout this procedure.

1. Create a region and assign it a password on the Process Engine and a corresponding region ID and connection point on the Content Engine.

NOTE When creating the region ID on the Content Engine Specify the DNS name of the clustered IP resource for the Process Engine server.

2. Log on to Workplace.
3. In Site Preferences select a connection point to use for Tasks functionality.
4. Create and launch a workflow.
5. Initiate a failover of Process Engine in the cluster.
6. Access the user's inbox and retrieve workflow information during/after the failover.

NOTE There might be a slight delay, but eventually the connection point automatically connects to the Process Engine after the failover.

7. Complete a step in a workflow.
8. Create and launch another workflow.
9. Verify that you can access the user's inbox and retrieve workflow information.

Install Process Engine in a Microsoft cluster server

Microsoft Cluster Server (MSCS) provides a framework to create “Groups” of resources that collectively provide a highly available service. The following instruction details how to add to a Group of resources in order to provide a highly available Process Engine.

Deployment tasks

NOTE If you want to perform this installation in a domain that is running in Windows Server 2003 mode, or want to perform the installation as a specific user, you must pre-create the domain users and groups for the Process Engine. Refer to the configuration steps outlined in [“Appendix A – Setting up a secure native mode domain installation” on page 165.](#)

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49.](#)
2. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Install Process Engine.”

NOTE Run the Process Engine installer as a domain user who has at least Account Operator privileges or as the domain user created in [“Create FileNet Users.” on page 165.](#)

- a. Verify that the cluster group is online on the system where you are installing.
- b. Start the installation. Enter the same installation information for all nodes.

NOTE The NCH domain name should be set to `virtual cluster name:FileNet`. The installer will verify that this parameter matches a network name resource in the cluster and the installation will fail if this parameter does not match.

- c. Install FNSW (executables) on the local drive.
- d. Install FNSW_LOC (local files) on the clustered disk drive.
- e. During the first part of the installation the installer will reboot the server.

CAUTION The rebooting will most likely fail the cluster group to another node. Before you login to the machine after the reboot to complete the Process Engine installation, login to another machine in the cluster and fail the cluster group back to the node where the current installation is being run. Once that is done you can login and continue the installation.

- f. Verify that the Process Engine Services are working.
 - i. Start the IMS Control Service and Process Engine Services Manager services from the Windows Service Control Manager.
 - ii. Open the Process Task Manager and verify that the Process Service has started without error.

If the IP Address being used by the Process Service isn't the cluster IP address listed in the PE cluster group then:

- From the Process Task Manager, stop the Process Service.

- From the General Tab change the **IP Address** field to the clustered IP resource's address and click **Apply**.
 - Start the Process Service from task manager to verify that the service starts correctly.
- iii. Stop the Process Service from the Process Task Manager.

Stop the IMS Control Service and Process Engine Services Manager services from the Windows Service Control Manager. Kill any additional Process Engine processes still active on this node.
 - iv. Installation of the Process Engine software on the first node is complete.
 - g. Move the cluster group to subsequent nodes in your cluster and install the Process Engine using the steps outlined above.
3. After installation is complete, from one of the nodes configure the Process Engine using the highly available Content Engine URL.

For information, see ["Connecting using Content Engine Web Service Transport \(CEWS\)" on page 170](#). For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task Configure the Process Task Manager.
 4. Configure the Process Engine for the cluster.

On all nodes perform the following steps:

 - a. In the Process Task Manager, check the box to automatically start the Process Services.
 - b. Modify the registry to set the NCHBroadcast value.
 - i. Navigate to:

```
HKEY_LOCAL_MACHINE\Software\FileNET\IMS\CurrentVersion
```
 - ii. Add a new DWORD Value:
 - Name: NCHBroadcast
 - Value: 0
 5. Add resources to the cluster to control the Process Engine.
 - a. Add a resource for the IMSService.
 - i. Set the Resource Type: Generic Service
 - ii. Set the Service name to: IMSService
 - iii. Check the **Use Network Name for computer name** box.
 - iv. Make the resource dependent on the following resources:
 - Clustered network name
 - Clustered disk resources

- v. Add values to the Registry Replication tab:

```
software\filenet\ims\currentversion  
system\CurrentControlSet\Services\IMSService
```

- b. Add a resource for the Process Engine Services Manager.

- i. Set the Resource Type: Generic Service
- ii. Set the Service name to: VWservicesPE
- iii. Check the **Use Network Name for computer name** box.
- iv. Make the resource dependent on the following resource:
 - IMSService

- v. Add values to the **Registry Replication** tab:

```
system\CurrentControlSet\Services\VWservicesPE
```

- c. Bring the cluster group online.

Deployment verification

You will use Application Engine (Workplace) to verify your deployment. A user should be able to use the Process Engine components through Application Engine and not lose the session connection or be prompted for login credentials throughout this procedure.

1. Create a region and assign it a password on the Process Engine and a corresponding region ID and connection point on the Content Engine.

Perform the tasks “Create a Process Engine Isolated Region” and “Create a Process Engine Connection Point” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE When creating the region ID on the Content Engine Specify the DNS name of the clustered IP resource for the Process Engine server.

2. Configure the Process Engine Connection Point for Application Engine

Perform the task “Configure the Process Engine Connection Point for Application Engine” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

3. Create and launch a workflow.
4. Initiate a failover of Process Engine in the cluster.
5. Access the user's inbox and retrieve workflow information during/after the failover.

NOTE There might be a slight delay, but eventually the connection point automatically connects to the Process Engine after the failover.

6. Complete a step in a workflow.
7. Create and launch another workflow.
8. Verify that you can access the user's inbox and retrieve workflow information.

Process Engine database reconnection

Each PE background process (VWRs, VWKs, VWSs, and vwtime) maintains an open database connection. Once that database connection is lost, due to database failover or other unexpected db connection problem, each of the background processes attempts to create a new connection for up to 9 minutes. Once a new connection is retrieved (the network connection is reinstated, the database comes back on-line, a replacement highly available database provides a new connection, etc.) or the retry connection time limit has been reached, an error is logged and that error might be sent back to the application level.

- If a new connection is successfully acquired, an error message is sent back to the calling routine telling the internal program routine that requested information from the database to retry the operation as the connection was re-established.
- If a new connection cannot be created within that time frame the returned error will be a “database not available” message.

NOTE Each background process only knows there is a lost connection when some task tries to access the database; therefore, each open transaction at the time of failure will have to be retried. All existing transactions for that process are rolled back and the transaction information will be reset

The transaction retries will either take place automatically within the PE software, or at the application level where the software can either display an error message for the user to handle or the program can catch the error and execute some routine (auto retry, etc).

Examples of errors returned on database disconnect

Possible errors returned to the application level:

Tuple:	213,109,173
Message:	Invalid or Lost database connection.
Description:	Database connection was lost. Could be a network error or other problems on the RDBMS server. This error is returned when a connection was once valid, but is no longer valid.
Action:	Contact your DBA or System Administrator.

Tuple:	213,109,174
Message:	Retry last DB transaction.
Description:	Database connection was lost but then connection was re-established. You should be able to now retry the last or current transaction where error occurred. This error is returned after a connection was re-established to the RDBMS.
Action:	With this error you can retry the last step or transaction you were running.

Additional errors or messages logged in the application event viewer or FileNet system log-files, showing the progress of reconnecting to the database.

Tuple:	121,y,xxx (xxx being the database error condition, y being the database type - 0 = generic all, 1 = Oracle, 7 = MSSQL and 9 = DB2)
Example:	Shows error message from database. 2005/02/17 10:14:08.651 121,1,12535 <i>fnsw</i> VW/eProcess (8697) ... [SERIOUS] Oracle Error ORA-12535: Unable to recover error message., File: ../src/GDBO.c, Line 2463
Example:	Shows progress message where PE software in continuing to reconnect to the database. The number '3' in the second line shows the number of reconnection attempts 2005/02/17 10:14:08.760 121,1,12535 <i>fnsw</i> VW/eProcess (8697) ... VW (Region=222): Lost DB connection reconnecting - 3

Known issues

1. If the following error shows up during a recovery scenario, it can be ignored:

```
2005/02/18 14:28:52.343 121,0,7 fnsw VW/eProcess (3592.804.484 0xe08.324) ...
VW (Region=500): Can't free GDB statement handle
```

2. If a failover occurs during a PE database initialization, you should restart the PE software before attempting to re-initialize.
3. If the PE failover occurs in a farm environment and the user is trying to save data to the region or transfer workflows, it can take up to 90 seconds before the user can retry the tasks.

Install a highly available Application Engine/Workplace XT

This section covers the tasks and requirements for making Application Engine (AE) or Workplace XT highly available in a FileNet P8 environment.

ATTENTION This topic covers installation of Workplace XT in a supported FileNet P8 HA environment. Do NOT install this component unless it is supported at the release levels of your FileNet P8 environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure Application Engine/Workplace XT

Installation tasks

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49.](#)
2. (JBoss) Verify that the Content Engine and Application Engine clusters have different names. For information, see [“JBoss clustering of Content Engine and Application Engine” on page 52.](#)
3. Mount the shared configuration directory from all nodes in the application server configuration.

Once the location of the shared configuration directory is decided, that directory must be mounted and accessible from all servers in the application server configuration. Also, ensure that the same directory path is used on all systems.

Example: If the path on the first system is:

Application Engine

```
/home/AE/Config
```

Workplace XT

```
/home/WebClient/Config
```

make sure this same path is available on all other systems.

For more information, see [“Create a local or shared directory for the shared configuration files” on page 58.](#)

4. (WebSphere and WebLogic) Verify that your application farm or cluster configuration is running.
NOTE JBoss deploys applications on startup. There is no need to start the application server until you have installed the application.

5. Install Application Engine/Workplace XT on all nodes.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task Install and configure Application Engine, or the *IBM FileNet Workplace XT Installation Guide* task Install IBM FileNet Workplace XT with the following additional information:

- (On the Content Engine API Configuration screen) If the Content Engine has been configured for high availability, make sure to specify the proper Content Engine URLs following the format and examples listed in [“Content Engine in an application server cluster using EJB transport” on page 169](#).
- (On the Specify Configuration Directory screen) Specify the location to hold the Application Engine/Workplace XT shared configuration data.

NOTE You must install Application Engine/Workplace XT on all nodes in the application server configuration. Even in an application server cluster configuration the Application Engine/Workplace XT provides components that run outside of the application server such as the Component Manager. These are only available through an Application Engine/Workplace XT installation.

6. Synchronize the UTCryptoKeyFile.properties file across your Application Engine/Workplace XT nodes.

Each participating Application Engine or Workplace XT server must use the same encryption key file. Copy the `UTCryptoKeyFile.properties` file installed on the first Application Engine or Workplace XT server in your farm to all the other servers in your farm.

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

7. Configure and deploy the Application Engine/Workplace XT for your application server configuration.

NOTE JBoss clusters do not have an administrative server. You must configure and deploy Application Engine/Workplace XT on each individual node per [“\(Farm of independent application server instances\) Configure and deploy the Application Engine/Workplace XT on each node.” on page 97](#).

- (Application server cluster) Configure and deploy Application Engine/Workplace XT from the administrative server instance.

i. Configure your web application.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure Application Engine (*Application server*)”, or the *IBM FileNet Workplace XT Installation Guide* task “Configure IBM FileNet Workplace XT(*Application server*)” to configure the Workplace/Workplace XT application using the Administration server.

NOTE For steps that require Java Virtual Machine (JVM) settings to be made make sure to make these changes for every node in the application server configuration.

ii. Deploy your web application.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Deploy Application Engine (*Application server*)”, or the *IBM FileNet Workplace XT*

Installation Guide task “Deploy IBM FileNet Workplace XT (*Application server*)” to deploy the Workplace/Workplace XT application using the Administration server.

NOTE For steps that require JVM settings to be made make sure to make these changes for every node in the application server configuration.

- web.xml file

If you need to edit the web.xml file as part of deployment, for example if you are using single sign-on (SSO), make these changes from the administrative server node only. In addition, to include these changes you must recreate the WAR or EAR file before deploying the application, see “To recreate the WAR or EAR file” procedure in the “Deploy Application (*Application server*)” task.

Most application server cluster configurations deploy an application to managed servers by copying the application files from the administrative node to the managed server, and only changes made on the administrative node are deployed to the configuration.

NOTE An additional method for deploying applications to an application server cluster is to maintain a copy of the application on all servers as in a farm of independent servers instead of copying files. In this type of configuration, edits to files should be performed on all nodes in the application server cluster.

- (WebSphere and WebLogic) When installing the Application Engine/Workplace XT through the administrative server make sure to deploy the Application Engine/Workplace XT to the cluster of server instances:

WebSphere

For WebSphere make sure you select the cluster and/or managed nodes that will be running the application when mapping modules to servers (the admin node is usually the default, so it is typical to change this). Click apply so that the application shows as deployed to the proper nodes instead of the default single node.

UTF-8 encoding parameters should be set on every node.

NOTE (WebSphere 6.x) If using an IBM HTTP Server to proxy/load-balance application requests to WebSphere, make sure to select the webserver node when mapping modules to the application.

WebLogic

For WebLogic make sure the targets for the application include the WebLogic cluster and/or managed nodes that will be running the application.

- (WebSphere only) Perform any LTPA configuration steps on the administrative server only.

NOTE The scope of this action will affect the entire application server configuration.

- (Farm of independent application server instances) Configure and deploy the Application Engine/Workplace XT on each node.

When installing in farm of independent application server instances the steps listed in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Deploy Application Engine (Application Server)”, or the *IBM FileNet Workplace XT Installation Guide* task “Deploy IBM FileNet Workplace XT (Application server)” must be performed on every node.

NOTE Since there is no central administration server every server acts as a separate server instance and must be configured as such.

8. Configure Server Ports

Perform the “To configure the server ports” section in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure Application Engine (Application server),” or the *IBM FileNet Workplace XT Installation Guide* task “Configure IBM FileNet Workplace XT (Application server)” on every node in the application server configuration.

9. (JBoss only) Configure jndi.properties to allow CE failover.

On each node, edit the following file:

`JBoss_HOME/server/all/conf/jndi.properties`

adding the line:

```
java.naming.provider.url=jnp://
CE_server_1:1100,CE_server_2:1100,...,CE_server_n:1100
```

where `CE_server_1` ... `CE_server_n` are your Content Engine nodes.

CAUTION This edit will cause Jboss to display a number of exceptions and error messages of type “NameAlreadyBoundException” at startup. These errors do not impact AE performance and can be ignored.

10. Configure your environment to connect to a highly available Content Engine.

Perform the “Enable the Process Engine Component Integrator” task in *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE Make sure you configure your environment to connect to a highly available Content Engine, see “[Appendix B – Connecting to a highly available Content Engine](#)” on page 168.

Application Engine/Workplace XT deployment verification

1. Using the application server administration tool, verify that the Workplace/Workplace XT application is running.
2. Open a web browser and type in the URL for the Workplace/Workplace XT application.
 - If using a load-balancer or proxy server the URL has the form:

Application Engine

`http://virtual_name:port_number/Workplace`

Workplace XT

`http://virtual_name:port_number/WorkplaceXT`

- If connecting to an individual cluster node the URL would have the form:

Application Engine

`http://clustered_server_name:port_number/Workplace/`

Workplace XT

`http://clustered_server_name:port_number/WorkplaceXT/`

NOTE Sign-in using the virtual name to properly set the base Workplace/Workplace XT URL. If it is not set correctly you must change it in Site Preferences.

3. Verify that the sign-in page displays, and log in as a user.
4. Fail one of the nodes.
5. Verify that you can reload the sign-in page, and log in as the same user.

Enabling the Process Engine Component Integrator

When enabling the Process Engine Component Integrator (optional) in an HA environment that contains more than one Application Engine node, you must configure Process Engine to allow the component integrator running on the active Application Engine node to query Process Engine for work.

If you configure Process Engine to talk to a specific Application Engine node, it would only attempt to communicate with the component integrator on that node, and will not be able to switch to another node in case of an Application Engine failover.

Follow the instructions in the Enable the Process Engine Component Integrator task of the *IBM FileNet P8 Platform Installation and Upgrade Guide* with the following modification:

- In the “To specify connection between Process Engine and Component Manager (on a Process Engine server)” procedure, leave the Host field empty.

Install a highly available Rendition Engine

This section covers the tasks and requirements for making Rendition Engine (RE) highly available in a FileNet P8 environment.

ATTENTION This topic covers installation of Rendition Engine in a supported FileNet P8 HA environment. Do NOT install this component unless it is supported at the release levels of your FileNet P8 environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure Rendition Engine

Installation tasks

Follow the instructions in the *IBM FileNet Rendition Engine Installation and Upgrade Guide* “Installation Tasks” section with the following additional information:

1. Install Content Engine, Process Engine, and Application Engine.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

For high availability installation of these components, see [“Install a highly available Content Engine” on page 70](#), [“Install a highly available Process Engine” on page 79](#), and [“Install a highly available Application Engine/Workplace XT” on page 94](#).

2. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).

3. Install and configure the first Rendition Engine.

- a. Perform task “Install Rendition Engine.”

- b. Perform task “Configure Rendition Engine connections on Content Engine.”

NOTE If CE is clustered or farmed use the virtual IP address for the CE cluster or farm when configuring RE.

- c. Perform task “Set the site preferences on the Application Engine to use path-based URLs.”

NOTES

- You only need to do this once per HA environment. configuring RE.
- Setting the site preferences is not required on an upgrade. However, you should verify that your settings are valid.

- d. Perform task “Enable word processing macros for rendering.”

NOTE Do this on each Rendition Engine server.

- e. Perform task “Install Rendition Engine software updates.”

- f. Perform the optional task “(Optional) Verify Installation.”

Before setting up any additional Rendition Engine servers you should verify that the newly installed environment works without problems.

4. Install additional Rendition Engines.

- a. Perform task “Install Rendition Engine.”

In the “SQL Server Database Information” screen, for a farm environment, make sure all nodes use the same database information. If you are using a highly available database make sure you provide the virtual name of the database server:

CAUTION After configuring the VistaTempDir in step “Set security on the VistaTempDir shared folder” you must verify that you can access this directory from each existing Rendition Engine server using the *FNRE_admin* user account, and that you can access the VistaTempDir on each existing Rendition Engine from the current server.

To verify accessibility:

- i. On the Rendition Engine server, select Start > Run and then type
`\\Other_RE_Server\VistaTempDir.`
 - ii. Verify that you can create a text document in the directory by right-clicking and selecting **New > Text Document**.
- b. Perform task “Enable word processing macros for rendering.”
NOTE Do this on each Rendition Engine server.
 - c. Perform task “Install Rendition Engine software updates.”

Rendition Engine installation verification

Perform the optional task “(Optional) Verify Installation.”

With the addition of each new Rendition Engine server you should verify that the environment works without problems.

Install a highly available Content Search Engine

This section covers the tasks and requirements for making Content Search Engine (CSE) highly available in a FileNet P8 environment.

Install and configure Content Search Engine

Installation tasks

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
2. Verify that your cluster is running.
3. Verify that the shared directories are mounted on each node.
 - Binary Share (UNIX, optional)
 - Temp Directory Share

For information, see [“Create required shared directories” on page 57](#).

4. Install CSE on all cluster nodes.

NOTE In a highly available environment, you must run the CSE installer on each node, selecting the shared directory as the “Content Search Engine Installation Path.”

- a. Install CSE on a cluster node.

Follow the instructions in the “Install Content Search Engine” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

- Use the shared binary location as installation location (*CSE_install_path*).

UNIX

On UNIX, this may either be a local shared cluster disk or a highly available shared storage location

Windows

On Windows this must be a local shared cluster disk.

- In the “Select Autonomy K2 Server Type” screen, select Master Administration Server.
- In the “Specify Master administration Server information” screen, use the hostname of the cluster rather than the hostname of the local cluster node.
- When the installer has finished, K2 will be running. Stop the application before failing to another node.

UNIX

Execute:

```
k2adminstop
```

Windows

Stop the "Verity K2 Administration Server" and "Verity K2 Administration Web Server" services.

b. Fail over to another node and repeat the install process with the same parameters.

5. Configure local files to point to cluster address.

Once installation has been completed on all nodes, the following files will contain local hostnames and must be edited to have the local node names changed to the cluster virtual hostname:

- `CSE_install_path/_cseuninst/installvariables.properties`
- `CSE_install_path/config.vcnf`
- `CSE_install_path/data/docs/WEB-INF/web.xml`
- `CSE_install_path/k2/common/base.xml`
- `CSE_install_path/k2/common/verity.cfg`
- `CSE_install_path/data/host/admin/adminN.xml`

Where the value of *N* is stored in `CSE_install_path/data/admin/admin.dat`.

6. Configure Veritas or Microsoft Cluster Server (MSCS) CSE Resource Group

After installing CSE on your cluster you must configure the following cluster resources. These should all be placed in the same Resource or Service group so that they fail over together:

- Network name (Windows) / IP Address (UNIX)
- Binary File Share (if local to the cluster rather than mounted from a remote highly-available storage location)
- Temporary File Share (if local to the cluster rather than mounted from a remote highly-available storage location)
- Collections File Share

For the preceding two bullets, follow the instructions in the first two steps of the "Create a Content Search Engine collections directory" task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, to create a temp directory and to create a collections directory.

- CSE services

UNIX

Configure the Veritas Application Agent as follows:

- User - `k2_os_user`
- StartProgram - `/CSE_install_path/_platform/bin/k2adminstart`
- StopProgram - `/CSE_install_path /k2/_platform/bin/k2adminstop`

- PidFiles - /CSE_install_path/data/host/k2admin.pid /opt/verity/data/host/catalina.pid

NOTE It is possible to start and stop the K2 Administration Server and K2 Administration Web Server separately on UNIX, but doing so requires custom scripting. See the *FileNet P8 High Availability Solutions Redbook* for more information.

Windows

- Verity K2 Administration Server (required)
 - Verity K2 Administration Web Server (optional)
7. Add the following resource dependencies:
 - Collections File Share
 - IP Address (UNIX)/Network name (Windows)
 - Collections mount point
 - K2 Administration Server / Administration Web Server
 - Binary and temporary file shares (if these are clustered resources. If they are mounted from a remote highly available location, this is not necessary)
 - Collections File Share
 - IP Address (UNIX)/Network Name (Windows)
 8. Configure CSE per install guide instructions, for your cluster.

Follow the instructions in the “Configure Content Search Engine” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

 - Use the cluster name for the K2 Dashboard URL entry.
 9. Create a Content Search Engine collections directory

Follow the instructions in the “Create a Content Search Engine collections directory” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

 - The steps to create a collections directory, and to create a temp directory were done as part of [“Configure Veritas or Microsoft Cluster Server \(MSCS\) CSE Resource Group” on page 102.](#)
 10. Configure Content Engine for content-based retrieval.

Follow the instructions in the “Configure Content Engine for content-based retrieval” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following additional information:

 - When configuring the Verity Domain, use the cluster hostname.

Install a highly available CFS-ICI

CAUTION This topic covers installation of Content Federation Services for IBM Content Integrator (CFS-ICI) in a supported FileNet P8 HA environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure CFS

Installation tasks

For more information about CFS installation, see the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide*.

1. Install Content Engine, Process Engine, and Application Engine.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

For high availability installation of these components, see [“Install a highly available Content Engine” on page 70](#), [“Install a highly available Process Engine” on page 79](#), and [“Install a highly available Application Engine/Workplace XT” on page 94](#).

2. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
3. Install and configure IBM Content Integrator (ICI) on each node in the cluster.

Follow the installation instructions provided in the *Content Integrator Enterprise Edition Support Web site* at <http://www.ibm.com/software/data/content-management/support/content-integrator/>.

4. Configure the ICI/CFS environment.

Perform the configuration tasks as outlined in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* with the following HA related modifications:

- “Install updated IBM Content Integrator connectors” - Do this on all nodes.
- “Create and enable session pools” - Do this on all nodes.
- “Create Content Engine to IBM Content Integrator connector”

When configuring the FileNet P8 Content Manager Connector use the following URLs to point to the CE Cluster:

WebSphere

Remote Server URL:

```
cemp:iiop://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

Remote Server Upload URL:

```
cemp:iiop://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```


Remote Server Download URL:

```
cemp:iiop://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

This configuration requires the WebSphere cluster name in addition to the node names as part of the URL. The default bootstrap port (e.g. 9810) is different from the non-default (e.g. 2809) on a non cluster configuration (i.e. a standalone WebSphere installation).

WebLogic**Remote Server URL:**

```
cemp:t3://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

Remote Server Upload URL:

```
cemp:t3://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

Remote Server Download URL:

```
cemp:t3://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

JBoss**Remote Server URL:**

```
cemp:jnp://node1 name:node1 port, node2 name:node2 port/FileNet/Engine
```

Remote Server Upload URL:

```
cemp:jnp://node1 name:node1 port, node2 name:node2 port/FileNet/Engine
```

Remote Server Download URL:

```
cemp:jnp://node1 name:node1 port, node2 name:node2 port/FileNet/Engine
```

- “Start RMI Bridges“ - Do this on all nodes.

5. Install and Deploy CFS.

Perform the installation and deployment tasks outlined in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* with the following HA related modifications:

- Install CFS

Perform either the “Install CFS interactively” or “Install CFS silently” tasks outlined in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* with the following modifications:

- In the ICI Configuration panel set the SSO Server URL and Configuration Server URL as follows:

ICI SSO Server URL:

```
rmi://node1 name:node1 port/SSOServerName,rmi://node2 name:node2 port/SSOServerName
```

IICE Configuration Server URL:

```
rmi://node1 name:node1 port/ConfigurationServerName,rmi://node2 name:node2 port/ConfigurationServerName
```

- “Deploy CFS Federator Administration application”

- Deploy to a cluster

WebSphere and WebLogic

Perform the deployment on the administrative server.

JBoss

Deploy to each farm node.

- “Start and Stop CFS Exporter” - Do this on all nodes.
- “Create Data Source for CFS Federator Database”

WebSphere and WebLogic

Do this on the administrative server.

JBoss

Deploy to each farm node.

- “Configure Content Engine to Support CFS” - Enter the Configuration Server URL as follows:

ICI SSO Server URL:

```
rmi://node1 name:node1 port/SSOServerName,rmi://node2 name:node2 port/SSOServerName
```

IICE Configuration Server URL:

```
rmi://node1 name:node1 port/ConfigurationServerName,rmi://node2 name:node2 port/ConfigurationServerName
```

6. Synchronize Connector, Data Maps, and Session Pools across the nodes.

NOTE This is not a task in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide*.

After configuring the ICI Connectors, Data Maps and Session Pools on one node, export the settings to XML files.

Copy these files to the other node and import them using the ICI Admin Tool. Refer to the ICI documentation for more details on the export/import options.

NOTE Make sure the JAR files in the IICE_HOME/lib directory are identical on all nodes.

CAUTION You must do this each time you make changes to your environment, such as “Set Up Additional External Repositories” below.

7. Perform Post Installation Steps.

Perform the tasks outlined in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* with the following HA related modifications:

- “Set up additional external repositories”

NOTE After making configuration changes like these you must redo [Step 6 “Synchronize Connector, Data Maps, and Session Pools across the nodes.”](#) on page 106.

Deployment verification

To verify the installation follow the instructions in the “Verify the CFS Installation” task in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide*.

To verify correct behavior on failover, follow the instructions below.

1. Using the application server administration tool, verify that the CFS Federator Administration application is running.
2. Open a web browser and type in the URL for the CFS Federator Administration application.
 - If using a load-balancer or proxy server the URL has the form:
`http://virtual_name:port_number/FedAdmin`
 - If connecting to an individual cluster node the URL would have the form:
`http://clustered_server_name:port_number/FedAdmin`
3. Verify that the sign-in page displays, and log in with ICI Administrator credentials.
4. Fail one of the nodes.
5. Verify that you can reload the sign-in page, and log in as the same user.

Install a highly available eForms for P8

This section covers the tasks and requirements for making eForms for P8 highly available in a FileNet P8 environment.

CAUTION This topic covers installation of FileNet eForms for P8 in a supported FileNet P8 HA environment. Do NOT install this functional expansion unless it is supported at the release levels of your FileNet P8 environment. For information, see the *FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure eForms

Installation and deployment tasks

1. Install Content Engine, Process Engine, and Application Engine or Workplace XT.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* and *IBM FileNet Workplace XT Installation and Upgrade Guide*.

For high availability installation of these components, see [“Install a highly available Content Engine” on page 70](#), [“Install a highly available Process Engine” on page 79](#), and [“Install a highly available Application Engine/Workplace XT” on page 94](#).

2. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
3. Install eForms on each Application Engine or Workplace XT node.

Follow the instructions in the “Installation Tasks” section of the *IBM FileNet eForms for P8 Installation and Upgrade Guide* with the following additional information:

WebSphere or WebLogic

When deploying the Application Engine or Workplace XT web application to the application server make sure to deploy it to the application server cluster.

JBoss

When deploying the Application Engine or Workplace XT web application to the application server, copy it to the deploy folder for each server in a JBoss cluster.

4. Configure and start eForms.
5. Follow the instructions in the “Configuration/Startup Tasks” section of the *IBM FileNet eForms for P8 Installation and Upgrade Guide* with the following additional information:

You only need to perform the configuration and startup tasks for eForms once from an Application Engine or Workplace XT application if the same region and site preference is used for all instances of Application Engine and Workplace XT.

eForms deployment verification

NOTE For this procedure you will need a form template created using the Form Designer application.

1. Make sure Application Engine or Workplace XT are running.
2. Login to Workplace or Workplace XT.

In a web browser, enter the URL for the Application Engine or Workplace XT application, and log in.

3. Add a form template:

NOTE Add the form template as you would any other document but specify it as a document class type of "ITX Template."

4. After you add it, verify that it successfully displays in Application Engine or Workplace XT.

Install a highly available Process Analyzer

CAUTION This topic covers installation of Process Analyzer (PA) in a supported FileNet P8 HA environment. Do NOT install this functional expansion unless it is supported at the release levels of your FileNet P8 environment. For information, see the *FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure Process Analyzer

Microsoft Cluster Server (MSCS) provides a framework to create “Groups” of resources that collectively provide a highly available service. The following instruction details how to add to a Group of resources in order to provide a highly available Process Analyzer.

Installation tasks

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49.](#)
2. Install Process Analyzer Engine on all nodes.

Follow the instructions in the *IBM FileNet Process Analyzer Installation and Upgrade Guide* task Install and configure Process Analyzer Engine, with the following additional information:

- Database Location should be remote, using virtual IP/name of SQL Server.
 - The installer creates the PA database. When you install on the second node and supply the same database name, the installer will detect that the database already exists.
 - If this is a new install of PA, the installer can safely overwrite the existing database.
 - If this is not a new install and the PA database already contains information, specify a new name for a temporary database that can be created. When the installation completes, edit `analyzer.properties` and change the `analyzer.db.name` property to the name of the correct database.
 - The Content Engine API Configuration screen should use the virtual IP/name of the CE.
 - Don't select any components to be started automatically when the installation is complete.
3. Installing the latest Content Engine and Process Engine Client files on Process Analyzer server

Follow the instructions in the *IBM FileNet Process Analyzer Installation and Upgrade Guide* task Configure Process Analyzer Engine, with the following additional information:

- Make sure that the PA Resource Group is taken offline manually using the Cluster Administrator tool prior to running the client installers.
- When installing the PE client software, uncheck the box in the PE client installer that asks if you want the installer to stop the BPM services. The PA service will already be offline.

NOTES

- If you let the client installer stop the Process Analyzer Services Manager, the cluster will restart it immediately.
- The PE client installer will automatically start the Process Analyzer Services Manager at the end of installation

4. Configure Process Analyzer Engine

Follow the instructions in the *IBM FileNet Process Analyzer Installation and Upgrade Guide* tasks “Install the latest Content Engine Client files on Process Analyzer server” and “Install the latest Process Engine Client files on Process Analyzer server“, with the following additional information:

- In Process Task Manager, be sure the set the Temp Directory UNC path to the virtual IP/name of the PA Temp Share, not to the IP/name of an individual cluster node.
- When configuring the Process Engine Database connection, use the virtual IP/name of the PE database cluster.
- When configuring the Content Engine URI, use the virtual IP/name of the CE cluster or CE farm load balancer.
- Check "Automatically Started" in Process Task Manager and hit "Apply". Otherwise PA won't be started when the VWServicesPA service is brought online and PA will need to be manually started when the cluster fails over.

5. Installing the PA Client

You can install PA client in a single location that multiple end users will access or on each end user workstation. The determining factors should be the number of end users who will access the reports and the sophistication of the end users.

Follow the instructions in the *IBM FileNet Process Analyzer Installation and Upgrade Guide* task Install Process Analyzer Client, with the following additional information:

- When running the Process Analyzer Client Configuration Tool, use the virtual IP/name of SQL Server Analysis Services server.

Install a highly available Records Manager

CAUTION This topic covers installation of Records Manager (RM) in a supported FileNet P8 HA environment. Do NOT install this functional expansion unless it is supported at the release levels of your FileNet P8 environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure Records Manager

Pre-installation tasks

The following procedure assumes that the application server already has a cluster instance configured and is ready for applications to be deployed. Aside from the configuration data (`bootstrap.properties`), the remainder of the Records Manager software does not require network or clustered storage resources. The Records Manager software can be installed onto local storage for each machine in the cluster.

1. Install Content Engine, Process Engine, and Application Engine.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

For high availability installation of these components, see [“Install a highly available Content Engine” on page 70](#), [“Install a highly available Process Engine” on page 79](#), and [“Install a highly available Application Engine/Workplace XT” on page 94](#).

2. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 49](#).
3. Verify that the cluster nodes and load balancers are started if the application server used to cluster Records Manager is WebLogic or WebSphere.
4. Mount the location of the `bootstrap.properties` file (a highly available share is recommended) from each server in the cluster.

This location should be mounted to the same path on every server. For example, if the `bootstrap.properties` file is mounted to `/opt/FileNet/bootstrap` on one node then it should be mounted to the same directory on each subsequent node.

A location for the `bootstrap.properties` file should already exist as part of the Application Engine high availability configuration. To ensure that each Records Manager instance functions identically in the cluster, it is important that each instance uses the same `bootstrap.properties` file. The Records Manager installer allows you to select the location of the `bootstrap.properties` file during installation.

CAUTION Do not host the `bootstrap.properties` file on an application server cluster node. Although the `bootstrap.properties` file can be installed locally with its location shared or NFS exported for other nodes to access, this configuration introduces a single point of failure. If the node hosting the file fails, then the remainder of the cluster nodes will cease to function.

Host this file on a highly available share exclusive of the Application Engine/Records Manager cluster.

Installation tasks

For more information about Records Manager installation, see the “Install Records Manager” task in the *IBM FileNet Records Manager Installation and Upgrade Guide*.

1. Install Records Manager on each node of the application server cluster.

Refer to the *Install Records Manager* task.

NOTES

- These instructions assume that the Application Engine has not yet been deployed and that the site preferences have not yet been set. If Records Manager is to be installed on an existing Application Engine cluster, the site preferences object in the Content Engine’s object store must be checked out via Enterprise Manager to verify the correct URL for Workplace. The XML element `<setting key="workplaceBaseURL" >` must be set to URL of the load-balancer or proxy server in the Application Engine farm. If the workplaceBaseURL is not set, check out the XML document, correct the setting, and check it back in.
- In the *Install Records Manager* task you are required to specify the server name of the Content Engine. If Content Engine is clustered, specify the cluster IP address of the service group that the API Listener belongs to. This is to ensure that all Records Manager cluster nodes are able to connect to the active node of the Content Engine cluster.
- This task also requires the path to the `bootstrap.properties` file. Select the location created to host the Application Engine configuration data here (likely a Windows share or NFS export location).

CAUTION Do not install Records Manager nodes with independent `bootstrap.properties` files.

2. Configure the application server cluster and deploy Application Engine/Workplace XT according to the *IBM FileNet P8 Platform Installation and Upgrade Guide* or *IBM FileNet Workplace XT Installation Guide*.
3. Deploy Records Manager.

WebSphere

Refer to the *Deploy the Records Manager Web Application* task, *IBM WebSphere 5.x, 6.x, or 7.x* procedures in the *IBM FileNet Records Manager Installation and Upgrade Guide* after the cluster and load balancer instances have been created using the WebSphere administrative console. Using these steps deploy Records Manager on the cluster and proxy server instances.

WebLogic

Refer to the *Deploy the Records Manager Web Application* task, *BEA WebLogic 8.1.x, 9.x, or 10.x* procedures in the *IBM FileNet Records Manager Installation and Upgrade Guide* after the WebLogic domain, administration server and cluster have been created. Use these steps to deploy Records Manager on the cluster.

JBoss

On JBoss you will configure and deploy the RM Web application in a farm of independent application server instances.

NOTE In this configuration, since there is no central administration server, every server acts as a separate server instance and must be configured as such.

When installing in farm of independent application server instances the steps listed in the *IBM FileNet Records Manager Installation and Upgrade Guide* task “Deploy on JBoss 4.0.x” must be performed on every node.

4. Running the RMCreationTool and RMAccessRole tool in an HA environment.

To successfully add Add Marking Sets and AddOns to the FileNet P8 domain and set up the Classified Records Security Personnel Access Role for a highly available Records Manager you must run the RMCreationTool and the RMAccessRole tool on one Records Manager node only once each per highly available Records Manager environment.

The tools use the `RemoteserverURL` entry in the local `WcmApiConfig.properties` file to complete the configuration. However, to successfully resolve the Content Engine URL from this entry, the entry must be pointing to a single active Content Engine node, not to a highly available URL. To achieve this you must temporarily reconfigure the `WcmApiConfig.properties` file to point to a single active Content Engine node when running the tool.

To run the RMCreationTool and RMAccessRole tool in an environment with a highly available Content Engine

- a. On your Records Manager server, create a backup version of the `WcmApiConfig.properties` file located in `RM_install_path/RM/RecordsManager/WEB-INF`.

Example:

Make a copy of the file, using the following filename:
`WcmApiConfig.properties.cluster`.

- b. Edit the `WcmApiConfig.properties` file.
 - i. Change the `RemoteserverURL` entry to point to the active Content Engine node using the following format:

WebSphere

```
comp:iiop://node name:WSI_port/FileNet/Engine
```

WebLogic

```
comp:t3://node name:WSI_port/FileNet/Engine
```

JBoss

```
comp:jnp://node_name:WSI_port/FileNet/Engine
```

For information about the `RemoteServerURL` value, see information about the Content Engine URL in [“Appendix B – Connecting to a highly available Content Engine” on page 168](#).

Example:

For WebSphere, change the value from

```
cemp:corbaloc::server1:port,:server2:port/cell/cluster/CECluster/FileNet/Engine
```

to

```
cemp:iiop://server1:port/FileNet/Engine
```

- c. Run the RMCreationTool or RMAccessRole tool on your Records Manager server.
Follow the instructions in the “Add Marking Sets and AddOns to the FileNet P8 domain” or “(DoD Classified only) Create the Classified Records Security Personnel Access Role” task of the *IBM FileNet Records Manager Installation and Upgrade Guide*.
 - d. After the tool completes successfully, reinstate the original `wcmApiConfig.properties` file that you backed up.
5. If a load-balancer or proxy server is used in this configuration make sure to use this URL for the remaining Records Manager installation steps that require an application URL.

Example:

If the application is deployed to a farm of application servers, and a load-balancer is configured for this farm then a URL for Application Engine could be:

```
http://proxy_url:proxyport/Workplace
```

Records Manager would have a URL of the form:

```
http://proxy_url:proxyport/RecordsManager
```

Use this URL when completing the installation steps for Records Manager.

NOTES

- You must perform installation and configuration steps that utilize the Records Manager or Application Engine/Workplace XT web applications on one node per HA environment only, and not for every node in the farm.
- You must perform the Transfer Workflow Definitions task from the command prompt using `WorkflowTransfer.bat/.sh`. You need to do this on one node per HA environment only, and not for every node in the farm.
- You must perform installation and configuration steps that utilize the Process Task Manager (to make configuration changes to the Component Manager) or the Records Manager Sweep applications on each node in the farm of application servers.

Deployment verification

1. Make sure Application Engine/Workplace XT and Records Manager applications are running on the cluster.
2. Open a web browser and type in the URL for the Records Manager application

NOTE If using a proxy server the URL has the form:

`http://proxy_server_name:port#/context_root`

Example:

`http://appserver-proxy.server.com:8080/RecordsManager`

3. Verify that the Sign-in page displays.
4. Fail one of the nodes.
5. Verify that you can reload the sign-in page.

Install a highly available IBM FileNet Connector for SharePoint Web Parts 2.2

CAUTION This topic covers installation of IBM FileNet Connector for SharePoint Web Parts 2.2 in a supported FileNet P8 HA environment. Do NOT install this functional expansion unless it is supported at the release levels of your FileNet P8 environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20.](#)

Install and configure IBM FileNet Connector for SharePoint Web Parts 2.2

Installation tasks

For general instructions on installing IBM FileNet Connector for SharePoint Document Libraries, see the *IBM FileNet Connector for SharePoint Web Parts Libraries Installation and Upgrade Guide, Version 2.2* (PDF) with modifications as listed below.

NOTE Before you begin the installation, ensure that all prerequisites are completed. For details, see “Installation planning considerations” in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide*.

To install IBM FileNet Connector for SharePoint Web Parts 2.2 in a farm environment

1. Install and configure a SharePoint farm.

To set up a farmed SharePoint environment, follow the instructions provided by Microsoft. For example, see the Microsoft TechNet article “Install Office SharePoint Server 2007 in a server farm environment.”

NOTE Ensure that your farm meets all requirements as outlined in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide*.

2. (IBM FileNet Credential Store Authentication only) Enable FileNet Credential Store Authentication

Only do this if you are using IBM FileNet Credential Store authentication with your installation.

- a. Note the following information for the remote MS SQL Server database required to set up the Microsoft SQL Server User Profile Database and configure IBM FileNet Credential Store Authentication.

- Database Server and Instance
- Database Port
- Database User
- Database Password

- b. (Optional) Manually create a Microsoft SQL Server User Profile Database

Follow the instructions in the “(Optional) Create a Microsoft SQL Server User Profile Database” task in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.2* with the following modifications:

NOTE This step is optional, you can use the IBM FileNet Connector for SharePoint Web Parts setup program to create the database, see [“Install IBM FileNet Connector for SharePoint Web Parts on each farm node.” on page 118.](#)

Make sure you enter the database server and instance name for the `<hostname>` property. For example:

`database_server/database_instance`

NOTE Enter only the database server name if you use the default database instance.

Note the following information required to configure IBM FileNet Credential Store Authentication below.

- Database Name

- 3. Install IBM FileNet Connector for SharePoint Web Parts on each farm node.

For each node, follow the instructions in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide* with the following modifications:

- (IBM FileNet Credential Store Authentication only) Configure the SharePoint farm to use the Microsoft SQL Server User Profile Database.
 - If you choose to have the setup program create the database for you, only do this when running the setup program on the first farm node. For the consecutive nodes select the option to use an existing database, and point to the database you created while running the setup program on the first node.
 - If you have created the database manually, select the option to use an existing database, and point to the database you created in [“\(Optional\) Manually create a Microsoft SQL Server User Profile Database” on page 118.](#)
 - For a remote database, make sure you enter the database server and instance name for the `<hostname>` property. For example:

`database_server/database_instance`

NOTE Enter only the database server name if you use the default database instance.
 - For port number enter the instance port number associated with your database instance.
 - Encryption Key

CAUTION Each SharePoint node must use the same encryption key.
- Choose a deployment method that suits your SharePoint setup.

NOTE In a farm environment the auto-deployment will fail, and an error will display in a splash screen at the end of the installer. You must manually deploy the components after installation. For information, see [“Manually Deploy Components” on page 120.](#)

- For each node after the first, make sure the information you enter for “Specify Web Part Properties” is identical to what you entered for the first node.

NOTE If your FileNet P8 environment is highly available, make sure you point to the virtual names of the highly available Content Engine, Application Engine, and Process Engine installations.

For example, for Content Engine use:

```
http://CE_farm_name:load_balancer_port#/wsi/FNCEWS40MTOM/
```

Where:

CE_farm_name is the virtual name of your Content Engine farm.

load_balancer_port# is the port for the load balancer for your Content Engine farm.

4. Verify that the installation was successful.

Check the install log at:

```
C:\Program Files\IBM\FileNetConnectorForSharePointWebParts\  
sharepoint_web_parts_install_log_2.2.0.txt
```

NOTE The log will contain errors due to the solution(s) not being deployed. This is normal in the case of a farm.

Example of error:

```
EXECUTE_DEPLOYWEBPARTS: -2147024894. Due to error certain actions may need to be  
done by hand
```

```
Status: ERROR
```

```
Additional Notes: ERROR - Batch File error deploying WebParts may need to be done  
by hand
```

Make sure there are no other errors besides the example above.

5. Verify that the encryption key specified in the `web.config` file for each deployed web application is identical for all nodes.

- a. Open the `web.config` file for each web application on each farm node.

The file is typically located in `C:\Inetpub\wwwroot\wss\VirtualDirectories\port#`

Where `port#` is unique for each installation.

- b. Verify that the entry for `<machineKey>` is identical for each web applications on all nodes.

NOTE If you installed the component using identical input these entries should not differ. If they do differ, select one of the entries and copy it to the `web.config` files on all the other nodes.

6. Manually Deploy Components

For each web application, follow the instructions in the “Manually Deploy Components” task *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide* with the following modifications:

- For a SharePoint farm, use the virtual host name of the farm instead of “localhost.”

CAUTION Performing these procedures will only be successful if you have already run the IBM FileNet Connector for SharePoint Web Parts 2.2 installer on each SharePoint server in the SharePoint farm.

7. (IBM FileNet Credential Store Authentication only) Specify Application Pool Security Account

For each farm node, follow the instructions in the “Specify Application Pool Security Account” task *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide*.

8. Add and Configure IBM FileNet Connector for SharePoint Web Parts

For each applicable SharePoint Web application server, follow the instructions in the “Add and Configure IBM FileNet Web Parts” task *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide* with the following modifications:

- Use the virtual URLs for CE, PE and AE.

Installation verification

1. Make sure SharePoint and FileNet P8 are running.

- a. Verify that FileNet P8 is running.

Verify that you can log in to Workplace or Enterprise Manager.

- b. Verify that SharePoint can connect to FileNet P8.

Log in to SharePoint and verify that the Web Parts can connect to IBM FileNet P8, and that you can browse your object store using the Browse Web Part.

Add a document to the object store using SharePoint.

- c. IBM FileNet Connector for SharePoint Web Parts

Log in to Workplace or Enterprise Manager and verify that you can see the document you just added via the SharePoint Web Part.

2. Verify that IBM FileNet Connector for SharePoint Web Parts connects after Content Engine failover.

- a. Fail over your highly-available Content Engine.

- b. Verify that SharePoint can connect to FileNet P8.

Log in to SharePoint and verify that Web Parts connects to FileNet P8, and that you can browse, view, and open documents in your object store through SharePoint Web Parts.

3. Verify that IBM FileNet Connector for SharePoint Web Parts connects after Application Engine failover.
 - a. Fail over your highly-available Application Engine.
 - b. Verify that SharePoint can connect to FileNet P8.

Log in to SharePoint and verify that Web Parts connects to FileNet P8, and that you can open documents in your object store through SharePoint Web Parts.
4. Verify that SharePoint clients have access after a SharePoint failover.
 - a. Fail over your highly available SharePoint server.
 - b. Verify that the SharePoint client can still connect to SharePoint and that SharePoint can connect to FileNet P8.

Install a highly available IBM FileNet Connector for SharePoint Document Libraries 2.2

CAUTION This topic covers installation of IBM FileNet Connector for SharePoint Document Libraries 2.2 in a supported FileNet P8 HA environment. Do NOT install this functional expansion unless it is supported at the release levels of your FileNet P8 environment. For information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Install and configure IBM FileNet Connector for SharePoint Document Libraries 2.2

NOTE For information on how to install IBM FileNet Connector for SharePoint Web Parts in a highly available environment, see [“Install a highly available IBM FileNet Connector for SharePoint Web Parts 2.2” on page 117](#).

Installation tasks

For general instructions on installing IBM FileNet Connector for SharePoint Document Libraries, see the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* (PDF) with modifications as listed below.

NOTE Before you begin the installation, ensure that all prerequisites are completed. For details, see “Installation Planning Considerations” in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide*.

To install IBM FileNet Connector for SharePoint Document Libraries 2.2 in a highly available environment

1. Install and configure a SharePoint farm.

To set up a farmed SharePoint environment, follow the instructions provided by Microsoft. For example, see the Microsoft TechNet article “Install Office SharePoint Server 2007 in a server farm environment.”

NOTE Ensure that your farm meets all requirements as outlined in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide*.

2. Install and configure a Microsoft active/passive cluster.

This cluster will host the Doc Lib Connector Remote Admin Services. To set up a Microsoft active/passive cluster, follow the instructions provided by Microsoft.

3. (IBM FileNet Credential Store Authentication only) Enable FileNet Credential Store Authentication

Only do this if you are using IBM FileNet Credential Store authentication with your installation.

- a. Note the following information for the remote MS SQL Server database required to set up the Microsoft SQL Server User Profile Database and configure IBM FileNet Credential Store Authentication.

- Database Server and Instance
- Database Port
- Database User
- Database Password

- b. (Optional) Manually create a Microsoft SQL Server User Profile Database

Follow the instructions in the “(Optional) Create a Microsoft SQL Server User Profile Database” task in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* with the following modifications:

NOTE This step is optional, you can use the IBM FileNet Connector for SharePoint Document Libraries setup program to create the database, see [“Install the Connector Services on the SharePoint Farm nodes.” on page 123.](#)

Make sure you enter the database server and instance name for the `<hostname>` property. For example:

```
database_server/database_instance
```

NOTE Enter only the database server name if you use the default database instance.

Note the following information required to configure IBM FileNet Credential Store Authentication below.

- Database Name

4. Install the Connector Services on the SharePoint Farm nodes.

For each node in the SharePoint farm, follow the instructions in the “Install IBM FileNet Connector for SharePoint Document Libraries” task in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* with the following modifications:

- Choose “Connector Services” only.

NOTE Deselect the option to install the Connector Administration Services; these will be installed on the active/passive cluster.

- (IBM FileNet Credential Store Authentication only) Configure the SharePoint farm to use the Microsoft SQL Server User Profile Database.

- If you choose to have the setup program create the database for you, only do this when running the setup program on the first farm node. For the consecutive nodes select the option to use an existing database, and point to the database you created while running the setup program on the first node.

- If you have created the database manually, select the option to use an existing database, and point to the database you created in “(Optional) Manually create a Microsoft SQL Server User Profile Database” on page 123.
- For a remote database, make sure you enter the database server and instance name for the <hostname> property. For example:

database_server/database_instance

NOTE Enter only the database server name if you use the default database instance.

- For port number enter the instance port number associated with your database instance.
- Verify that no installation errors occurred.

View the install log at:

```
Drive:\Program Files\IBM\FileNetConnectorForSharePointDocumentLibraries\  
IBM_FileNet_Connector_for_SharePoint_Document_Libraries_InstallLog_2.2.0  
.0.txt
```

NOTE The install log will contain an error stating that the solution did not deploy; this is normal for a farm deployment. Error example:

```
EXECUTE_DEPLOYREDIR: -2147024894. Due to deploying Redirector error certain  
actions may need to be done by hand
```

```
Status: ERROR
```

```
Additional Notes: ERROR - Batch File error deploying Redirector may need to be  
done by hand
```

5. (IBM FileNet Credential Store Authentication only) Specify Application Pool Security Account

Perform this task on the remote database server. Follow the instructions in the “Specify Application Pool Security Account (FileNet Credential Store Authentication only)” task in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2*.

6. Install the Connector Administration Services remotely on each cluster node.

For each cluster node in the Microsoft active/passive cluster you created in [Step 2 “Install and configure a Microsoft active/passive cluster.” on page 122](#), follow the instructions in the of “(optional) Install the Connector Administration Services on a Remote Server” task in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* with the following modifications:

- Change the drive letter on the default install folder to your active/passive cluster shared drive.
- On the cluster server, choose "Connector Administration Services." The Web Services are already installed on the SharePoint farm.

NOTE To correctly configure the Connector Administration Services you must run the setup program on the currently active cluster node. Before you install on the next cluster node you must fail over your cluster group to that node, and then repeat the installation procedure with identical input.

7. Create Connector Administration Services Database and Data Store

Follow the instructions in the “Create Connector Administration Services Database and Data Store” task in the *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2*. with the following modifications:

- Copy the appropriate database script from the active node of the Document Libraries Connector Remote Admin Services active/passive cluster to the database server where you will install the database.

Default location on the cluster node:

```
Drive#:\Program Files\IBM\UFI\Database  
Scripts\db_ufi_400_database_type.sql
```

- Launch UFI Configuration Manager to create a new data store from the active server in the active/passive cluster.

8. Configure the Web Service on each node in the SharePoint farm.

For each node in the SharePoint farm, follow the instructions in the “Set IIS Manager Account Properties” task in the *IBM FileNet Connector for SharePoint Doc Lib Installation and Upgrade Guide, Version 2.2*.

9. Add Content Engine Object Store Mapping Keys to the web.config Files.

For each node in the SharePoint farm, follow the instructions in the “Add Content Engine Object Store Mapping Keys to the web.config Files” task in the *IBM FileNet Connector for SharePoint Doc Lib Installation and Upgrade Guide, Version 2.2* with the following modifications:

NOTE If your FileNet P8 environment is highly available, make sure you point to the virtual names of the CE farm.

For example, use:

```
http://CE_farm_name:load_balancer_port#/wsi/FNCEWS40MTOM/
```

Where:

CE_farm_name is the virtual name of your CE farm.

load_balancer_port# is the port for the load balancer for your CE farm.

10. Deploy the IBM FileNet User Administration Web Part

On one node in the SharePoint farm, follow the instructions in the “Deploy the IBM FileNet User Administration Web Part” task in the *IBM FileNet Connector for SharePoint Doc Lib Installation and Upgrade Guide, Version 2.2* with the following modifications:

- For a SharePoint farm, use the virtual host name of the farm instead of “localhost.”

11. Manually deploy the IBM FileNet Redirector on each node in the SharePoint farm.

On one node in the SharePoint farm, follow the instructions in the “Deploy the IBM FileNet Redirector” task in the *IBM FileNet Connector for SharePoint Doc Lib Installation and Upgrade Guide, Version 2.2* with the following modifications:

- For a SharePoint farm, use the virtual host name of the farm instead of “localhost.”
12. Configure the Remote Admin Server Cluster for fail over.
 - a. Make sure the following services are set to manual startup on both servers in the active/passive cluster:
 - FileNet SharePoint Connector Service
 - UFI FileNet P8 Content Engine 4.0 Connector
 - UFI Task Routing Engine
 - UFI Utility Service
 - b. Open Cluster Administrator for the Microsoft active/passive cluster.
 - c. Right-click your group name and choose New > Resource.
 - d. In the New Resource Box, give the new resource a name. For example: “Doc Lib Connector.”
 - e. Under Resource Type, choose “Generic Service.”
 - f. Click **Next**.
 - g. Make sure both servers are listed under “Possible Owners.” Click **Next**.
 - h. For Dependencies, move the drive that you installed the Doc Lib Connector Remote Admin Services on from “available resources” to “resource dependencies.”
 - i. Under Service Name, enter:
`ibm.ctms.taskrouteservice`
 - j. Leave start parameters blank. Click **Next**.
 - k. On the registry replication page, click **Add** and enter:
`SOFTWARE\IBM\EMRC\4.0`
 - l. Click **OK**, then **Finish**.
 - m. Note that the new resource is created off-line.

Installation verification

For more information on how to use the Configuration Manager for the tasks described below, see *IBM FileNet P8 Help* topic Expansion Products > IBM FileNet Connector for SharePoint > IBM FileNet Connector for SharePoint Document Libraries.

1. Configure your test environment.
 - a. Make sure SharePoint and IBM FileNet P8 are running.
 - b. Use the Configuration Manager to create a task route.
 - Configure the task route to collect documents from a verification folder.

- On the schedule tab, configure the task to run once, and that the end date is set to “when task completes.”
 - Make sure post-processing states “move and delete.”
- c. Make sure the proper target folder exists in the selected object store (create it using Enterprise Manager or Workplace) and that it is empty.
2. Run verification tests on the environment

There are four functional tests available to verify that your highly available IBM FileNet SharePoint Connector for Document Libraries is working correctly.

- [“To verify the base document transfer capabilities.” on page 127.](#)
- [“To verify failover of the Document Libraries Connector Remote Admin Services on the active/passive cluster.” on page 127.](#)
- [“To verify the Document Libraries Connector Services on the SharePoint server farm nodes” on page 128.](#)
- [“To verify document transfer after Content Engine failover” on page 128.](#)

To verify your IBM FileNet Connector for SharePoint Document Libraries installation

- To verify the base document transfer capabilities.
 - i. Upload a small number of documents to the verification folder referenced in the task route.
 - ii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.
 - iii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)
- To verify failover of the Document Libraries Connector Remote Admin Services on the active/passive cluster.
 - i. Upload a small number of documents to the verification folder referenced in the task route.
 - ii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.
 - iii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)
 - iv. Leave the services running.
 - v. Upload a small number of documents to the verification folder referenced in the task route.
 - vi. Fail the cluster from the active node to the passive node.

vii. When the cluster starts on the new node, it will start the doc lib resource.

NOTE The task route you created above is configured to run once. Restarting the doc lib service on the new server will cause the task to start again, and run to completion.

viii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)

- To verify the Document Libraries Connector Services on the SharePoint server farm nodes

i. Upload a small number of documents to the verification folder referenced in the task route.

ii. Configure your load balancer to point to the first node of your SharePoint server farm.

iii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.

iv. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)

v. Stop the Task Routing Engine service.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and take it off-line.

vi. Upload a small number of documents to the verification folder referenced in the task route.

vii. Configure your load balancer to point to the second (or consecutive) node of your SharePoint server farm.

viii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.

ix. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)

- To verify document transfer after Content Engine failover

i. Upload a small number of documents to the verification folder referenced in the task route.

ii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.

iii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)

iv. Stop the Task Routing Engine service.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and take it off-line.

- v. Fail over your Content Engine.
- vi. Upload a small number of documents to the verification folder referenced in the task route.
- vii. Start the Task Routing Engine service to run your task route.

From the Cluster Administrator, right click on the doc lib resource you created in [Step 12 “Configure the Remote Admin Server Cluster for fail over.” on page 126](#) and bring it on-line.

- viii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 129.](#)

To verify transfer of documents from SharePoint to IBM fileNet P8

1. Verify that SharePoint Connector for Document Libraries is running.
 - On the active node of the cluster, from Windows Task Manager, verify that the following services are running:
 - FileNet SharePoint Connector Service
 - UFI FileNet P8 Content Engine 4.0 Connector
 - UFI Task Routing Engine
 - UFI Utility Service
 - From Enterprise Manager or Workplace, verify that documents are being added to the folder specified by the File in Folder task.
 - From SharePoint, verify that documents are being migrated and deleted from the verification folder referenced in the task route.
2. Verify that the file transfer completed successfully:
 - From Enterprise Manager or Workplace, verify that all documents were added to the folder specified by the File in Folder task.
 - From SharePoint, verify that documents were migrated and deleted from the verification folder referenced in the task route.

Plan and prepare for high availability upgrades

This section, Plan and Prepare for High availability upgrades, provides instructions for preparing your system to upgrade the core IBM FileNet P8 components in a highly available environment. Installation instructions themselves are contained in the IBM FileNet P8 Platform Installation and Upgrade Guide, and separate add on installation guides.

Plan the upgrade

This section describes how to prepare your environment for the installation of a highly available 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 independent software vendor (ISV) software.

Plan your upgrade

Content Engine

- The 4.5 (4.0 and later) release of Content Engine does not use separate object store and file store services, instead a Content Engine consists solely of the Content Engine application running on a Java application server.
- The 4.5 Content Engine can reside on the same servers as the 3.5.x Content Engine or on separate servers, running independently of each other.
- Once the object stores and file stores have been converted the 4.5 Content Engine will take ownership of these, and they will no longer be available to the 3.5.x Content Engine.

Content Engine upgrade procedure overview

The 4.5 version of Content Engine is an enterprise application running on a Java application server. Content Engine can be made highly available using the processes for making any application highly available in a Java application server, as documented in [“Install a highly available Content Engine” on page 70](#).

NOTE The clustered services that comprised a highly available 3.5.x Content Engine will be disabled during the upgrade and can be left offline or brought online after the upgrade as needed.

The Content Engine upgrade process in a highly available environment is as follows:

1. Prepare Content Engine for Upgrade.

Gathering information needed for performing the upgrade and installing and configuring Content Engine in a highly available environment.

- a. Gather information needed for the upgrade.
- b. Create resources needed for the new Content Engine software.
- c. Install FileNet P8 4.5 Content Engine in a highly available environment.

CAUTION When deploying Content Engine in an HA environment you must specify the cluster as the deployment target.

2. Upgrade Content Engine Items.

- a. Migrate and upgrade GCD information from the sysinit file to the FileNet P8 4.5 GCD database repository.
- b. Migrate and upgrade object stores and file stores to the FileNet P8 4.5 Content Engine.

The 3.5.x/4.5 Content Engine farm/cluster after upgrade

After the 4.5 upgrade is performed you can bring the Content Engine farm/cluster resources online if needed. The 3.5.x Content Engine and 4.5 Content Engine can operate independently and will not conflict with each other. Once the upgrade process is complete the 3.5.x Content Engine cluster resources (if cluster is being used) can be uninstalled as needed, assuming that the 3.5.x environment is no longer needed.

Object store availability in a mixed 3.5.x/4.5 environment

Object stores will only be visible and accessible in the environment for which they are configured. Object stores that have not yet been upgraded to 4.5 will still be accessible from the 3.5.x environment. When a 3.5.x object store is upgraded to 4.0.x it will no longer show up in the 3.5.x environment and will only be accessible from the 4.5 environment.

Process Engine

Process Engine upgrade procedure overview

Server farm environments

Process Engine farming was initially supported with the FileNet P8 4.0.0 Platform GA release. The instructions in this technical notice only cover migration of a PE cluster to a PE farm environment, and not upgrades of existing PE server farms.

The following high-level steps are necessary to upgrade a Process Engine cluster to a Process Engine farm:

1. Move all cluster services to the first node.
2. Upgrade the first node to PE 4.5.
3. Uninstall PE 3.5.x from the second node in the cluster.
4. Install PE 4.5 on the second node using a unique NCH name.

Cluster environments

To upgrade Process Engine you must upgrade the software on all nodes in the cluster, and upgrade the Process Engine database objects once only from the first node upgraded.

The following high-level steps are necessary to upgrade the Process Engine for high availability:

1. Stop all Process Routers in the FileNet P8 environment.
2. Stop all Process Engine services in the cluster.
3. Disable or remove any cluster resources monitoring and controlling the Process Engine.
4. Upgrade the Process Engine software on one node.
5. Enable or recreate cluster resources to monitor and control the Process Engine.
6. Start Process Engine services.
7. Complete Post-Upgrade Process Engine Configuration steps
8. Upgrade the software on the remaining nodes.

NOTE There is no need to complete the post-upgrade steps on any node except the first node that was upgraded.

Application Engine/Workplace XT

Application Engine/Workplace XT upgrade procedure overview

An upgrade of an Application Engine cluster/server farm includes the following steps:

1. Stop all Process Routers in the FileNet P8 environment.
2. Undeploy the current Application Engine.
3. Upgrade the Application Engine software.
4. Deploy Workplace.
5. Configure Application Engine.
6. Restore the original bootstrap settings.

Records Manager

Rendition Engine

Information for upgrading the Rendition Engine from the 4.0 release to the 4.5 release is covered in this topic. The upgrade involves uninstalling and reinstalling the software.

- You will reinstall Rendition Engine on the existing Rendition Engine servers.
- Rendition Engine 4.5 and later does not use the publishing plugin. This component is uninstalled as part of the CE upgrade.

Rendition Engine upgrade procedure overview

The following high-level steps are necessary to upgrade and configure a highly available Rendition Engine for high availability:

1. Upgrade Content Engine, Process Engine, and Application Engine.
2. Perform prerequisite tasks as outlined in [“Perform required preparation tasks” on page 135](#).
3. Uninstall previous version of Rendition Engine
4. Upgrade and configure an initial Rendition Engine.
5. Upgrade and configure additional Rendition Engines.
6. (Optional) Verify Installation.

IBM FileNet Connector for SharePoint Document Libraries

IBM FileNet P8 platform, version 4.5, supports the same version of IBM FileNet Connector for SharePoint Document Libraries as IBM FileNet P8 platform, version 4.0. No changes to the highly available SharePoint environment are required during an upgrade of IBM FileNet P8 platform.

IBM FileNet Connector for SharePoint Web Parts

IBM FileNet P8 platform, version 4.5, supports the same versions of IBM FileNet Connector for SharePoint Web Parts as IBM fileNet P8 platform, version 4.0. No changes to the highly available SharePoint environment are required during an upgrade of IBM FileNet P8 platform.

Perform required preparation tasks

ATTENTION In addition to the preparation tasks listed below, follow the plan and prepare instructions in the *IBM FileNet P8 Platform Plan and Prepare Your Environment for IBM FileNet P8* document for the components you are installing. For add-on components, follow the plan and prepare instructions in the corresponding installation guide. Note that some of the tasks listed below might replace the tasks described for a non highly available environment.

To prepare your environment for a highly available installation

Perform the required preparation tasks for the components you plan to install:

- [“Installation administrator tasks” on page 136](#)
- [“Application server administrator tasks” on page 137.](#)
- [“IT administrator tasks” on page 138](#)
- [“Security administrator tasks” on page 140](#)
- [“Database administrator tasks” on page 141](#)

In addition, if you are using the Installation and Upgrade Worksheet for your environment, all administrators should fill out the worksheet with the relevant information as described in [“Appendix F - Installation and upgrade worksheet” on page 181.](#)

Installation administrator tasks

The items below are the required tasks for the Application Installer Administrator.

Content Engine

In addition to the tasks listed below, review and perform the required Content Engine tasks in the “Perform the required preparation tasks” topic of the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Process Engine

In addition to the tasks listed below, review and perform the required Process Engine tasks in the “Perform the required preparation tasks” topic of the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Application Engine/Workplace XT

In addition to the tasks listed below, review and perform the required Application Engine or Workplace XT tasks in the “Perform the required preparation tasks” topic of the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide.

Rendition Engine

In addition to the tasks listed below, review and perform the required Rendition Engine tasks in the “Perform the required preparation tasks” section of the “Plan and Prepare for upgrade” topic in the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

Application server administrator tasks

The items below are the required tasks for the Application Installer Administrator.

Perform all required application server administrator tasks for your environment, as listed in the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide, with the following HA specific details:

Content Engine

Configure the application server for Content Engine (upgrades from 3.5.x)

Perform the applicable tasks for your setup. For pre-installation, verification, and configuration steps needed for a highly available Content Engine, see [“Install a highly available Content Engine” on page 70](#). The steps listed are needed both for upgrades and new installations.

Application Engine/Workplace XT

Perform the task “Prepare for Application Engine upgrade” of the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following modifications.

- Backup the shared bootstrap.properties file in the shared network location.
- (AE 3.5 only) A shared configuration directory will be created beneath the original location of the bootstrap file.

Example: If your 3.5.x shared bootstrap location was Z: the newly created shared configuration directory will be Z:\AE\.

IT administrator tasks

The items below are the required tasks for the Information Technology (IT) administrator.

Process Engine

Configure the Process Engine nodes hosts file

Server Farm

For each Process Engine server, configure the hosts file so that on each node in the farm it:

- Includes the NCH name for every other node in the farm.
- Identifies the loadbalancer virtual as itself.

The hosts file entries should look like this:

	Hosts file entry
<i>PE_farm1</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name loadbalancer_virtual_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name</i>
<i>PE_farm2</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name loadbalancer_virtual_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name</i>
<i>PE_farm3</i>	1.2.3.4 <i>PE_farm1_server PE_farm1_nch_name</i> 1.2.3.5 <i>PE_farm2_server PE_farm2_nch_name</i> 1.2.3.6 <i>PE_farm3_server PE_farm3_nch_name loadbalancer_virtual_name</i>

Example:

loadbalancer_virtual_name (1.2.3.10) is the load balancer virtual name, and *PE_farm1* (1.2.3.4), *PE_farm2* (1.2.3.5), and *PE_farm3* (1.2.3.6) are three PE servers.

The format of *PE_farm1_nch_name* is:

```
PEfarm1-filenet-nch-server
```

The NCH name in the hosts file is the NCH name with -filenet-nch-server appended, and any dashes or underscores removed.

CAUTION (UNIX) Verify that the nsswitch.conf file (AIX netsvc.conf and/or irs.conf) is configured to resolve names via the local hosts file first.

Server Cluster

Configure the hosts file on each cluster node server to follow the following format:

```
PEvirtual_ip PEvirtual_dns_name PEclustered_nch_name local_node_name
```

Server	Hosts file entry
<i>PE_node</i>	<i>1.2.3.4 virtual_dns_server clustered_nch_server local_node_server</i>

Example

```
1.2.3.4 virtual_dns_server.ibm.com clustered_nch_server local_node_server
```

The format of PEclustered_nch_server is:

```
PEvirtual_dns_name-filenet-nch-server
```

The NCH name in the hosts file is the NCH name with -filenet-nch-server appended, and any dashes or underscores removed.

Security administrator tasks

The items below are the required tasks for the Security Administrator.

Perform all required security administrator tasks for your environment, as listed in the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide, with the following HA specific details:

Content Engine

Specify IBM FileNet P8 Accounts for upgrades

NOTE This step is required for upgrades from CE 3.5 only.

Records Manager

(UNIX) Verify write permission on the /bin directory

Verify that the `/bin` directory is write-enabled before running the RM installer. The directory need only be write-enabled during installation, and can be set back to whatever security permission you choose after the installation completes.

Database administrator tasks

The items below are the required tasks for the database administrator.

Perform all required database administrator tasks for your environment, as listed in the “Plan and prepare for IBM FileNet P8 upgrade” section of the *Plan and Prepare Your Environment for IBM FileNet P8* guide, with the following HA specific details:

Content Engine

Create the GCD data source for Content Engine (upgrades from 3.5.x)

For more information about making databases highly available with P8 components, see [“Database services” on page 34](#).

Upgrade and configure for high availability

This section provides the information needed to upgrade an existing FileNet P8 3.5.x or 4.x highly available environment to FileNet P8 4.5.

Depending on your configuration, you might have to install at least a portion of the additional service pack and interim fix software while you work your way through the standard IBM FileNet P8 platform upgrade.

To upgrade IBM FileNet P8 components in a highly available environment

NOTE For an HA environment FileNet recommends that you perform the required installations/upgrades across the whole FileNet P8 environment at the same time, bringing the whole environment down before starting the installation/upgrade, and bringing it back up once all installation/upgrade tasks are completed.

To upgrade an IBM FileNet P8 platform in an HA environment, follow the instructions found in the *Plan and Prepare Your Environment for IBM FileNet P8* document, and in the “Upgrade and configure IBM FileNet P8 Platform” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

1. Upgrade the IBM FileNet P8 platform documentation.
2. Ensure that you are running the minimum required release of IBM FileNet P8 platform software and service packs to perform an upgrade.
3. Upgrade the highly available Content Engine farm or cluster.
For more information, see [“Upgrade a highly available Content Engine” on page 143](#).
4. Upgrade the highly available Process Engine farm or cluster.
For more information, see [“Upgrade a highly available Process Engine” on page 146](#).
5. Upgrade the highly available Application Engine farm or cluster.
For more information, see [“Upgrade a highly available Application Engine/Workplace XT” on page 158](#).
6. Upgrade any other add-on components per your environment.
 - [“Upgrade a highly available Rendition Engine” on page 160](#).
 - [“Upgrade a highly available CFS-ICI” on page 161](#).
 - [“Upgrade a highly available Records Manager” on page 162](#).

Upgrade a highly available Content Engine

This section covers the tasks and requirements for upgrading a highly available Content Engine in a FileNet P8 environment.

Content Engine upgrade

CAUTION If you install version 4.5 of FileNet Enterprise Manager on a same server where a 3.5.x version of FileNet Enterprise Manager is installed you will no longer be able to use the 3.5.x version of the FileNet Enterprise Manager. Install the new version of FileNet Enterprise Manager only if you will no longer need access to the FileNet 3.5.x P8 domain from this server.

To prepare Content Engine for upgrade

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 135](#).
2. Prepare Content Engine for Upgrade involves performing the following tasks as described in [“Complete pre-upgrade Content Engine configuration”](#) from the Upgrade Content Engine task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

“To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain (3.5.x only).”

In a cluster environment take the windows services for Content Engine offline, this includes the clustered resources for:

- Apache
- FileNet Object Store Service
- FileNet File Store Service
- Process Services Manager

CAUTION You must leave the GCD share and virtual network resources (network name and IP) online. The upgrader requires access to the GCD sysinit file. If the share is offline it will not be accessible to the upgrader.

In a highly available environment where file store shared folders are highly available make sure they are kept online. The upgrader requires access to the shared folders using the same shared path that is stored in the GCD for the file store folder.

To upgrade Content Engine

1. Depending on your environment, follow either of the following two procedures “Upgrade Content Engine software from 3.5.x” or “Upgrade Content Engine software from 4.0.x” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - Installing the Upgrader Tool.

Install the Upgrader tool on at least one Windows server by running the Content Engine installer on a Windows machine, and selecting the Tools option in the Choose Components screen.

 - The Windows system does not need to be a part of a Windows Active Directory Domain.
 - The upgrader tool does not have to be collocated with Content Engine 3.5.x or 4.0.x to be installed and run on a server.
2. Configure Content Engine instances upgraded from 4.0.x

Perform “Configure Content Engine instances upgraded from 4.0.x” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

 - Object Stores.

You can create object stores at this point, but make sure no object stores are created with the same name as an object store that will be upgraded. If an object store already exists and the upgrade tool is run to upgrade an object store of the same name the upgrade will fail for that object store.

To Deploy upgraded Content Engine instances

CAUTION When deploying Content Engine in an HA environment you must specify the cluster as the deployment target.

To upgrade Content Engine data

Perform the “Upgrade Content Engine Data” task from the *IBM FileNet P8 Platform Installation and Upgrade Guide* from a single node where you installed the Upgrader Tool.

- Upgrading using the command line tool.

Perform the tasks in the “Command Line Interface to the Upgrader Tool” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide* from a single node where you installed the Object Store 3.5 to 4.0 Upgrade Tool.

 - “To create an XML upgrade status file.”
 - “To run the Upgrader tool using the CLI.”

- Upgrading using the Graphical User Interface

Run the upgrade tool and perform the tasks in the “Upgrading data using the graphical user interface” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide* from a single node where you installed the Upgrader Tool.

- “To run the Upgrader tool.”
- “To create the XML upgrade status file.”
- “To configure authentication, FCD, and object store parameters.”

CAUTION You must specify a URI to a single node instance as the application will make several updates to the GCD in short succession. During the upgrade the upgrader must not be load-balanced to another node in the highly available Content Engine configuration.

Upgrade a highly available Process Engine

This section covers the tasks and requirements for upgrading a highly available Process Engine in a FileNet P8 environment.

Migrating Process Engine 3.5.x clusters to Process Engine 4.5 farms

Process Engine 4.5 supports highly available server farms. If you are currently running Process Engine in a cluster environment, you might migrate to a server farm environment. The following procedure documents a migration of a two-node cluster to a two server farm.

To migrate a Process Engine cluster to a Process Engine farm

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 135](#).
2. Upgrade Process Engine on the first node.
 - a. Using the cluster admin tools, take the PE Resource offline.
 - b. Using the cluster admin tools, remove the PE Resource.
 - c. (Optional, UNIX only) If you plan to re-use the node's IP address for the load balancer, using the cluster admin tools, remove the cluster IP resource.
 - d. (VCS only) From the cluster admin tools, freeze the existing PE cluster on one the first node.
 - e. Upgrade Process Engine on the first node to version 4.5.
 - i. Upgrade the first Process Engine in the cluster as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Upgrade and configure Process Engine.”

CAUTION During this procedure make sure the shared configuration data partitions remain mounted at all times.

 - /fnsw/local
 - fn_sec_db0
 - fn_sec_rl0
 - Drive:\FNSW_LOC (Windows)
 - ii. (Windows) Reboot the server after you have successfully installed PE.
 - iii. (Windows) After reboot move the cluster resources back to the first node.

On Windows the cluster will move to the next node during the reboot.
 - iv. (Windows) Remove the NCHBroadcast registry key on the first node only.
 1. Open Windows registry

2. Navigate to:

```
HKEY_LOCAL_MACHINE\Software\FileNET\IMS\CurrentVersion
```

3. Remove the NCHBroadcast registry key.

f. (UNIX) Reset IP address and NCH name.

At this point, the upgraded PE 4.5 will be still be listening to the virtual IP and NCH (Network clearinghouse) name.

From Process Task Manager, do the following:

- Change the PE IP address from the virtual IP to the local IP address of the server.
- Verify that the IP and hostname in Process Task Manager (vtaskman) match and are correct.
- Verify that the DNS entry for the NCH alias is correct.

3. Verify that there is only one entry in the vwserver database table for the upgrade system.

Oracle

a. Login to the PE database as SYS and run the following command:

```
select * from f_sw.vwserver;
```

b. Verify that the name of the server you just upgraded is the only entry, and that it only appears once in this table.

SQL Server

a. Launch Microsoft SQL Server Enterprise Manager and navigate to the vwserver table:

```
PE_database > Tables > vwserver
```

b. Verify that the name of the server you just upgraded is the only entry, and that it only appears once in this table.

4. Uninstall Process Engine 3.5.x from the next node in the farm as described in the *IBM FileNet P8 platform 3.5.x Installation and Upgrade Guide* task “Remove Process Engine (platform).”

NOTE After uninstalling, verify that everything got removed and then reboot the next node server.

5. (Windows) Evict the second node from the cluster.

- a. On the first node, open Microsoft Cluster Administrator.
- b. Select the second node, right-click **Evict Node**.

6. Install Process Engine 4.5 on the second node.

Install and configure the second Process Engine in the farm as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Install and Configure Process Engine “ with the following additional comments:

- a. (Windows only) Log in to the second node server as a local administrator, not as a Windows domain user.

- b. (UNIX only) Create local resources for:
 - /fnsw/local
 - fn_sec_db0
 - fn_sec_rl0
 - c. During the installation, do the following:
 - i. When prompted, enter a unique NCH name set in [“\(UNIX\) Reset IP address and NCH name.” on page 147.](#)

NOTE Use a unique value different from the NCH name used by the first node.
 - ii. Configure all additional Process Engines to use the same database as the initial Process Engine:

CAUTION Make sure your database settings are configured identical to what you initially set up during your FileNet PE 3.5.x installation.
 - iii. When running SQL scripts choose the following option: “I have already run the pre-install scripts manually.”
 - iv. In the “Complete post-upgrade Process Engine configuration“ do the following:

UNIX

No configuration required.

Windows

Perform “To configure contiguous free memory for Process Engine (Windows only).”
 - d. Start Process Engine on the second node.
7. Verify that both servers exist in the vwserver database table.

Oracle

 - a. Login to the PE database as SYS and run the following command:


```
select * from f_sw.vwserver;
```
 - b. Verify that the name of both server one and server two appear in this table.

SQL Server

 - a. Launch Microsoft SQL Server Enterprise Manager and navigate to the vwserver table:


```
PE_database > Tables > vwserver
```
 - b. Verify that the name of both server one and server two appear in this table.
 8. Configure the hardware load balancer.
 - a. Verify that your hardware load balancer is set up to allow ALL PROTOCOLs traffic, not just TCP.

- b. Verify that the load balancer virtual name is in the host file of the local PE server and that it resolves to that PE server IP address.

(UNIX) To enable the vworbbroker process to start, the name resolution must come from the host file first.

- c. Configure the Process Task Manager for PE farm

NOTE This configuration is only required for the first PE. These are steps that should be performed in connection with the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure Process Task Manager.”

- i. Using Process Task Manager (vwtaskman), on the Servers node, add the hardware load balancer virtual name registered in DNS.
 - ii. On the Process Engine node, select the **Advanced** tab and add the following property:
vworbbroker.endPoint = giop:tcp:loadbalancer_virtual_name:32777
- d. Configure the Isolated Region connection points on the CE. All connection points for the farm should use the load balancer virtual name as the DNS name of the PE server.

Process Engine pre-upgrade tasks

Perform the following steps before upgrading your highly available Process Engine environment.

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 135](#).
2. (Process Engine 3.5.x upgrade) Stop all Process Routers in the FileNet P8 environment.
3. Stop the clustered Process Engine resource.
4. Delete the FileNet Process Engine resource from the cluster.
5. Perform all other pre-upgrade tasks listed in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, “Complete pre-upgrade Process Engine configuration” task with the following additional comments:

Process Engine 3.5.x upgrade

- “To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x.”
Complete once per environment.
- “To verify that all Process Analyzer 3.5.x events have been published.”
Complete once per environment.
- “To purge event logs and statistics records.”
Complete once per environment.
- “To stop all Process Engine-related services and applications.”
Complete this on each node.

Process Engine 4.0.x upgrade

- “To purge event logs and statistics records.”
Complete once per environment.
- “To stop all Process Engine-related services and applications.”
Complete this on each node.

Upgrading server farms (UNIX, Windows)

Perform the following steps to upgrade your highly available Process Engine in a server farm environment.

Upgrade tasks

1. Make sure the farm nodes are running.
2. Update the Process Engine per “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
NOTE In the “Content Engine API Configuration” screen make sure to input information for a highly available Content Engine. For more information, and for examples, see [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#).
3. Perform post-update tasks as documented in the “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

CAUTION Do not re-enable Oracle Password Complexity Verification until after all nodes have been updated.

- “To set the f_maint and f_sw passwords.”
Complete this on each node.
 - “To restore any custom modifications for root and fnsu users (UNIX).”
Complete this on each node.
 - “To edit the /etc/inittab file.”
Complete this on each node.
 - “To edit the pe_start file (HP).”
(HP only) Complete this on each node.
4. Repeat [Step 2](#) and [Step 3](#) for all farm nodes that will run Process Engine services.
 5. On each node, perform the “Install Process Engine software updates” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 6. On each node, perform the “Upgrade the Content Engine Client files on Process Engine servers” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE Make sure to input information for a highly available Content Engine. For more information, and for examples, see [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#).

7. Perform the tasks listed in the “Complete post-upgrade Process Engine configuration” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

Upgrade from Process Engine 4.0.2-001, 4.0.3, or higher

- “To update the database to version 4.5.”

Perform this step once per environment from one of the farm nodes.

- “To update isolated regions.”

Perform this step once per environment from one of the farm nodes.

8. Start the Process Engine services on all nodes.

Upgrading Process Engine in a Veritas cluster server (UNIX)

Perform the following steps to upgrade your highly available Process Engine in a Veritas cluster server environment.

Upgrade tasks

1. Make sure the cluster group is running on the node being upgraded.
All cluster resources for the Process Engine group (with the exception of the Process Engine resource itself, must be online.
2. Update the Process Engine as documented in the “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
3. Perform post-update tasks as documented in the “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

CAUTION Do not re-enable Oracle Password Complexity Verification until after all nodes have been updated.

- “To set the f_maint and f_sw passwords.”
Complete this on each node.
 - “To restore any custom modifications for root and fnsw users (UNIX).”
Complete this on each node.
 - “To edit the /etc/inittab file.”
Complete this on each node.
 - “To edit the pe_start file (HP).”
(HP only) Complete this on each node.
4. Move the cluster group to another node and repeat [Step 2](#) and [Step 3](#). Do this for all nodes in the cluster that will run Process Engine services.
 5. On each node, perform the “Install Process Engine software updates” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 6. On each node, perform the “Upgrade the Content Engine Client files on Process Engine servers” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE Make sure to input information for a highly available Content Engine. For more information, and for examples, see [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#).

7. Perform the tasks listed in the “Complete post-upgrade Process Engine configuration” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

Upgrade from Process Engine 3.5.x

- “To update the Process Engine database objects.”
 Perform this step once per environment from the active cluster node.
- “To update the Process Engine security.”
 Perform this step once per environment from the active cluster node.
- “To migrate routers and update isolated regions.”
 Make sure to use the virtual name of the Process Engine cluster when creating the Process Engine Region ID from FileNet Enterprise Manager.
 Perform To configure new regions in Process Task Manager once per environment on the active cluster node.
- “To update email notification.”
 Perform this step once per node.
- “To update isolated regions.”
 Perform this step once per environment from the active cluster node.
- “(SQL Server and DB2 only) To run post-upgrade scripts.”
 Perform this step once per environment from the active cluster node.
- “To back up the database and restart software.”
 Perform this step once per environment from the active cluster node.

Upgrade from Process Engine 4.0.2-001, 4.0.3, or higher

- “To update the database to version 4.5.”
 Perform this step once per environment from the active cluster node.
- “To update isolated regions.”
 Perform this step once per environment from the active cluster node.

8. Stop the Process Engine services prior to putting them under Veritas control.

Execute:

```
killfns w -DAyS
```

Perform this step once per environment from the active cluster node.

9. Install the FileNet P8 4.5 Process Engine Agent.
 - a. Untar the FileNet Process Engine Agent package located under the Veritas folder on the FileNet P8 Process Engine software package.
 - b. Copy the files to all nodes in the cluster.

- c. On all nodes perform the following steps:
- i. Create the directory `/opt/VRTSvcs/bin/FN_ProcessEngine`.
 - ii. Copy the `FN_ProcessEngineAgent` agent binary to the following directory (overwriting the existing file):

```
/opt/VRTSvcs/bin/FN_ProcessEngine
```

- iii. Change the permissions on the directory as follows:

```
chmod -R 755 /opt/VRTSvcs/bin/FN_ProcessEngine
```

- iv. Copy the file `FN_Types.cf` of the following directory (overwriting the existing file):

```
/etc/VRTSvcs/conf/config
```

- d. Restart the Veritas cluster service on each node to read the new FileNet Agent types.
- e. Start the FileNet Process Engine Agent for Veritas on all nodes
- f. A Process Engine resource should now be listed as one of the available resource types:
Type the command:

```
hatype -list
```

Verify that one of the listed resources is "FN_ProcessEngine."

- g. Create a Process Engine resource.
Enable the resource and verify that it shows as "offline."
- h. Make the Process Engine resource dependent on the clustered IP and the Process Engine partitions on clustered storage:
 - local partition resource
 - `fn_sec_db0` partition resource
 - `fn_sec_rl0` partition resource

10. Bring the cluster online.

Upgrade Process Engine in a Veritas or Microsoft cluster server (Windows)

Perform the following steps to upgrade your highly available Process Engine in a Veritas cluster server or Microsoft cluster server environment.

Upgrade tasks

1. Make sure the cluster group is running on the node being upgraded.

All cluster resources for the Process Engine group (with the exception of the Process Engine resources), must be online.

2. Update the Process Engine as documented in the “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
3. Perform post-update tasks as documented in the “Upgrade Process Engine interactively” or “Upgrade Process Engine silently” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

CAUTION Do not re-enable Oracle Password Complexity Verification until after all nodes have been updated.

- “To set the f_maint and f_sw passwords.”

Complete this on each node.

- “To edit the /etc/inittab file.”

Complete this on each node.

4. Move the cluster group to another node and repeat [Step 2](#) and [Step 3](#). Do this for all nodes in the cluster that will run Process Engine services.
5. On each node, perform the “Install Process Engine software updates” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
6. On each node, perform the “Upgrade the Content Engine Client files on Process Engine servers” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE Make sure to input information for a highly available Content Engine. For more information, and for examples, see [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#).

7. Perform the tasks listed in the “Complete post-upgrade Process Engine configuration” task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

Upgrade from Process Engine 3.5.x

- “To update the Process Engine database objects.”

Perform this step once per environment from the active cluster node.

- “To update the Process Engine security.”

Perform this step once per environment from the active cluster node.

- “To migrate routers and update isolated regions.”

Make sure to use the virtual name of the Process Engine cluster when creating the Process Engine Region ID from FileNet Enterprise Manager.

Perform To configure new regions in Process Task Manager once per environment on the active cluster node.

- “To update email notification.”

Perform this step once per node.

- “To update isolated regions.”

Perform this step once per environment from the active cluster node.

- “(SQL Server and DB2 only) To run post-upgrade scripts.”

Perform this step once per environment from the active cluster node.

- “To back up the database and restart software.”

Perform this step once per environment from the active cluster node.

- “To configure continuous free memory for Process Engine (Windows only).”

Perform this step once per node.

Upgrade from Process Engine 4.0.2-001, 4.0.3, or higher

- “To update the database to version 4.5.”

Perform this step once per environment from the active cluster node.

- “To update isolated regions.”

Perform this step once per environment from the active cluster node.

- “To configure continuous free memory for Process Engine (Windows only).”

Perform this step once per node.

8. On all nodes add the NCHBroadcast Value to the Windows registry.

- For Veritas Cluster Server, see [“Modify the registry to set the NCHBroadcast value.” on page 86.](#)

- For Microsoft Cluster Server, see [“Modify the registry to set the NCHBroadcast value.” on page 90.](#)

9. Enable or recreate the cluster resources and bring the cluster online.

If you are recreating the cluster resources see the Process Engine section of this document for details on the parameters and service names to configure.

Upgrade a highly available Application Engine/Workplace XT

This section covers the tasks and requirements for upgrading a highly available Application Engine or Workplace XT in a FileNet P8 environment.

NOTE The Application Engine or Workplace XT upgrade process is similar to a fresh installation in terms of configuration and deployment. Application Engine is the last component to be upgraded in a IBM FileNet P8 platform environment.

Upgrade tasks

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 135](#).
2. Upgrade your application.

Application Engine

- a. Perform the steps in the “Prepare for Application Engine upgrade” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following modifications.
 - Perform the steps on all nodes in your cluster.
 - In the “Copy modified files to the installed Workplace directory.” task, copy the configuration files to the shared configuration directory, if you are using one.
- b. Perform the “Upgrade and configure Application Engine” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following modifications.
 - Perform the steps on all nodes in your cluster.

Workplace XT

Perform the “Upgrade IBM FileNet Workplace XT” task of the *IBM FileNet Workplace XT Installation and Upgrade guide*, with the following modifications.

- Perform the steps on all nodes in your cluster.
3. Perform the “Install Application Engine software updates” (Application Engine) or “Install IBM FileNet Workplace XT software updates” (Workplace XT) task on all nodes.
 4. Perform “Install the latest Content Engine Client files on Application Engine servers” (Application Engine) or “Install the latest Content Engine Client files on Workplace XT servers” (Workplace XT) on all nodes.
 5. Perform “Install the latest Process Engine Client files on Application Engine servers” (Application Engine) or “Install the latest Process Engine Client files on Workplace XT servers” (Workplace XT) on all nodes.
 6. (Application Engine only) Perform “Configure Application Engine upgrades from 3.5.x (*appserver*)” on all nodes.
 7. (Application Engine only) Perform “Manually copy custom data” on all nodes.

8. Deploy Application Engine or Workplace XT to the cluster nodes.

Follow the instructions in [“Configure and deploy the Application Engine/Workplace XT for your application server configuration.”](#) on page 95 and [“\(Farm of independent application server instances\) Configure and deploy the Application Engine/Workplace XT on each node.”](#) on page 97.

9. Application Engine/Workplace XT deployment verification

Follow the instructions in [“Application Engine/Workplace XT deployment verification”](#) on page 97.

Upgrade a highly available Rendition Engine

This section covers the tasks and requirements for upgrading a highly available Rendition Engine in a FileNet P8 environment.

Rendition Engine upgrade

Follow the instructions in the *IBM FileNet Rendition Engine Installation and Upgrade Guide* “Upgrade and configure Rendition Engine” section.

Uninstall previous version of Rendition Engine

Perform the upgrade tasks as outlined in the *IBM FileNet Rendition Engine Installation and Upgrade Guide* “Uninstall previous version of Rendition Engine” task on each Rendition Engine node.

Upgrade Rendition Engine

Perform the upgrade tasks as outlined in [“Install and configure Rendition Engine” on page 99](#).

Rendition Engine upgrade verification

Perform the optional the *IBM FileNet Rendition Engine Installation and Upgrade Guide* “(Optional) Verify Installation” task.

With the addition of each new Rendition Engine server you should verify that the environment works without problems.

Upgrade a highly available CFS-ICI

This section covers the tasks and requirements for upgrading a highly available CFS-ICI in a FileNet P8 environment.

To upgrade an existing CFS-ICI system

1. Uninstall the existing CFS-ICI environment.
2. Install CFS-ICI in your environment, following the procedure in [“Install a highly available CFS-ICI” on page 104](#).

Upgrade a highly available Records Manager

The Records Manager upgrade process is similar to a fresh installation in terms of configuration and deployment.

NOTES

- The following procedure covers upgrades of the Records Manager Web Application only.
- This document covers the following upgrade paths:
 - Version 3.7 to 4.5
 - Version 4.0 to 4.5
- This document does not include upgrade instructions for Records Manager version 3.5.x to Records Manager version 4.0 or 4.5.
- For instructions on upgrading other components such as the RM Sweep Application or Crystal Reports, see the “Upgrade Tasks” section of the *IBM FileNet Records Manager Installation and Upgrade Guide*, version 4.5.

Upgrade Tasks

To upgrade the Records Manager Web application, follow the tasks in the “Upgrade Tasks” section of the *IBM FileNet Records Manager Installation and Upgrade Guide*, version 4.5, with HA specific modifications to specific tasks as outlined below.

1. Verify that you have performed the required preparation tasks as listed in [“Perform required preparation tasks” on page 135](#).
2. Upgrade Records Manager Software.

To upgrade the Records Manager software, on each node, follow the instructions in the “Upgrade Records Manager Software” task with the following HA specific modifications:

- “Identify the following information, which you will be required to supply during the upgrade process.”

Use the same information for the `bootstrap.properties` file and the `WEB-INF` folder you used when you installed Records Manager version 3.7 or 4.0. See [“Install a highly available Records Manager” on page 112](#).

- “Complete the Setup screens as follows:”

When prompted, enter the HA specific information you identified above.

3. Configure report viewing if required.

If your setup is using reports, follow the instructions in the “(Optional) Add Report Viewing Support” task with the following HA specific modifications:

You must copy the Crystal Reports JAR files to all server instances/nodes in the farm/cluster.

4. Install Content Engine and Process Engine Client Files.

Follow the instructions in the task “Install Content Engine and Process Engine client files” on all server instances/nodes in the farm/cluster.

5. Redeploy Workplace or Workplace XT.

Follow the instructions in the “Redeploy Workplace or Workplace XT” task of the *IBM FileNet Records Manager Installation and Upgrade Guide*, version 4.5. For a highly available Workplace/Workplace XT installation, see deployment instructions in [“Upgrade a highly available Application Engine/Workplace XT” on page 158](#).

Perform these steps on all server instances/nodes in the farm/cluster.

6. Redeploy the RM Web application.

Follow the instructions in the “Deploy the Records Manager Web application” task.

NOTE When deploying the Records Manager Web Application, specify the cluster as the deployment location instead of individual server instances.

WebSphere

Refer to the “IBM WebSphere 5.x, 6.x, and 7.x” procedures Using these steps re-deploy Records Manager on the cluster and proxy server instances.

WebLogic

Refer to the “BEA WebLogic 8.1.x, 9.x, and 10.x” procedures. Use these steps to deploy Records Manager on the cluster.

7. Restart Component Manager.

From the Application Engine or Workplace XT server, launch the Process Task Manager. For information, go to the *FileNet P8 online help* and navigate to [System Administration > Enterprise-wide Administration > Process Task Manager > Getting Started > Start Process Task Manager](#).

Appendixes

The appendixes section contains the following major topics:

- [“Appendix A – Setting up a secure native mode domain installation” on page 165.](#)
- [“Appendix B – Connecting to a highly available Content Engine” on page 168.](#)
- [“Appendix C – Using load balancers with IBM FileNet P8” on page 171.](#)
- [“Appendix D – Recompiling the FileNet Process Engine Veritas agent” on page 176.](#)
- [“Appendix E – Managing Workplace/Workplace XT settings in a highly available environment” on page 178.](#)
- [“Appendix F - Installation and upgrade worksheet” on page 181.](#)

Appendix A – Setting up a secure native mode domain installation

Completion of the following configuration steps enables you to proceed with a Process Engine cluster install without having to perform the installation as a Domain Administrator. You can instead use these procedures to create the “installer” user (with limited rights), setup other required FileNet users and groups, and configure the node 1 and node 2 cluster servers.

Configure the domain controller

1. Create FileNet Groups.

Create the following domain local security groups on the domain controller:

Table 1: FileNet Groups

Group Name	Group Description
FNADMIN	Members have all privileges on FileNet files and databases
FNOP	Members can start/stop and execute FileNet software
FNUSR	Members have normal privileges on FileNet files and databases

2. Create FileNet Users.

On the Domain Controller, create the following users:

Table 2: FileNet Users

User Name	User Description	Member of
<i>Installer</i> ^a	User that will install the FileNet P8 Process Engine software.	FNADMIN, FNOP, and FNUSR
fns	Primary FileNet software user	FNADMIN, FNOP, and FNUSR
oracle ^b	Primary Oracle software user	N/A

a.The user name *Installer* is arbitrary; you may use any name you wish. If you plan to retain the *Installer* user after installation you can for example use the fns user for installation.

b.The oracle user name can be local and the user account can remain disabled during the installation.

CAUTION The oracle user name is required even if you have a Microsoft SQL Server relational database. The Process Engine installer checks for this name and will fail if it is not present.

NOTE If you are planning to install Microsoft SQL Server, you do not need to create a special user for the RDBMS software. The SQL Server installer configures the system for use with the SQL Server software.

3. Add Nodes to Pre-Windows 2000 Compatible Access Properties
 - a. On the Active Directory Users and Computers window, click the **Builtin** folder.
 - b. Right-click the Pre-Windows 2000 Compatible Access object and select **Properties**, as shown above.
 - c. Click on the **Members** tab, and click the **Add** button.
 - d. Select **Node 1** and click the **Add** button.
 - e. Select **Node 2** and click the **Add** button.
 - f. Click **OK** on the **Select Users, Contacts, Computers, or Groups** dialog box.

The Pre-Windows 2000 Compatible Access Properties window displays again and shows that the two nodes have been added.

- g. Click **Apply** on the Pre-Windows 2000 Compatible Access Properties dialog box to apply the changes.
 - h. Click **OK** to close the Pre-Windows 2000 Compatible Access Properties window.
 - i. Click **OK** to close the Active Directory Users and Computers window.

Configure node 1 and node 2 servers

Perform these procedures on the node 1 server first, and then repeat them on the node 2 server.

1. Create the LocalAdminInstall File

Use this procedure to create the LocalAdminInstall file in the C:\temp directory.

- a. Turn-on the node 1 and node 2 servers and logon as Domain Administrator on each server.
 - b. Open a Command Prompt window.
 - c. From the C: drive, change to the \temp directory by entering:

```
cd \temp
```

NOTE If the temp directory does not exist on the C: drive, use the mkdir command to create one.

- d. At the \temp directory, type the following command and press Enter:

```
copy con LocalAdminInstall
```

- e. Press and hold **Ctrl** key, and press the **Z** key.
 - f. Press **Enter**.

g. Verify that the LocalAdminInstall file was successfully created in the C:\temp directory.

2. Create Local FileNet Groups

On the node servers, create the following local groups:

Table 3: Local FileNet Groups

Group Name	Group Description
dba	Members should be domain fnsw and domain oracle
ora_dba	Members should be domain fnsw and domain oracle

3. Add Users to Local Admin Group

Add the domain users Installer and fnsw to the local “Administrators” group.

After both the installer and fnsw users have been added to the local “Administrators” group, close the Computer Management window and continue to the next section.

4. Modify the Local Security Policy for the Domain Account (fnsw)

Modify the local security policy to give the fnsw user domain account permissions for the following local policies:

- Windows 2003
 - Act as part of the operating System
 - Log on as a service
 - Adjust memory quotas for a process
 - Replace a process level token

After all Policy selections have been modified, close the Local Security Settings window and return to [“Create the LocalAdminInstall File” on page 166](#) to repeat these procedures on the node 2 server.

5. Return to Main Body of the HA technical notice

After you have performed these procedures on both nodes, return to the Process Engine section to continue with your cluster server system installation.

Appendix B – Connecting to a highly available Content Engine

The Content Engine connection in a highly available environment differs depending on which application you are configuring to connect.

Administrative applications

Administrative applications, such as FileNet Enterprise Manager, modify metadata and GCD information such as property templates, class and property definitions, and domain level objects such as sites, object stores, and file storage areas.

Due to the built-in lag in metadata synchronization across the Content Engine nodes, changes made on one node are not synchronized immediately after a configuration change is made via an administrative application. For more information, see [“Content Engine cache refresh delays” on page 38](#).

To ensure your administrative applications always get a consistent view of the system metadata, you must configure them all to point to a single Content Engine node (the same node), rather than a virtual server.

In the event the configured Content Engine fails you will need to reconfigure your administrative applications to point to a different Content Engine node. For information, go to the *FileNet P8 online help* and navigate to [System Administration > Content Engine Administration > FileNet P8 Domain > How to... > Logon to a domain dialog box](#).

User applications

User applications such as Workplace and Workplace XT are content and process centric applications used in everyday production to check documents in and out and access workflows, for example.

For user applications, the Content Engine connection in a highly available environment differs depending on your HA configuration, but in any case must be configured to connect to the virtual server, so that it will be redirected automatically to another Content Engine node in the event of a Content Engine node failure.

- If a Java application server cluster is used, EJB connections must use a specific format for establishing connections to Content Engine. See [“Content Engine in an application server cluster using EJB transport” on page 169](#).
- Applications using the web service (CEWS) protocol to connect to Content Engine must use the virtual server name of the hardware or software that is load balancing the connection. See [“Connecting using Content Engine Web Service Transport \(CEWS\)” on page 170](#).

Content Engine in an application server cluster using EJB transport

For clients such as Application Engine that use an EJB transport to communicate with Content Engine in an application server cluster configuration an URL with a format different from non-highly available configurations must be used.

When Content Engine is made highly available through an application server cluster configuration the Content Engine URL should have the following form:

WebLogic

SSL

```
cemp:t3S://node1 name:node1 SSL port,node2 name:node2 SSL port/FileNet/Engine
```

Non-SSL

```
cemp:t3://node1 name:node1 port,node2 name:node2 port/FileNet/Engine
```

NOTE Even though the individual server names are listed out WebLogic will load-balance requests through the EJB transport to the Content Engine if the server instances are part of a WebLogic cluster.

Example:

```
cemp:t3://testnode1:7001,testnode2:7001/FileNet/Engine
```

WebSphere.

```
cemp:corbaloc::node 1 hostname:bootstrap port,:node2 hostname:bootstrap port/  
cell/clusters/your websphere cluster name/FileNet/Engine
```

This configuration requires the WebSphere cluster name in addition to the node names as part of the URL. The default bootstrap port (e.g. 9810) is different from the non-default (e.g. 2809) on a non cluster configuration (i.e. a standalone WebSphere installation).

NOTE Only one URL is used regardless of SSL use. When using WebSphere an SSL connection is automatically established using this transport if the SSL port is enabled for the WebSphere application server instance.

Example:

```
cemp:corbaloc::testnode1:9810,:testnode2:9810/cell/clusters/testwascluster/  
FileNet/Engine
```

JBoss

SSL

For information, see JBoss documentation for SSL.

Non-SSL

```
cemp:jnp://node1_name:node1_port,node2_name:node2_port/FileNet/Engine
```

NOTE The behavior of JBoss is the opposite of WebLogic in the note above; instead of load-balancing through the listed servers, JBoss tries the list in order and connects to the first server that responds.

Example:

```
cemp:jnp://testnode1:1100,testnode2:1100/FileNet/Engine
```

Connecting using Content Engine Web Service Transport (CEWS)

Applications like Process Engine use the Content Engine Web Service (CEWS, previously referred to as WSI) transport to connect to the Content Engine. This transport can utilize load-balancers to balance requests to Content Engine servers in a highly available configuration. In this type of configuration a virtual server name is used by clients to access the Content Engine.

Use the virtual server name of the load-balancing device when configuring a connection to Content Engine to ensure that applications will function in the event of a Content Engine failure.

Use the following format for the Content Engine URL:

```
http://virtual server name:virtual port number/wsi/FNCEWS40MTOM
```

Example:

```
http://testvirtual11:7001/wsi/FNCEWS40MTOM
```

Appendix C – Using load balancers with IBM FileNet P8

Farming requires a mechanism to balance the load across all the nodes in a farm, and to redirect client connections to surviving nodes in case of failure. This section summarizes the load balancing options available for IBM FileNet P8 4.5.

Load balancer support for IBM FileNet P8

- Content Engine

Content Engine is stateless and supports both hardware and software load balancers.

NOTE Depending on which of the two transports the client applications use to Content Engine, certain restrictions apply. [Table 4 on page 171](#) outlines the types of load balancing that can be used with Content Engine for each transport.

Table 4: Content Engine

CE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	Microsoft Network Load Balancing
Content Engine Web Services clients (CEWS transport: HTTP) <ul style="list-style-type: none"> • Custom applications could use any of the APIs listed here • Pure web service clients • COM API clients <ul style="list-style-type: none"> – Capture 	Yes	No	Yes	Yes

Table 4: Content Engine

CE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	Microsoft Network Load Balancing
Content Engine Web Services clients (CEWS transport: HTTP) contd. <ul style="list-style-type: none"> • .NET API clients <ul style="list-style-type: none"> – IBM Content Collector (ICC) – Sharepoint Connector • Java API clients when using CEWS as transport <ul style="list-style-type: none"> – Process Engine 	Yes	No	Yes	Yes
Java clients (EJB transport: IIOP/T3/JNP) <ul style="list-style-type: none"> • AE/Workplace • Workplace XT • FileNet Records Manager, BPF • Custom Java API clients when using EJB transport 	No (see Caution on page 172)	Yes	No	No
FileNet Enterprise Manager and other administrative applications ^a	No	No	No	No

a.For information, see [“Administrative applications” on page 168.](#)

CAUTION To successfully configure your environment for Content Engine high availability for clients that use the EJB Transport ([Table 4 on page 171](#)), you must use the EJB cluster URL of the application server for EJB client communication with Content Engine. For configuration information, see [“Content Engine in an application server cluster using EJB transport” on page 169.](#) If you use a hardware load-balancer between a client and Content Engine using EJB transport, the load-balancer only performs the initial JNDI lookup. The client then locks to an individual Content Engine node, bypassing the load-balancer.

- Process Engine

Since the Process Engine does not run inside a Java 2 Enterprise Edition (J2EE) application server, a load balancer external to FileNet Process Engine is required. FileNet Process Engine 4.5 is qualified with a hardware load balancer.

NOTES

- If FileNet Content Engine and FileNet Process Engine are collocated, a hardware load balancer is required as FileNet Process Engine does not support software load balancing. If any FileNet Content Engine clients are using the EJB transport, software load balancing by the application server would still be needed for FileNet Content Engine. (See [Table 4 on page 171.](#))
- If you access your Process Engine farms through a load balancer configured to use a dynamic DNS server you must configure Time-To-Live (TTL) setting for your environment. Follow the instructions in [“Configure the Time-To-Live value for load balancers that use dynamic DNS servers” on page 174.](#)

Table 5: Process Engine

PE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	Microsoft Network Load Balancing
Java API AE/Workplace Workplace XT	Yes (TCP for client access, UDP for intraserver communication)	No	No	No
Process Engine Web Service clients (PEWS transport: HTTP)	Yes	No	Yes	Yes

- Application Engine

Application Engine supports both hardware and software load balancers. Note, however, that it is not stateless; you must use a load balancer that supports sticky sessions such as Cisco Local Director (tested by IBM FileNet). You can also use WebSphere or WebLogic HTTP plug-in load balancing in combination with an HTTP server.

Table 6: Application Engine

AE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	Microsoft Network Load Balancing
Web Browsers	Yes + sticky sessions	No	Yes + sticky sessions	Yes + sticky sessions
Application integration (Office)	Yes + sticky sessions	No	Yes + sticky sessions	Yes + sticky sessions

- Database

When using Oracle RAC as the database service, load balancing is provided through the built-in capabilities of both the Oracle RAC client and SQL NET.

Configure the Time-To-Live value for load balancers that use dynamic DNS servers

Process Engine farms accessed through a load balancer configured to use a dynamic DNS server might not be accessible after a failover. In this configuration, the load balancer name-to-address translation is cached by the JVM of the application that accesses the Process Engine API. After a failover, this cache might be pointing to an incorrect address for the active farm node.

To configure your applications for correct failover in this type of environment you must set the Time-To-Live (TTL) setting for host name caching on the application, such as Application Engine, to a value that is equal to or less than the TTL for the DNS server. This forces the application to look up the current name-to-address translation before attempting to connect to the Process Engine load balancer.

To modify the TTL setting

1. On each node in the application server cluster, edit the `java.security` file located in the host `../JVM/jre/lib/security` directory.
2. Update the TTL property:


```
networkaddress.cache.ttl = T
```

 where *T* is equal to or less than the value, in seconds, set for the dynamic DNS server.

Example:

In a highly available IBM FileNet P8 environment Application Engine runs on WebSphere, and the Process Engine farm load balancer is F5 with 3DNS configured for a TTL of 30 seconds.

Edit the following file on each of the Application Engine servers in the cluster:

```
WAS_Home/AppServer/java/jre/lib/security/java.security
```

Set the TTL value to 30 seconds in the following entry:

```
networkaddress.cache.ttl = 30
```

3. Restart the application server.

Appendix D – Recompiling the FileNet Process Engine Veritas agent

Use the following procedure to recompile the FileNet Process Engine Veritas Agent if needed.

NOTE The following procedure can be used to recompile the agent or to produce a new agent binary for use with Veritas Cluster Server.

To recompile the agent

1. Create or identify a directory to use when compiling the agent.
2. Copy the agent archive (for example HPUX_Agent.tar, AIX_Agent.tar, Solaris_Agent.tar) to the directory from the PE CD or ESD image.

NOTE For PE installation, use the PE CD or ESD image. For service pack or interim fix installation, use the corresponding CD or ESD image.

3. Unpack the archive using the following command:

```
tar -xvf agent archive filename
```

Where:

agent archive filename is the name of the agent .tar file.

4. Install a C++ compiler on your system.
Make sure that the compiler matches a supported compiler per Symantec's supported compilers for Veritas Cluster Server Agents, such as:
 - AIX XL C+
 - HPUX aC++
 - Solaris Forte C++
5. Verify that both the bin directory of the installed compiler and the “make” command are in the path of the user compiling the agent.
6. Recompile the agent.
 - a. Move to the directory where you unpacked the agent archive.
 - b. Execute the make command.
 - c. Verify that an agent .o file is created.
 - d. Verify that no compilation errors occurred.
 - e. Execute the make command to create the FN_ProcessEngine agent.

7. (If a FileNet Veritas Process Engine Agent is already installed) Update the current agent.

a. Stop the existing agent on each cluster node using the following command:

```
"haagent -stop FN_ProcessEngine" -sys system name
```

Where:

system name is the name of a system in the cluster running the FileNet Process Engine Veritas Agent.

b. Replace the existing agent binary with the newly created agent binary on all nodes that will run Process Engine.

The default location of the agent is:

```
/opt/VRTSvc/bin/FN_ProcessEngine
```

8. (If no agent is installed) Copy the newly created agent file to:

```
/opt/VRTSvc/bin/FN_ProcessEngine
```

Appendix E – Managing Workplace/Workplace XT settings in a highly available environment

This topic contains documentation for managing Workplace/Workplace XT settings in a load balanced Application Engine/Workplace XT environment consisting of a load balancer or proxy device, several HTTP servers and several application server instances.

When several application server instances host the Workplace/Workplace XT application the goal is to provide to end-users the same work environment from every server instance. Workplace/Workplace XT utilizes a number of preference settings and configuration information to customize the Workplace/Workplace XT application for user needs. In order to provide a seamless work experience in an environment with multiple application server instances residing behind a load-balancer or proxy device it is important that each instance access the same configuration data.

The installer allows you to select the location of a shared configuration directory during the installation process. This shared configuration directory is used to store the files that contain preference settings and configuration information for the Workplace/Workplace XT application. A shared configuration directory allows you to have a single Workplace/Workplace XT instance with the same configuration information and preferences across multiple server instances. The process for managing settings using a shared configuration directory is documented below.

The shared configuration directory should ideally be placed on a highly available share or NFS exported directory that is accessible by all systems in the application server configuration.

Some Component Manager settings are stored on local files on the individual Application Engine servers, and cannot be successfully modified during runtime. Configuration changes made through the user interface will only be applied to the local server, and the modified file(s) will not be distributed throughout the application server configuration. You must distribute the changes manually and restart Component Manager for any changes to take effect as documented in the procedure below.

CAUTION 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/Workplace XT servers.

- On Process Engine servers, the file is located in the `/fnsw/bin` directory.
- On Application Engine servers, the file is located in `C:\Program Files\FileNet\AE\Router` on Windows and in `/opt/FileNet/AE/Router` on UNIX.
- On Workplace XT, the file is located in `C:\Program Files\FileNet\WebClient\Router` on Windows and in `/opt/FileNet/WebClient/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. In an HA environment you must do this for every node that is running Component Manager.

NOTE To be able to perform these tasks you must be a user assigned the Application Engine Administrator access role, and have copy/overwrite permissions on the directories listed below on all nodes.

To update Workplace/Workplace XT settings in a load balanced environment

If you need to update your Workplace/Workplace XT Site Preferences you must refresh the configuration files on each node server.

1. Make your Workplace/Workplace XT configuration changes.
 - a. Sign into Workplace/Workplace XT as a user who is assigned the Application Engine Administrator access role.

Application Engine:

`http://virtual_name:port_number/Workplace`

Workplace XT:

`http://virtual_name:port_number/WorkplaceXT`

- b. Click **Admin**.
 - c. Click **Site Preferences**.
 - d. Update your settings.
 - e. Click **Exit**.
2. Reload the Workplace/Workplace XT configuration files on each node.

After changes are made you must reload the Workplace/Workplace XT configuration files on each node server.

CAUTION To load identical settings on all the nodes you must perform the following steps on each individual node server by logging in to the Workplace instance running on the server. Do not log in using the load balancer URL.

- a. Sign into Workplace/Workplace XT as a user who is assigned the Application Engine Administrator access role.

Use the following URL to sign in:

Application Engine:

`http://node_server_name:port#/Workplace`

Workplace XT:

`http://node_server_name:port#/WorkplaceXT`

- b. Click **Admin**.
 - c. Click **Site Preferences**.
 - d. Click **Refresh**.
 - e. From the Refresh page, click **Reload configuration files**.
 - f. Click **Exit**.

For more information and a list of all configuration files that can be reloaded, go to the on-line help and navigate to:

Application Engine

Developer Help > Workplace Development > Workplace Customization Guide > Appendixes > Reloading Workplace Configuration Files.

Workplace XT

Expansion Products > IBM FileNet Workplace XT > Configure Workplace XT > Manage Workplace XT configuration files, scroll down to Reloading Workplace XT configuration files.

To update Component Manager settings in a highly available environment

1. From one server make changes to the Component Manager as needed.
2. Replicate the vwtaskman.xml file throughout the application server configuration.

Copy this file from:

Application Engine

AE_install_path/FileNet/AE/Router

Workplace XT

install_path/FileNet/WebClient/Router

to the corresponding location on all nodes running Component Manager.

3. Restart all Component Managers to pick up the changes made.

Appendix F - Installation and upgrade worksheet

The *Installation and Upgrade Worksheet* is a spreadsheet that contains properties, parameters, and field names for installation, upgrade, and initial configuration of FileNet P8 Platform components and add-on products. Administrators who are preparing the environment for installation or upgrade of IBM FileNet P8 Platform can use the worksheet during their preparation tasks to record the appropriate values and provide them to the Installation Administrator who will run the installation or upgrade.

You can access the worksheet through the following link:

[Installation and Upgrade Worksheet](#)

or on the IBM product documentation page. To download this and other IBM FileNet product documentation from the IBM web site, see [“Access IBM FileNet documentation, compatibility matrices, and fix packs” on page 20](#).

Worksheet organization

Review the following considerations about the organization of the worksheet:

- The two highlighted columns, “Property or Parameter” and “ENTER YOUR VALUE HERE”, provide the simplest view of the requirement. The others add identifying information and help you sort the rows usefully.
- For descriptions of the non-obvious properties and parameters in the worksheet, see the “Installation and upgrade glossary” in the *Plan and Prepare Your Environment for IBM FileNet P8* document.
- The “Role” column uses the following acronyms for Administrator roles:
 - ITA: Information Technology Administrator
 - ASA: Application Server Administrator
 - DBA: Database Administrator
 - SA: Security Administrator
 - P8A: FileNet P8 Administrator

Optimizing the worksheet

The worksheet provides all possible properties, parameters, and field names for installation, upgrade, and initial configuration of P8 Platform components as well as add-on products. Only a small subset of these rows will apply to a particular installation or administrator role as part of that installation. For example, a database administrator who is only working with DB2 can receive a smaller spreadsheet with rows that have been sorted for the Admin Role of DBA, and the independent software vendor of DB2.

You may receive a simplified or more focused variation of this spreadsheet that is appropriate to your environment or role. Work with your install team and your IBM services representative to make the best use of this tool.

HINT With the **Data > Filter > AutoFilter** command enabled, as it is by default in the shipping worksheet file (`p8_worksheet.xls`), perform the following actions to quickly see only the installation properties you must specify for a given installer or configuration program:

- Click the **AutoFilter** drop-down arrow in the "Installation or Configuration Program" column header and select the program you are interested in (for example, PE installer).
- Click the **AutoFilter** drop-down arrow in the "Setup Type" column header, select Custom, and specify: Setup Type contains "Installation."
- Click the **AutoFilter** drop-down arrow in all other column headers and select (All).

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows 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.

Other company, product, and service names may be trademarks or service marks of others.

U.S. Patents Disclosure

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.



Program Number: 5724-R76, 5724-R81

Printed in USA

GC31-5610-02

