


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.

Typographical Conventions

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

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

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

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

Revision Log

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

Date	Revision
06/09	<p>Added "(Process Engine 4.0.1) Synchronize vwtaskman.xml for correct Process Engine high availability failover" on page 73.</p> <p>Updated BPF deployment procedures on JBoss in "Manually Deploy the BPF Web Application." on page 109 and "Manually Deploy the BPF Web Application." on page 166.</p> <p>Updates to "Appendix B – Connecting to a Highly Available Content Engine" on page 173 and "Appendix C - Using load balancers with IBM FileNet P8" on page 176.</p>
12/08	<p>Updated "Installing a Highly Available Records Manager" on page 95 with information about installing Records Manager Version 4.5.</p> <p>Added "Installing a Highly Available IBM FileNet Services for Lotus Quickr" on page 135.</p>
07/08	<p>Updated the Installing a Highly Available Records Manager topic by adding Step 4 "(DoD Chapter 4/DoD Classified only) Create the Classified Records Security Personnel Access Role." on page 98.</p> <p>Added "Configure the Time-To-Live value for load balancers that use dynamic DNS servers" on page 179.</p>
06/08	<p>Added support for Connector for SharePoint in "Installing a Highly Available IBM FileNet Connector for SharePoint Web Parts 2.2" on page 121 and "Installing a Highly Available IBM FileNet Connector for SharePoint Document Libraries 2.2" on page 126.</p> <p>Added overview topic for functional expansion components in "Functional Expansion Components Upgrade Tasks" on page 157.</p> <p>Added support for upgrading BPF installations in "Upgrading a Highly Available BPF" on page 161.</p>
05/08	<p>Added clarifying note regarding session replication in "Server Clusters" on page 19.</p> <p>Updated section "High Availability Products Tested by the IBM FileNet P8 platform" on page 26.</p>
03/08	<p>Added support for Content Federation Services for Information Integrator Content Edition (CFS-IICE) in "Installing a Highly Available CFS-IICE" on page 112.</p> <p>Updated information about load balancers in "Appendix C - Using load balancers with IBM FileNet P8" on page 176.</p>

Date	Revision
01/31/08	<p>Added support for Business Process Framework in “Installing a Highly Available BPF” on page 104.</p> <p>Removed instructions for deploying Content Engine in independent farms of application servers as this configuration is no longer recommended.</p>
01/08/08	<p>Reintroduced and updated Workplace XT information in “Installing a Highly Available Application Engine/Workplace XT” on page 76.</p>
12/07	<p>Updated information about session affinity in “Load-Balancer or Proxy Server” on page 35.</p> <p>Added support for Connector for SharePoint Web Parts in “Installing a Highly Available Connector for SharePoint Web Parts 2.1” on page 118.</p>
11/07	<p>Updated web-site locations for IBM FileNet product information, documentation, and software updates.</p> <p>Added “Access IBM FileNet Documentation” on page 13 to clarify how to find product documentation on the IBM web site.</p> <p>Added support for JBoss and WebSphere 6.1.x in “Installing a Highly Available Content Engine” on page 34 and “Installing a Highly Available Application Engine/ Workplace XT” on page 76.</p> <p>Added “Installing a Highly Available Rendition Engine” on page 101.</p> <p>Added “Appendix C - Using load balancers with IBM FileNet P8” on page 176.</p>
09/07	<p>Added support for Workplace XT in “Installing a Highly Available Application Engine/ Workplace XT” on page 76.</p> <p>Added “Installing a Highly Available Records Manager” on page 95.</p>
03/07	<p>Added DB2 details throughout the document.</p> <p>Added “Service Pack, Fix Pack and Interim Fix Installation Tasks” on page 88.</p> <p>Added information on PE farming. See “Server Farm Environments” on page 146.</p> <p>Added the following appendixes:</p> <ul style="list-style-type: none"> • “Appendix B – Connecting to a Highly Available Content Engine” on page 173. • “Appendix D – Recompiling the FileNet Process Engine Veritas Agent” on page 181.
01/07	<p>Initial revision of the <i>IBM FileNet P8 High Availability Technical Notice.</i></p>

Table of Contents

Typographical Conventions	5
Revision Log	6
IBM FileNet P8 High Availability	13
Intended Audience for this Document	13
Access IBM FileNet Documentation	13
Gather Auxiliary Documentation	14
Terms and Concepts	16
What is High Availability?	16
High Availability and Disaster Recovery	16
Disaster Recovery with IBM FileNet P8 platform	16
Server Virtualization and High Availability	17
The FileNet P8 Architecture in a Highly Available Environment	17
Building Blocks for High Availability	17
Server Farms	18
Server Clusters	19
Geographically Dispersed Clusters	20
Cluster Configurations	21
Asymmetric Clusters	21
Symmetric Clusters	21
N + 1 Clusters	22
N to M Clusters	23
Choosing Between a Farm or a Cluster	24
Industry Terminology	24
High Availability Products Tested by the IBM FileNet P8 platform	26
Load Balancing/Farming	26
Server Clusters	26
Database Services	27
Content Search Engine High Availability	27
Load Balancing/Farming	28
WebLogic Clusters (Farms)	28
WebSphere Clusters (Farms)	29
JBoss Clusters	29
Server Clusters	29
Veritas Cluster Server	30
Microsoft Cluster Server	30
Database Services	30
Oracle Real Application Clusters	30
Core Components Installation Tasks	32
Overview	32
Installing a Highly Available Content Engine	34
Overview	34

Requirements and General Information	35
Load-Balancer or Proxy Server	35
Content Engine Installation Location	35
Oracle RAC Configuration	35
Content Engine Cache Refresh Delays	36
Procedure overview for Content Engine	37
Pre-installation tasks	38
Installation and Deployment Tasks	40
Creating Object Stores in FileNet P8 4.0.x	47
General Information	47
Enable Application Server Connections after Database Failover	48
Deployment Verification	51
File Store HA Configuration	52
Installing a Highly Available Process Engine	53
Overview	53
Procedure Overview for Process Engine	54
UNIX Platforms	54
Server Farm Environments	54
Cluster Environments	54
Windows Platforms	55
Server Farm Environments	55
Cluster Environments	55
Server Farm (UNIX, Windows)	57
Veritas Cluster Server (UNIX)	59
Pre-Deployment Tasks	59
Deployment Tasks	59
Deployment Verification	64
Veritas Cluster Server (Windows)	65
Pre-Deployment Tasks	65
Deployment Tasks	65
Deployment Verification	68
Microsoft Cluster Server	69
Pre-Deployment Tasks	69
Deployment Tasks	69
Deployment Verification	72
(Process Engine 4.0.1) Synchronize vwtaskman.xml for correct Process Engine high availability failover	
73	
Process Engine Database Reconnection	74
Examples of Errors Returned on Database Disconnect	74
Known Issues	75
Installing a Highly Available Application Engine/Workplace XT	76
Overview	76
Requirements and General Information	78
Load-Balancer or Proxy Server	78
Workplace/Workplace XT Settings	78
Shared Storage/Configuration Directories	79
Component Manager	79

Procedure overview for Application Engine/Workplace XT	80
High-Level Steps	80
Pre-Deployment Tasks	80
Deployment Tasks	81
Deployment Verification	83
Managing Workplace/Workplace XT Settings in a Highly Available Environment	85
Service Pack, Fix Pack and Interim Fix Installation Tasks	88
Applying Content Engine Service Packs or Interim Fixes	89
Fix Pack/Interim Fix Installation Tasks	89
Applying Process Engine Service Packs or Interim Fixes	90
Service Pack/Fix Pack/Interim Fix Installation Tasks - Server Farms	90
Service Pack/Fix Pack/Interim Fix Installation Tasks - UNIX Clusters	90
Service Pack/Fix Pack/Interim Fix Installation Tasks - Windows Clusters	91
Applying Application Engine/Workplace XT Service Packs or Interim Fixes	92
Service Pack/Fix Pack/Interim Fix Installation Tasks	92
Functional Expansion Components Installation Tasks	93
Additional Documentation	93
Installing a Highly Available Records Manager	95
Overview	95
Procedure Overview for Records Manager	96
Records Manager Installation and Deployment Tasks	97
Pre-deployment Tasks	97
Deployment Tasks	97
Deployment Verification	100
Installing a Highly Available Rendition Engine	101
Overview	101
Procedure Overview for Rendition Engine	101
Pre-installation Tasks	102
Install and configure Rendition Engine	102
Installing a Highly Available BPF	104
Overview	104
Procedure Overview for BPF	105
High-Level Steps	105
Pre-installation Tasks	105
Installation Tasks	106
Installation Verification	111
Installing a Highly Available CFS-IICE	112

Overview	112
Procedure Overview for CFS	113
CFS Installation and Deployment Tasks	114
Pre-installation Tasks	114
Installation Tasks	114
Deployment Verification	117
Installing a Highly Available Connector for SharePoint Web Parts 2.1	118
Overview	118
Installation Tasks	119
Installation Verification	120
Installing a Highly Available IBM FileNet Connector for SharePoint Web Parts 2.2 ..	121
Overview	121
Installation Tasks	121
Installation Verification	125
Installing a Highly Available IBM FileNet Connector for SharePoint Document Libraries	
2.2.	126
Overview	126
Installation Tasks	126
Installation Verification	131
Installing a Highly Available IBM FileNet Services for Lotus Quickr	135
Overview	135
Procedure Overview for FileNet Quickr Services	135
IBM FileNet Services for Lotus Quickr Installation and Deployment Tasks	136
Pre-deployment Tasks	136
Deployment Tasks	136
Deployment Verification	137
Core Component Upgrade Tasks	139
Upgrading a Highly Available Content Engine	140
Overview	140
Content Engine Upgrade Overview	141
Content Engine Upgrade	142
The 3.5.x/4.0.x Content Engine Farm/Cluster after Upgrade	145
Object Store Availability in a Mixed 3.5.x/4.0.x Environment	145
Upgrading a Highly Available Process Engine	146
Procedure Overview for Process Engine Upgrade	146
Server Farm Environments	146
Cluster Environments	146
Migrating PE 3.5.x Clusters to PE 4.0.x Farms	147
Upgrading UNIX Clusters	151
Pre-Upgrade Tasks	151

Upgrade Tasks	151
Upgrading Windows Clusters	154
Pre-Upgrade Tasks	154
Upgrade Tasks	154
Upgrading a Highly Available Application Engine	156
Procedure overview for Application Engine Upgrade	156
Pre-Upgrade Tasks	156
Upgrade Tasks	156
Functional Expansion Components Upgrade Tasks	157
Additional Documentation	157
Upgrading a Highly Available Records Manager	158
Procedure overview for Records Manager Upgrade	159
Pre-upgrade Tasks	159
Upgrade Tasks	159
Upgrading a Highly Available BPF	161
Procedure Overview for BPF	162
High-Level Steps	162
Pre-upgrade Tasks	162
Upgrade Tasks	163
Upgrade Verification	168
Appendixes	169
Appendix A – Setting up a Secure Native Mode Domain Installation	170
Configure the Domain Controller	170
Configure Node 1 and Node 2 Servers	171
Appendix B – Connecting to a Highly Available Content Engine	173
Administrative applications	173
User applications	173
Content Engine in an Application Server Cluster using the EJB Transport	174
Connecting Using the Web Service Transport	175
Appendix C - Using load balancers with IBM FileNet P8	176
Load balancer support for IBM FileNet P8	176
Configure the Time-To-Live value for load balancers that use dynamic DNS servers	179
Appendix D – Recompiling the FileNet Process Engine Veritas Agent	181
Notices	183
Trademarks	185
U.S. Patents Disclosure	185

IBM FileNet P8 High Availability

This document describes the technical issues related to configuring the FileNet® P8 Platform for high availability. The first section defines the terms and concepts of high availability and describes some of the third-party products that could be deployed to make a IBM FileNet P8 platform highly available. Subsequent sections describe high availability configurations of the FileNet P8 4.0.0 Content Engine, Process Engine, and Application Engine, as well as expansion products such as Records Manager and Workplace XT. The final section provides information about how to upgrade an existing FileNet P8 3.5.x highly available environment to version 4.0.0.

CAUTION These instructions are provided as overview material relying on third-party and FileNet documentation for more detailed instructions. The document is *not* intended as a step-by-step guide detailing the use of the many third-party software and tools required to create and configure highly available environments. For general third-party information, see the documentation provided by these software and tool vendors.

Intended Audience for this Document

This technical notice is for experienced system and database administrators who will be using this documentation to implement their intended environments.

The procedures in this document require:

- A working administrative knowledge with the FileNet P8 system.
- Experience with BEA WebLogic Server, IBM WebSphere®, or JBoss for setting up application server clusters.
- Experience in high availability software, such as Microsoft® Cluster Server (MSCS) or Veritas Cluster Server to set up and configure fail-over configurations.
- Familiarity with hardware and/or software load balancer setup and configuration.
- Experience with supported operating systems and hardware platforms, volume management software, and database platforms.

Access IBM FileNet Documentation

To access documentation for IBM FileNet products

1. Access the [Product Documentation for FileNet P8 Platform](http://www.ibm.com/support/docview.wss?rs=3247&uid=swg27010422) support page (<http://www.ibm.com/support/docview.wss?rs=3247&uid=swg27010422>).
2. Select the appropriate product link.

To access compatibility matrices and fix packs for IBM FileNet products

1. Access the [Fix Packs for FileNet P8 Platform](http://www.ibm.com/support/docview.wss?&uid=swg27010146) support page (<http://www.ibm.com/support/docview.wss?&uid=swg27010146>)
2. From the Fix Pack page:
 - To access the compatibility matrix, under FileNet P8 Compatibility Matrices, click **Matrix**.
 - To access the fix pack you need, under a specific product name, click the release number.

Gather Auxiliary Documentation

NOTE For general instructions on how to navigate to this and other IBM FileNet product documentation on the IBM web site, see [“Access IBM FileNet Documentation” on page 13](#).

Before you begin installing the software, retrieve the documents listed below.

For IBM FileNet platform components:

- *IBM FileNet P8 Platform Installation and Upgrade Guide* (PDF)
- *IBM FileNet P8 Hardware and Software Requirements* (PDF)
- *IBM FileNet P8 Release Notes* (PDF)

For IBM fileNet Workplace XT:

- *IBM FileNet Workplace XT Installation Guide* (PDF)

For IBM fileNet BPF:

- *IBM FileNet Business Process Framework Installation Guide* (PDF)

For IBM FileNet CFS:

- *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* (PDF)
- *High availability deployment scenario for WebSphere Information Integrator Content Edition* (PDF)

NOTE To access this whitepaper, in the Select a category list above, choose “Websphere,” then “WebSphere Information Integrator Content Edition,” then, under Learn select “Whitepapers”.

For IBM FileNet Connector for SharePoint Web Parts:

- *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.1* (PDF)
- *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.2* (PDF)
- *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* (PDF)

For IBM FileNet Records Manager:

- *IBM FileNet Records Manager Installation and Upgrade Guide* (PDF)

For IBM FileNet Rendition Engine:

- *IBM FileNet Rendition Engine Installation and Upgrade Guide* (PDF)

You will also need documentation for third-party products such as web application servers, load balancers, and other tools that you will be using to set up your highly available environment.

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. This 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, 99.999%, at the high end, to basic availability at 95%. This is a percentage of scheduled up time for a system, so five nines requires a system to be up 99.999% of that scheduled time. 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. Analysts note that the cost of providing high availability increases exponentially as the target moves from 95% to 99% to 99.999%, so prudent system owners take into account the risk of downtime to their business when selecting their high availability targets.

Even a highly available system can still fail for a number of reasons, including human errors and process problems, in addition to hardware or software failures. Making the hardware and software highly available is a necessary component in high availability, but professional and reliable system administration and well designed applications are equally necessary, if not more so. This technical notice addresses only the hardware and software issues, but FileNet customers need to consider all the components in providing high availability.

The goal of high availability is to continue to provide a user with a working system as seamlessly as possible in the event of a component failure. 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.

High Availability and Disaster Recovery

High availability solutions provide business continuity in the face of localized failures, such as a single server failure or hard disk crash. Disaster recovery solutions, on the other hand, provide for business continuity in the face of natural or man-made disasters that cause the loss of an entire production system. While the goal of both high availability and disaster recovery solutions is the same—keeping your Enterprise Content Management (ECM) system available for continued business operations—the solutions themselves are quite different. A disaster recovery solution must provide a complete alternate system, with current or near-current data, typically at a geographically remote site unaffected by the disaster. Disaster recovery solutions may also include an alternate working site for users of the system, if their primary work location is no longer available due to a disaster. In contrast, a high availability solution typically provides for an alternate system component that takes over for a failed component at the same site.

Disaster Recovery with IBM FileNet P8 platform

Technical issues related to Disaster Recovery for the IBM FileNet P8 platform are covered in the *FileNet P8 Disaster Recovery Technical Notice*. To download this document from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Server Virtualization and High Availability

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

One enabling technology for this consolidation is logical partitioning, where one server can be configured to act like several completely 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 FileNet 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.

The FileNet P8 Architecture in a Highly Available Environment

In a highly available environment FileNet P8 components such as Application Engine and Content Engine are each installed and deployed as a single application that is spread to multiple running instances. These FileNet P8 components are not separated into a presentation component and business logic component; each component exists as an all-inclusive application.

Example:

When setting up Application Engine you cannot separate the images and static content of the Workplace application and have them run on a separate server instance from the servlets and Java™ server pages of the Workplace application.

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 a load-balancer or proxy server is used in a configuration of 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” on page 13](#).

Building Blocks for High Availability

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.

Other essential 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.

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. For example, 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. Microsoft Windows also 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.

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. There are many competing load balancing products that should be considered to find the best combination of price and performance; Cisco and Microsoft are just two of many load balancing vendors.

Server farms are best suited to server tiers that are processing-centric rather than data-centric, because all the servers in the farm are clones of each other. Processing logic does not change often, so it is not difficult to keep all the servers identical in a processing-centric tier. Web servers, and application servers executing business logic, are both good candidates for server farms.

Data-centric tiers, such as file servers and data servers, are not well suited for farming, because their data content is constantly changing. Keeping this dynamic data identical across a group of cloned file or data servers, and managing the client accesses to the data to protect the integrity of

the data in the face of multiple writers, would be difficult in a farm, with a copy on each server. The better solution for this type of server is a server cluster, described next.

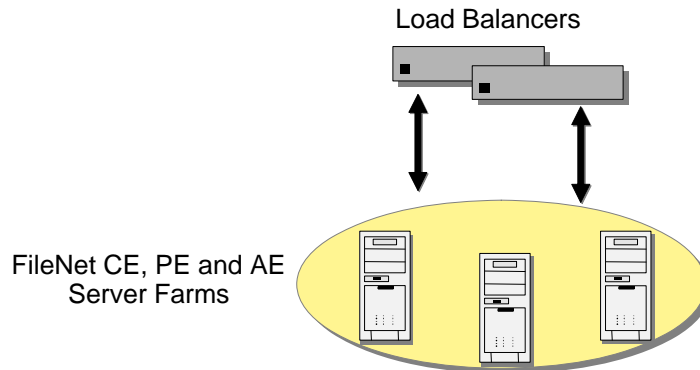


Figure 1: The basic server farm setup.

Server Clusters

Server clusters are based on the concept of shared software configuration data storage, in contrast to server farms, which do not share software configuration data among the servers in the farm. Server hardware and software vendors offer vendor-specific server clustering products as their high availability offering for these kinds of data-centric servers. These products all have the following general characteristics:

- 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 (a so-called failover). When the failed server is repaired and ready to run again, the workload is shifted back over from the other server (a so-called 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.

- The failover feature can mask both planned and unplanned outages from users. For instance, an intentional failover can be done to allow one of the servers to be upgraded or backed up and then brought back online in a failback.

- 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 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 2000 and Windows Server 2003, and Sun Cluster for Solaris. Symantec offers its Veritas Cluster Server product for all of these platforms.

Server clusters typically communicate through a broadcast or share a central repository to keep track of cluster information and cluster node status.

Symantec and Microsoft cluster technologies are similar. 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.

Geographically Dispersed Clusters

FileNet has not tested—and generally does not recommend—“geographically-dispersed” active-passive clustering configurations. In these 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.

This kind of server cluster is much more complex than a local cluster with both nodes in the same location, due to the data replication requirement and the greater potential for network problems between the servers. Server clusters depend on multiple reliable network connections to monitor each other's state and make correct decisions on when to trigger a failover from the active server

to the passive server. In the course of a server failover and failback, after a primary server failure and subsequent repair, data synchronization and integrity problems can arise in geographically-dispersed clusters that would be avoided altogether by using a local active-passive server cluster with a single, shared, highly available copy of the data.

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. By definition disaster recovery solutions assume the loss of the primary data center, necessitating data replication to insure data survival, and wide area networking to insure the recovery site is outside the impact zone of the disaster. For more information on disaster recovery configurations, see the *FileNet P8 Disaster Recovery Technical Notice*. To download this document from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Cluster Configurations

Cluster software vendors support several distinct types of cluster configurations. Using Symantec's terminology, the notable configurations tested by the FileNet P8 4.0.0 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.




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.

A simple example of a FileNet P8 4.0.0 N+1 cluster has three nodes:

1. Process Engine server node
2. Remote database server node
3. 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.

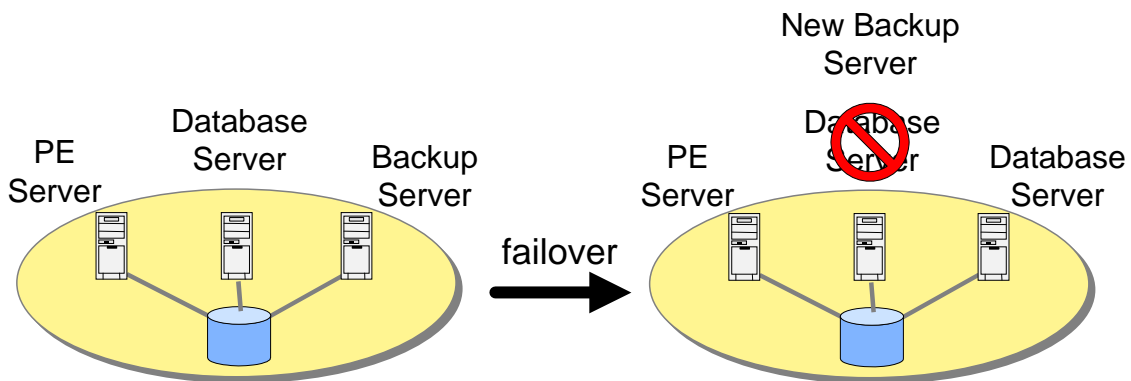


Figure 5: N+1 cluster.

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.

A FileNet P8 N to M cluster could for example have four nodes:

1. Process Engine server node (PE Server)
2. Content Engine Database server node (CE Database Server)
3. Process Engine Database server node (PE Database Server)
4. 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.

Choosing Between a Farm or a Cluster

The technology direction for FileNet products is to support load-balanced server farms wherever possible, in preference to active-passive server clusters.

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. As farming support is added to FileNet P8 products that previously only supported active-passive server clusters—such as the P8 4.0.0 Process Engine—we expect most FileNet P8 systems to deploy farms in order to take advantage of these benefits.

NOTE Over time, FileNet expects a rapidly shrinking demand for clustered configurations in favor of farmed configurations. Therefore we anticipate reducing and eventually eliminating the testing and qualification of active-passive server clusters for products that support farming.

Industry Terminology

Vendors and industry analysts use a variety of differing terms for these 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.

The table below provides a mapping of the various industry terms to the terminology used here.

	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 their WebSphere Clones (farms). For all intents and purposes this can be viewed as “Session Failover” in the scope of this document.

High Availability Products Tested by the IBM FileNet P8 platform

FileNet has tested the IBM FileNet P8 platform with high availability products from BEA Systems, IBM, Microsoft, Symantec, F5 networks, JBoss, and Oracle. These products are described below. There are many other high availability vendors, products, and configurations that should work with the IBM FileNet P8 Platform. In general, the IBM FileNet P8 Platform provides an infrastructure that can be leveraged by any high availability technology.

CAUTION For high availability configurations that differ significantly from those described here, however, it is prudent to do a proof of concept test, or even a pilot project, to confirm that the configuration functions as expected.

Load Balancing/Farming

For more information, see [“Load Balancing/Farming” on page 28](#).

Table 1: 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.

Server Clusters

For more information, see [“Server Clusters” on page 29](#).

Table 2: Tested server clusters

Vendor	Notes
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	Windows

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. Please see [“Veritas Cluster Server \(UNIX\)” on page 59](#) for more details on the Veritas high availability agent for the Process Engine.

Database Services

For more information, see [“Database Services” on page 30](#).

Table 3: Tested database services

Vendor	Notes
IBM DB2® UDB with Veritas Cluster Server	-
Oracle with Veritas Cluster Server	-
Oracle Real Application Clusters (RAC)	See “Oracle Real Application Clusters” on page 30 .
SQL Server with Microsoft Cluster Server (MSCS)	-

Content Search Engine High Availability

Content Search Engine is supported in a highly available Content Engine configuration, but cannot itself be made highly available in a cluster or farm configuration. The Content Search Engine component will be a single point of failure for an otherwise highly available Content Engine.

In a highly available Content Engine environment only one server at a time in the farm can be selected to work the Verity CBR dispatcher queue. If the server hosting that function goes down, another server must be manually configured to resume operation. To do this, the Enable Dispatcher property must be set for the selected Content Engine server using Enterprise Manager.

NOTE There is no crash recovery mechanism for index jobs in the case where a server goes down while an index job is running. If the server goes down, the index job is terminated, but it remains in the database and is still visible in the index jobs list control. You can manually delete the terminated index job to remove it from the list.

To configure a Content Engine server to work with the dispatcher

The Verity server configuration is set at the Domain, Site, Virtual Server (Node) and Server levels. For multiple server systems such as farms the Enable Dispatcher property must be set on one server only, not on multiple servers.

1. Log in to Enterprise Manager.
2. Navigate to: Enterprise Manager (domain) > Sites > <site_name> > Virtual Server

3. Configure one of the Content Engine servers with the Enable Dispatcher property set.
 - a. Expand the virtual server for the server you want to configure.
 - b. Right-click the <server_name> and select **Properties**.
 - c. Select the **Verity Server** tab.
 - d. Check “Override inherited settings.”
 - e. Check “Enable Dispatcher.”
 - f. Click **OK**.
4. Configure all other Content Engine servers with the Enable Dispatcher property unset.
 - a. Expand the virtual server for the server you want to configure.
 - b. Right-click the <server_name> and select **Properties**.
 - c. Select the **Verity Server** tab.
 - d. Check “Override inherited settings.”
 - e. Verify that “Enable Dispatcher” is unchecked.
 - f. Click **OK**.

Load Balancing/Farming

There are a number of hardware and software load-balancing products available for server farm configurations. For more information, see [“Appendix C - Using load balancers with IBM FileNet P8” on page 176](#).

- 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 173](#).

Server Clusters

Veritas Cluster Server and Microsoft Cluster Server are two examples of server cluster products.

Veritas Cluster Server

Symantec provides a software package called Veritas Cluster Server (VCS). Once installed and configured, VCS 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.

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.

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).

Oracle Real Application Clusters

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 [“Oracle RAC Configuration” on page 35](#) and [“\(WebSphere and WebLogic\) Oracle RAC JDBC connection URL.” on page 47](#).

Core Components Installation Tasks

This section provides the information needed to install a highly available FileNet P8 4.0.x environment.

Overview

A number of service packs, fix packs and interim fixes have been released since the initial release of 4.0.0 software that are required to reach a fully supported 4.0.x level. This is especially true for customers using DB2 as their database engine.

Depending on your configuration, you may 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 [“Installing a Highly Available Content Engine” on page 34](#).

- 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 [“Installing a Highly Available Process Engine” on page 53](#).

- Application Engine Server Farm.

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

As Application Engine 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 [“Installing a Highly Available Application Engine/Workplace XT” on page 76](#).

Installing a Highly Available Content Engine

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

Overview

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 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.

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

Requirements and General Information

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 176](#).

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. 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 173](#).

Content Engine Installation Location

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

NOTE File store data must still be shared. For more information, see [“File Store HA Configuration” on page 52](#).

Oracle RAC Configuration

For more Content Engine requirements, see the “Prerequisite Tasks” in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- 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

When using Oracle 10g 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.

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 objectstore 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.

Procedure overview for Content Engine

To install Content Engine in a highly available environment

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(s)/tablespace(s).

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 *IBM FileNet P8 Platform Installation and Upgrade Guide* 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. Install and Deploy Content Engine to your application server configuration.
4. Verify the data sources and connection pools for your GCD and object store databases.
5. Enable WebSphere/WebLogic/JBoss connection to your GCD and object store databases after failover.
6. Create a FileNet P8 domain.
7. Create an object store and verify the Content Engine Installation.

Pre-installation tasks

The following procedure assumes that you have created the highly available application server configuration and that it is running prior to installing and deploying FileNet P8 Content Engine.

NOTE Aside from the shared file store folders, you do not need to place any other Content Engine files or data on network or shared storage systems.

To prepare for Content Engine installation

1. 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 P8 Platform Installation and Upgrade Guide* tasks:

- “Verify that Microsoft SQL Server is Installed for FileNet P8”
- “Verify that Oracle Server is Installed for FileNet P8”
- “Verify that DB2 Server is Installed for FileNet P8”
- “Verify that MS SQL Server Client software is Installed for FileNet P8”
- “Verify that Oracle Client is Installed for FileNet P8”
- “Verify that DB2 Client is Installed for FileNet P8”

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 with your application server.

3. Configure the Application Server for Content Engine.

To configure the Application Server on each node follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* with the following modifications:

- WebSphere:
Do the task “Configure an Application Server for Content Engine (WebSphere).”
NOTE Set the JDBC parameters for the nodes using the administrative console, not on the individual servers.
- WebLogic:
Do the task “Configure an Application Server for Content Engine (WebLogic).”
- JBoss:
Do the task "Configure an Application Server for Content Engine (JBoss)", except use JBOSS_DIST/all as the base configuration file set instead of JBOSS_DIST/server.

4. 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.

5. Verify that all instances in your highly available application server configuration are running.

Installation and Deployment Tasks

To install Content Engine

1. Install Content Engine.

The Content Engine installer accepts input to configure and deploy the Content Engine to a Java application server.

CAUTION Do not select the option to create an Authentication provider during CE installation if you have already configured the authentication providers on the application server(s).

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.

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.

After successful installation, undeploy the Content Engine from the Administrative server and redeploy it to the managed nodes in the cluster. For more information, see [“\(Application server cluster only\) Deploy the Content Engine application to the cluster nodes.”](#) on page 40.

NOTE As 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. See [“To deploy on JBoss”](#) on page 45 for details of which files to copy to other nodes.

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

- (Application Server Cluster configurations only) When the installer prompts for the application server instance where Content Engine will be installed select:
 - (WebLogic) The Administrator node
 - (WebSphere) The Deployment Manager node.
 - (JBoss) JBoss cluster, then select the directory containing the JBoss configuration file set.

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.

- 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.

2. (Application server cluster only) Deploy the Content Engine application to the cluster nodes.

To deploy on WebLogic

- a. Restart WebLogic on all instances including the administrative server.
- b. Open the Administrative console and go to the targets tab for the Connection Pools and Data Sources created for the GCD database.

NOTE (WebLogic 9.x) Go to the targets tab for Data Sources.

- c. Select the Weblogic cluster that will be running Content Engine.
- d. Deselect the administration server.

NOTE This will prevent the resources from being deployed to the administration server.

- e. Select the **ContentEngine** application from the WebLogic administration console.
- f. Select the **Targets** tab for the application and select the Weblogic cluster that will be running Content Engine.
- g. Deselect the administration server.

NOTE This will prevent the Content Engine from being deployed to the administration server.

- h. Save all changes.

NOTE (WebLogic 9.x) Lock and release the configuration.

- i. Restart all administration server instances and managed nodes.

To deploy on WebSphere

- a. Undeploy the Content Engine from the deployment manager node.
 - i. Open the WebSphere Administrator console for the configuration and select the **Enterprise Applications** link.
 - ii. Select the FileNet Engine application and click **Uninstall**.
 - iii. Save the changes and synchronize the configuration with all nodes.
 - iv. Restart the deployment manager instance, all node agents, and application server instances in the cell which will host the Content Engine application.
 - v. Depending on your WebSphere version, continue with [Step b](#), [Step c](#), or [Step d](#) once all instances have successfully restarted.

- b. (WebSphere 5.x) Deploy the Content Engine application to the managed nodes in the application server configuration.

- i. Open the WebSphere Administrator console for the configuration and select the **Install new Application** link.
- ii. Browse to and select the **Engine-ws.ear** file, or type in the full path to the **Engine-ws.ear** file then click **OK**.

The file is located in the default location:

<Content Engine Install Path>/FileNet/Content Engine/lib

- iii. Accept the default for the “Preparing for the application installation/You can choose to generate default bindings and mappings” screen.
- iv. Accept the defaults for the “Install New Application\Step 1:Provide options to perform the installation” screen.

- v. Select the database type used in the configuration from the “Deploy EJBs Option - Database Type” drop-down box on the “Install New Application\Step 2:Provide options to perform the ejb deploy” screen.
 - vi. Accept the defaults for the “Install New Application\Step 3:Provide JNDI Names for Beans” screen.
 - vii. Accept the defaults for the “Install New Application\Step 4:Map EJB references to beans” screen”
 - viii. Accept the defaults for the “Install New Application\Step 5:Map virtual hosts for web modules” screen.
 - ix. Select the cluster names and server names that will provide access to Content Engine from the Clusters and Servers selection box on the “Install New Application\Step 6:Map modules to servers” screen.

Select each of the modules listed and click **Apply**. The server field for each module should now show up with the values selected in the Clusters and Servers selection box.
 - x. Review the summary and click **Finish** on the “Install New Application\Step 7:Summary” screen.
- c. (WebSphere 6.0.x) Deploy the Content Engine application to the managed nodes in the application server configuration.
- i. Open the WebSphere Administrator console for the configuration and select the **Install new Application** link.
 - ii. Browse to and select the **Engine-ws.ear** file or type in the full path to the **Engine-ws.ear** file then click **OK**.

The file is located in the default location:
<Content Engine Install Path>/FileNet/Content Engine/lib
 - iii. Accept the defaults for the “Preparing for the application installation/Choose to generate default bindings and mappings” screen.
 - iv. Accept the defaults for the ”Install New Application\Step 1:Select installation options” screen.
 - v. Select the cluster name, server name, and any web servers that will provide access to Content Engine from the Clusters and Servers selection box on the “Install New Application\Step 2:Map modules to servers” screen.

Check the box next to each of the modules listed and click **Apply**. The server field for each module should now show up with the values selected in the Clusters and Servers selection box.
 - vi. Select the database type used in the configuration from the “Deploy EJBs Option - Database Type” drop-down box on the “Install New Application\Step 3:Provide options to perform the ejb deploy” screen.

- vii. Accept the defaults for the “Install New Application\Step 4:Provide JNDI Names for Beans” screen.
 - viii. Accept the defaults for the “Install New Application\Step 5:Map JCA resource references to resources” screen.
 - ix. Accept the defaults for the “Install New Application\Step 6:Map EJB references to beans” screen.
 - x. Accept the defaults for the “Install New Application\Step 7:Map virtual hosts for web modules” screen.
 - xi. Review the summary and click **Finish** on the “Install New Application\Step 8:Summary” screen.
- d. (WebSphere 6.1.x) Deploy the Content Engine application to the managed nodes in the application server configuration.
- i. Open the WebSphere Administrator console for the configuration and select the **Install new Application** link.
 - ii. Browse to and select the **Engine-ws.ear** file or type in the full path to the Engine-ws.ear file then click **OK**.

The file is located in the default location:

<Content Engine Install Path>/FileNet/Content Engine/lib/<Select Path>
 - iii. At "Install New Application", Step 1, accept the defaults and click **Next**.
 - iv. At "Install New Application", Step 2, Map modules to servers, specify the WebServer you are planning to use. Check the box next to each of the modules listed and click **Apply**. The server field for each module should now show up with the values selected in the Clusters and Servers selection box
 - v. At "Install New Application", Step 3, Map virtual hosts for Web Modules, select both Engine-init.war and FileNet CE-MP WebService, and set their values to default_host.
 - vi. At "Install New Application", Step 4, verify your configuration and click Finish. Once the configuration is saved, click **Save**.
- e. Perform post-deployment tasks.
- i. After the application has been loaded and configured in WebSphere, save the changes and synchronize all nodes.
 - ii. (If you are going to access Centera fixed content devices with your FileNet P8 system) Install Centera shared libraries.

NOTE To access Centera fixed content devices you must install Centera shared libraries on each machine where you deploy Content Engine Server. Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Install and Deploy Content Engine,” “To install Centera shared libraries.”

- iii. (WebSphere 5.1.x or 6.0.x) Configure the Classloader settings.
 - 1. Expand the Applications link and select **Enterprise Applications** on the left pane of the WebSphere Administrator console.
 - 2. Click the **FileNetEngine** application.
 - 3. Change the classloader to "ParentLast."
 - 4. Change the war classloader policy to "Application."
 - 5. Click **Apply**.
- iv. (WebSphere 6.1.x) Configure the Classloader settings.
 - 1) From the Administrative Console, expand Applications. Click **Enterprise Applications**, and click your application (default FileNetEngine).
 - 2) From the Configuration tab, click **Class loading and update detection**.
 - 3) Change Classloader order from "Classes loaded with parent class loader first" to "Classes loaded with application class loader first."
 - NOTE** Do this only for the specific web application. Do not change the similar settings for the entire application server.
 - 4) Change the polling interval to a number appropriate for your environment.
 - 5) Click **Apply**.
 - 6) From the WebSphere administrative console, navigate to Applications > Enterprise Applications > FileNetEngine > Manage Modules > FileNet CEMP Web Service.
 - 7) Change the class loader order to Classes loaded with application class loader first.
 - 8) Click **Apply**.
- v. Add a connection factory for the Content Engine application.
 - 1. Navigate to the Resource Adapter interface.
 - WebSphere 6.0:
 - a) From the WebSphere administrative console, click **Applications**.
 - b) Navigate to Enterprise Applications > FileNetEngine > Connector Modules > engine.rar > Resource Adapter
 - WebSphere 6.1:
 - a) From the WebSphere administrative console, click **Applications**.
 - b) Navigate to Enterprise Applications > FileNetEngine > Manage Modules > FileNet P8 connector > Resource Adapter
 - 2. Click **J2C Connection Factories**.
 - NOTE** One connection factory should already exist, this procedure will a second one.

3. Click **New**.
 4. For the name enter: FileNetConnectionFactory.
 5. For the value enter: FileNet/Local/ConnectionFactory.
 6. Set the three drop down boxes at the bottom of the screen to “none.”
 7. Click **Apply**.
- vi. Set Connection Pool properties.
- 1 Click the **Connection Pool Properties** link from the current page of the WebSphere Administrator console.
 - 2 Change the “Max” property to 100.
 3. Change the “Min” property to 10.
 4. Click **Apply**.
 5. Save the changes to the master repository.
 6. Synchronize all nodes.
- vii. Start the Content Engine application.
1. Expand the Applications link and select **Enterprise Applications** on the left pane of the WebSphere Administrator console.
 2. Select the check box next to FileNetEngine and click **Start**.

To deploy on JBoss

CAUTION As JBoss clusters communicate state over multicast it is critical that you configure clusters for Content Engine to segregate multicast traffic from each other and from other JBoss clusters. Instructions regarding JBoss cluster configuration can be found at <http://www.jboss.com/>

To deploy the Content Engine application to the cluster nodes, do the following:

- a. Leave Content Engine deployed on the application server where you installed it.
 - b. Copy the following files to each JBoss node you are setting up:
 - All data source (**<JBoss_HOME>/server/all/deploy/*-ds.xml**) files created as part of the install process.
- NOTE** Depending on how you installed JBoss there might be datasource files not created by the CE installer, such as **hsqldb-ds.xml**, in the same location. Do not copy these files.
- **<JBoss_HOME>/server/all/deploy/Engine-ds.xml**
 - The JDBC driver in **<JBoss_HOME>/server/all/lib**
 - **<JBoss_HOME>/server/all/conf/login-config.xml**
 - **<JBoss_HOME>/server/all/deploy/Engine-jbc.ear**

NOTE Clustered JBoss instances are deployed from the “all” directory and not the “default” directory. Instructions regarding JBoss cluster configurations can be found at <http://www.jboss.com>.

- c. Start or restart the JBoss application server.
- d. Verify that the application deployed successfully.

Verify that the **server.log** file located in `<JBOSS_home>/server/all/log` lists deployment of the WAR or EAR file you used.

To configure Content Engine post installation

1. Verify the GCD data sources and connection pools, and configure them for database failover.

Follow the instructions in “To verify the GCD data sources and connection pools” in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Complete Post-Install Content Engine Configuration.”

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

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

3. 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 “Install and Deploy Content Engine.”

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 173.

- 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 “Install and Deploy Content Engine.”

- 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 “Install and Deploy Content Engine.”

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.

4. Perform post-deployment configuration steps.

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

5. (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 4.0.x

To create object stores follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* tasks “Create Object Stores” and “Verify Content Engine Installation”.

General Information

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

- If you are using an Application Server Cluster configuration set the scope or deployment targets of the database resources you create for use with Content Engine as follows:
 - Set the scope to the Cell level when configuring JDBC providers and data sources in a WebSphere Network Deployment or Deployment Manager configuration.

Setting the scope to the cell level makes the resources you create available through the cell without having to create the same entry at every node.
 - Set the targets of the data sources to the WebLogic cluster in which Content Engine is deployed when configuring data sources and connection pools using WebLogic.
- **NOTE** Older versions of WebLogic will require you to set the targets for data sources and connection pools, in which case both should be set to the WebLogic cluster to which Content Engine is deployed.
- If you are using a highly available database cluster use the virtual name of the cluster when entering the server/host name during the creation of data sources for your database server.
- (WebSphere and WebLogic) Oracle RAC JDBC connection URL.

If you are using Oracle RAC for your back end 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>.

- WebLogic 8.x - Search for “Configuring and using MultiPools.”
- WebLogic 9.x - Search for “Configuring JDBC Multi Data sources.”

Enable Application Server Connections after Database Failover

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

- WebSphere
 - “To enable WebSphere connection to the object store database after failover” on page 48.
 - “To enable WebSphere connection to the GCD database after failover” on page 48.
- WebLogic
 - “To enable WebLogic connection to the object store database after failover” on page 49.
 - “To enable a WebLogic 8.1.x connection to the GCD database after failover” on page 49.
 - “To enable a WebLogic 9.2 connection to the GCD database after failover” on page 49.
- JBoss
 - “To enable JBoss connection to the GCD and object store databases after failover” on page 50.

To enable WebSphere connection to the object store database 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 5.x

Follow the instructions in the *WebSphere 5.1.x documentation* topic “[Ensuring applications obtain valid connections](#)” at www.ibm.com > Support & downloads > WebSphere > WebSphere Application Server > Information center > V 5.1.x > WebSphere Application Server Version 5.1.x Information Center.

Search for “Ensuring applications obtain valid connections” to find the instructions.
- WebSphere 6.x
 - a. Log on to WebSphere Administrative Console.
 - b. Navigate to JDBC Providers > JDBC Providers for GCD Datasource > Data sources > *Data Source Name* > WebSphere Application Server data source properties.
 - c. 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.
 - d. Save your changes.

To enable WebSphere connection to the GCD database after failover

To enable WebSphere 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 WebSphere Administrative Console.
2. Navigate to JDBC Providers > JDBC Providers for GCD Datasource > Data sources > *GCD Data Source Name* > WebSphere Application Server data source properties
3. Select 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. Save your changes.

To enable WebLogic connection to the object store database after failover

To enable WebLogic to connect to the object store 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 *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, navigate to Services > JDBC > Data Sources.
3. On the Summary of Data Sources page, click *Data Source Name*.
4. (WebLogic 8.1.x) Click the **Connections** tab, click **Advanced**, and then select the Test Reserved Connections check box.
5. (WebLogic 9.2) In the Configuration tab, click the Connection Pool tab, and then select the Test Connections on Reserve check box.
6. Save your changes.

To enable a WebLogic 8.1.x connection to the GCD database 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 left pane, expand Services > JDBC > Connection Pool to display the connection pools in the current domain.
3. Click the Configuration tab, and then click the Connections tab.
4. Click **Show** to display the advanced connection options.
5. Select the Test Connections On Reserve check box, and set a Test Table Name.
6. Save your changes.

To enable a WebLogic 9.2 connection to the GCD database 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* and click the **Configuration:Connection Pool** tab.
4. Click **Advanced** and select the Test Connections On Reserve check box.
5. 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. 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 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.

`<port>` is the port used to access CE on the load balancer or server running CE.

The “Startup Context” page displays certain data regarding CE, such as server uptime, product build number, and so on.

NOTE Depending on your configuration you can either 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, or you can connect manually to each CE node to verify that CE is running on the node. For JBoss, the second alternative must be used as JBoss only reports the name of the cluster and not the individual servers.

3. 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.”

4. 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 HA Configuration

Unlike previous versions FileNet P8 Content Engine version 4.0.x 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.

Installing a Highly Available Process Engine

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

Overview

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.

Procedure Overview for Process Engine

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

UNIX Platforms

Server Farm Environments

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

1. 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 security settings, language packs, and so on.
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 [“Server Farm \(UNIX, Windows\)” on page 57](#).

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.
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).
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. Synchronize the **vwtaskman.xml** file across your servers. For information, see “(Process Engine 4.0.1) Synchronize vwtaskman.xml for correct Process Engine high availability failover” on page 73
 9. Failover to each node to ensure that Process Engine resources startup properly on each node. For detailed instructions, see “Veritas Cluster Server (UNIX)” on page 59.

Windows Platforms

Server Farm Environments

See “Server Farm Environments” on page 54. For detailed instructions, see “Server Farm (UNIX, Windows)” on page 57.

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.
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).
6. Add resources to the cluster to control Process Engine.
 - IMS Control Service
 - Process Engine Services Manager
 - Registry replication

7. Synchronize the **vwtaskman.xml** file across your servers. For information, see [“\(Process Engine 4.0.1\) Synchronize vwtaskman.xml for correct Process Engine high availability failover”](#) on page 73
8. Failover to each node to ensure that Process Engine resources startup properly on each node.
For detailed instructions, see [“Veritas Cluster Server \(Windows\)”](#) on page 65 or [“Microsoft Cluster Server”](#) on page 69.

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. 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 security settings, language packs, and so on.

2. Configure the hardware load balancer.
 - a. Verify that your hardware load balancer is set up to allow ALL PROTOCOLs traffic, not just TCP.
 - b. For each Process Engine server, in the hosts file, set the hardware load balancer virtual name to the IP address of the PE server itself.

Example:

<loadbalancer_virtual_name> (10.15.4.48) is the load balancer virtual name, and <PE_farm_1> (10.15.4.67), <PE_farm_2> (10.15.4.106), and <PE_farm_3> (10.15.4.111) are three PE servers.

The hosts file entries should look like this:

Server	Hosts file entry
<PE_farm_1>	10.15.4.67 <PE_farm_1> <loadbalancer_virtual_name>
<PE_farm_2>	10.15.4.106 <PE_farm_2> <loadbalancer_virtual_name>
<PE_farm_3>	10.15.4.111 <PE_farm_3> <loadbalancer_virtual_name>

CAUTION (UNIX) Verify that the nsswitch.conf file is configured to resolve names via the local hosts 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.

3. 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.
CAUTION As farm nodes act as independent servers they must each have a unique NCH name.
 - When running SQL scripts choose 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 4.0.0 installation.
 - You do not need to perform the task “Configure Process Task Manager” as this was done when you configured the first Process Engine.
4. Start Process Engine on the additional farmed servers.
5. Refresh the Process Engine node on the Process Task Manager.
You should now see all the farmed PE servers under the servers node.
6. Synchronize the **vwtaskman.xml** file across your servers. For information, see [“\(Process Engine 4.0.1\) Synchronize vwtaskman.xml for correct Process Engine high availability failover”](#) on page 73

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 fnsw 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 181](#).

Pre-Deployment Tasks

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.

Deployment Tasks

1. Review the Notes section for the platform specific Process Engine installation task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

2. 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 /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).
 - 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 /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).
 - 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).
3. 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 “To install the Process Engine interactively” or “To install the Process Engine silently” found in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - b. Once the installation completes perform system specific post-configuration steps on each system:
 - AIX
 - “(Optional: Oracle databases only) To remove fnsw from the Oracle database administrators group.”
 - “To restore any custom modifications for root and fnsw users.”
 - “To verify /etc/services file settings.”
 - “To clean before starting Process Engine.”
 - “To edit the /etc/inittab file.”

- HPUX
 - “(Optional: Oracle databases only) To remove fnswh from the Oracle database administrators group.”
 - “To restore any custom modifications for root and fnswh users.”
 - “To verify /etc/services file settings.”
 - “To clean before starting Process Engine.”
 - “To edit the /etc/inittab file.”
 - Solaris
 - “(Optional: Oracle databases only) To remove fnswh from the Oracle database administrators group.”
 - “To enable ports.”
 - “To restore any custom modifications for root and fnswh users.”
 - “To verify /etc/services file settings.”
 - “To clean before starting Process Engine.”
 - “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 /fnswh/local. Repeat the installation procedure starting at [Step a on page 61](#) above.
- NOTE** Repeat, one server at a time, for all Process Engine nodes.
4. Configure Process Engine for the cluster.
- NOTE** Do 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.
5. Enable automation 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.

6. Configure the Process Engine using the highly available Content Engine URL. For more information, see [“Connecting Using the Web Service Transport” on page 175](#).

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

7. Install the FileNet P8 4.0.0 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 *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/VRTSvcs/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.
8. Synchronize the **vwtaskman.xml** file across your servers. For information, see [“\(Process Engine 4.0.1\) Synchronize vwtaskman.xml for correct Process Engine high availability failover” on page 73](#)

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.
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.

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.

Pre-Deployment Tasks

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.

Deployment 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 170](#).

1. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Install Process Engine (Windows).”

NOTE Run the Process Engine installer as a domain user created for installing Process Engine software as detailed in [“Create FileNet Users.” on page 170](#). You can also use the domain “fnsw” 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.
 - 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.
2. 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 the Web Service Transport” on page 175](#). For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure the Process Task Manager.”
 3. 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

4. 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.
5. Bring the cluster group online.
 6. Synchronize the **vwtaskman.xml** file across your servers. For information, see [“\(Process Engine 4.0.1\) Synchronize vwtaskman.xml for correct Process Engine high availability failover” on page 73](#)

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.

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.

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.

Pre-Deployment Tasks

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.

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 170](#).

1. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task *Install Process Engine (<Operating System>)*.

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 170](#).

- 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.

2. 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 the Web Service Transport" on page 175](#). For more information, see the *IBM FileNet P8 Platform Installation and Upgrade Guide* task Configure the Process Task Manager.

3. 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

4. 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.

5. Synchronize the **vwtaskman.xml** file across your servers. For information, see [“\(Process Engine 4.0.1\) Synchronize vwtaskman.xml for correct Process Engine high availability failover”](#) on page 73

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.

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.

(Process Engine 4.0.1) Synchronize vtaskman.xml for correct Process Engine high availability failover

The **vtaskman.xml** file is used by Process Engine to get configuration details regarding the RMI registry port, and the "Automatically Started" setting. The file is stored on the PE server in the **/fnsw/bin** folder (UNIX and Windows). For correct failover of a highly available Process Engine this file must be identical on all servers.

To enable correct failover you must synchronize the **vtaskman.xml** file across all PE nodes after the initial configuration, and after each Process Engine configuration change.

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 <fnsw> 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 <fnsw> 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.

Installing a Highly Available Application Engine/Workplace XT

CAUTION 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 *FileNet P8 Hardware and Software Requirements* and the *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

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

NOTE The procedure for making AE and Workplace XT highly available are similar. Both procedures are included in this chapter, with differences highlighted where appropriate.

Overview

Application Engine/Workplace XT provides front-end access to FileNet P8 content in an application server cluster. 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 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.

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

This section covers the tasks and requirements for making these components highly available.

Requirements and General Information

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 your configuration.

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

CAUTION 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 so 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:

- For Application Engine:

```
http://<loadbalancer URL>:<load balancer port>/Workplace
```

- For Workplace XT:

```
http://<loadbalancer URL>:<load balancer port>/WorkplaceXT
```

Use this URL when completing the installation steps.

Workplace/Workplace XT Settings

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

CAUTION

- If your Application Engine/Workplace XT is deployed 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 Application Engine/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.

Shared Storage/Configuration Directories

- 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).

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.

Procedure overview for Application Engine/Workplace XT

To install Application Engine/Workplace XT in a highly available environment

High-Level Steps

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.

The WebSphere and WebLogic implementations in this section are based on these steps.

Pre-Deployment Tasks

1. (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.

2. 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 [“Managing Workplace/Workplace XT Settings in a Highly Available Environment” on page 85](#).

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.

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.

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.

Deployment Tasks

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

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task Install 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 the EJB Transport” on page 174.](#)
- (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.

2. 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.](#)

3. 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 83.](#)

- (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 deploy 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 (<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.

4. Configure Server Ports

Perform the “To configure the server ports” section in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task “Configure and Deploy 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.

5. 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 Application Engine 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.

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 this 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.

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.

Service Pack, Fix Pack and Interim Fix Installation Tasks

This section provides the information needed to install FileNet P8 4.0 service packs, fix packs, and interim fixes in an existing FileNet P8 highly available environment.

Install any interim fixes required for the IBM FileNet P8 platform. To determine whether additional interim fixes are needed, contact your service representative.

WARNING Do NOT install a service pack until all your FileNet P8 functional expansions, such as Records Manager, support it. See the most recently released Service Pack edition of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

In addition to this technical notice and the documents listed in [“Access IBM FileNet Documentation” on page 13](#) you will need the following documents which provide instructions on how to apply the IBM FileNet P8 platform service packs:

- *Content Engine 4.0.x Service Pack or Interim Fix Readme* - included with the Content Engine service pack or interim fix software.
- *Process Engine 4.0.x Service Pack or Interim Fix Readme* - included with the Process Engine fix pack or interim fix software.
- *Application Engine 4.0.x Service Pack or Interim Fix Readme* - included with the Application Engine fix pack or interim fix software.

To install the FileNet P8 Service Packs or Interim fixes

CAUTION

The installation order for service packs and interim fixes across the IBM FileNet P8 platform differs depending on your setup and if you are installing service packs or interim fixes as part of a clean installation or an upgrade.

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.

See the *IBM FileNet P8 Platform Installation and Upgrade Guide* for instructions, with additional comments and instructions for HA environments documented in this guide as listed below.

- For Content Engine, see [“Applying Content Engine Service Packs or Interim Fixes” on page 89](#).
- For Process Engine, see [“Applying Process Engine Service Packs or Interim Fixes” on page 90](#).
- For Application Engine/Workplace XT, see [“Applying Application Engine/Workplace XT Service Packs or Interim Fixes” on page 92](#).

Applying Content Engine Service Packs or Interim Fixes

This section covers the installation of the Content Engine service packs or interim fixes in a highly available environment.

CAUTION The installation order for service packs and interim fixes across the IBM FileNet P8 platform differs depending on your setup and if you are installing service packs or interim fixes as part of a clean installation or an upgrade.

To install and configure Content Engine 4.0.x service packs or interim fixes in an HA environment, follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the HA specific modifications noted below.

Fix Pack/Interim Fix Installation Tasks

Follow the instructions listed in the "Installation Instructions" section of the *Content Engine 4.0.x Service Pack* or *Interim Fix Readme* - included with the Content Engine fix pack or interim fix software.

NOTES

- For an Application Server Cluster, install the service pack or interim fix on the administrative server and re-deploy the CE application to a cluster instead of to the standalone node server.
- (Service packs/interim fixes requiring installation of Content Engine Java Client files) Install the Content Engine Java Client files depending on your PE and AE setup:

For Process Engine, install the Content Engine Java Client files on all standalone node servers.

For Application Engine

- For a farm of Independent Application Server Instances install the Content Engine Java Client files on all standalone node servers.
- For an Application Server Cluster, install the Content Engine Java Client files on the administrative server only.

For more information about Content Engine Java Client files, see the *Content Engine 4.0.x Service Pack* or *Interim Fix Readme*.

To re-deploy Content Engine after successful installation

After successfully installing the service pack, fix pack, or interim fix, follow the instructions below to re-deploy your CE application.

- ["To deploy on WebSphere" on page 41.](#)
- ["To deploy on WebLogic" on page 40.](#)

Applying Process Engine Service Packs or Interim Fixes

This section covers the installation of Process Engine service packs, fix packs, or interim fixes in a highly available environment.

CAUTION The installation order for service packs and interim fixes across the IBM FileNet P8 platform differs depending on your setup and if you are installing service packs or interim fixes as part of a clean installation or an upgrade.

To install and configure Process Engine 4.0.x service packs or interim fixes in an HA environment, follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the HA specific modifications noted below.

NOTE (Service packs, fix packs, interim fixes requiring the Process Engine Client Installer) Run the Process Engine Client Installer depending on your AE setup:

- For a farm of independent application server Instances run the Process Engine Client Installer on all standalone node servers.
- For an Application Server Cluster, run the Process Engine Client Installer on the administrative server only.

For more information about Process Engine Client Installer, see the *Process Engine 4.0.x Service Pack* or *Interim Fix Readme*.

Service Pack/Fix Pack/Interim Fix Installation Tasks - Server Farms

1. Stop all Process Engine services and applications.
2. Install the Process Engine service pack in the *Interim Fix Readme*.
CAUTION Perform tasks involving installation to and configuration of components shared by all PE nodes, such as the PE database, when installing the first node only. Do not reconfigure these components when installing on the second and consecutive nodes.
3. Once the service pack installation is finished on the node, restart PE on the current node.
4. Repeat steps 2-3 for all nodes that run Process Engine services until the service pack or interim fix has been completely installed on all nodes in the farm.
5. Start Process Engine services.

Service Pack/Fix Pack/Interim Fix Installation Tasks - UNIX Clusters

1. Suspend cluster operations on the active node so that the cluster groups do not failover when services are stopped/started during the service pack or interim fix installation. The specific procedure for this varies by cluster vendor (for example VERITAS can "freeze" a node).
2. Stop all Process Engine services and applications.

3. Install the Process Engine service pack or interim fix in the cluster by following the UNIX specific steps under in the "Installation Instructions" section of the *Process Engine 4.0.x Service Pack* or *Interim Fix Readme*.

CAUTION Perform tasks involving installation to and configuration of components shared by all PE nodes, such as the PE database, when installing the first node only. Do not reconfigure these components when installing on the second and consecutive nodes.

4. Once the installation is finished on the node, restart PE and resume cluster operations on the current node. The specific procedure for this varies by cluster vendor (for example VERITAS can "unfreeze" a node).
5. Move the cluster group to another node.
6. Repeat steps 3-5 for all nodes in the cluster that run Process Engine services until the service pack or interim fix has been completely installed on all nodes in the cluster.
7. Start Process Engine services.

Service Pack/Fix Pack/Interim Fix Installation Tasks - Windows Clusters

1. Suspend cluster operations on the active node so that the cluster groups do not failover when services are stopped/started during the service pack or interim fix installation. The specific procedure for this varies by cluster vendor (for example MSCS can "Pause" a node).
2. Stop all Process Engine services and applications.

Install the Process Engine service pack or interim fix in the cluster by following the Windows specific steps under in the "Installation Instructions" section of the *Process Engine 4.0.x Service Pack*, *Service Pack*, or *Interim Fix Readme*.

CAUTION Perform tasks involving installation to and configuration of components shared by all PE nodes, such as the PE database, when installing the first node only. Do not reconfigure these components when installing on the second and consecutive nodes.

3. Once the installation is finished on the node, restart PE and resume cluster operations on the current node. The specific procedure for this varies by cluster vendor (for example MSCS can "Resume" a node).
4. Move the cluster group to another node.
5. Repeat steps 3-5 for all nodes in the cluster that run Process Engine services.
6. Start Process Engine services.

Applying Application Engine/Workplace XT Service Packs or Interim Fixes

This section covers the installation of Application Engine/Workplace XT service packs or interim fixes in a highly available environment.

CAUTION The installation order for service packs and interim fixes across the IBM FileNet P8 platform differs depending on your setup and if you are installing service packs or interim fixes as part of a clean installation or an upgrade.

To install and configure Application Engine/Workplace XT service packs or interim fixes in an HA environment, follow the order as documented in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the HA specific modifications noted below.

Service Pack/Fix Pack/Interim Fix Installation Tasks

Follow the instructions listed in the "Installation Instructions" section of the *Application Engine 4.0.x Service Pack* or *Interim Fix Readme* - included with the Application Engine service pack or interim fix software.

Follow the instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task "Install Application Engine Software Updates," or the *IBM FileNet Workplace XT Installation Guide* task "Install IBM FileNet Workplace XT Software Updates."

NOTES

- You must install the service pack or interim fix on each node to update the Component Manager.
- For a farm of Independent Application Server Instances re-deploy the AE application the standalone node servers.
- For an Application Server Cluster re-deploy the AE application to a cluster instead of to the standalone node server.

To re-deploy Application Engine/Workplace XT after successful installation

After successfully installing the service pack or interim fix, follow the instructions in "[Configure and deploy the Application Engine/Workplace XT for your application server configuration.](#)" on page 81 to re-deploy your AE application.

Functional Expansion Components Installation Tasks

This section covers high availability installation and configuration of FileNet functional expansion components. FileNet functional expansions can be installed and configured for high availability in any order.

CAUTION Before attempting to install and configure these components, make sure your FileNet P8 core environment (CE, PE, and AE) is up and running.

To install FileNet functional expansions components in a highly available environment

- [“Installing a Highly Available Records Manager” on page 95.](#)
- [“Installing a Highly Available Rendition Engine” on page 101.](#)
- [“Installing a Highly Available BPF” on page 104.](#)
- [“Installing a Highly Available CFS-IIICE” on page 112.](#)
- Connector for SharePoint, Version 2.1
 - [“Installing a Highly Available Connector for SharePoint Web Parts 2.1” on page 118.](#)
- Connector for SharePoint, Version 2.2
 - [“Installing a Highly Available IBM FileNet Connector for SharePoint Web Parts 2.2” on page 121.](#)
 - [“Installing a Highly Available IBM FileNet Connector for SharePoint Document Libraries 2.2” on page 126.](#)
- [“Installing a Highly Available IBM FileNet Services for Lotus Quickr” on page 135.](#)

Additional Documentation

NOTE For general instructions on how to navigate to this and other IBM FileNet product documentation on the IBM web site, see [“Access IBM FileNet Documentation” on page 13.](#)

Depending on your setup you will need additional documentation to set up your FileNet Functional Expansions for high availability:

- *FileNet Records Manager Installation and Upgrade Guide (PDF).*
- *FileNet Rendition Engine Installation and Upgrade Guide (PDF).*
- *IBM FileNet Business Process Framework Installation Guide (PDF).*
- *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide (PDF).*
- *High availability deployment scenario for WebSphere Information Integrator Content Edition (PDF).*
- *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.1 (PDF).*

- *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.2 (PDF).*
- *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2 (PDF).*
- *IBM FileNet Services for Lotus Quickr Installation and Upgrade Guide, Version 1.0 (PDF)*

Installing 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

You can install FileNet P8 Records Manager version 4.0 or 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 4.0 or 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 [“Installing a Highly Available Content Engine” on page 34](#) and [“Installing a Highly Available Process Engine” on page 53](#).
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 [“Installing a Highly Available Application Engine/Workplace XT” on page 76](#).

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.

NOTE The WebSphere and WebLogic implementations in this section are based on these steps.

Records Manager Installation and Deployment Tasks

Pre-deployment 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. Verify that the cluster nodes and load balancers are started if the application server used to cluster Records Manager is WebLogic or WebSphere.

2. (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.

3. 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. FileNet recommends hosting this file on a highly available share exclusive of the Application Engine/Records Manager cluster.

Deployment 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.

NOTE 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.
 - For WebSphere, refer to the *Deploy the Records Manager Web Application* task, *IBM WebSphere 5.x and 6.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.
 - For WebLogic, refer to the *Deploy the Records Manager Web Application* task, *BEA WebLogic 8.1.x and 9.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.
4. (DoD Chapter 4/DoD Classified only) Create the Classified Records Security Personnel Access Role.

To successfully set up the Classified Records Security Personnel Access Role for a highly available Records Manager you must run the RMAccessRole tool on one Records Manager node once only per highly available Records Manager environment.

The tool uses 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 Content Engine server, not to a highly available URL. To achieve this you must temporarily reconfigure the `WcmApiConfig.properties` file to point to a single Content Engine node when running the tool.

To run the 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 one of the Content Engine nodes 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 173](#)

Example:

For WebSphere, change the value from

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

to

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

- c. Run the `RMAccessRole` tool on your Records Manager server.

Follow the instructions in the “(DoD Chapter 4/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_number>/<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.

Installing a Highly Available Rendition Engine

You can install FileNet Rendition Engine in highly available configurations using the standard Rendition Engine installer. No extra software is needed to set up a highly available Rendition Engine.

Overview

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

Each RE, controlled by the Vista domain, polls the RE database for submitted rendition jobs in a round-robin fashion. If a job is available, the Vista domain assigns it to one of the available RE servers. In this scenario the 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

- FileNet P8 relies on the built in Lipient high availability for RE 4.0.x.
CAUTION This functionality is new for RE 4.0.1, 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.

Procedure Overview for Rendition Engine

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.

Pre-installation Tasks

1. Install Content Engine, Process Engine, and Application Engine. Follow the instructions in the *FileNet P8 Installation and Upgrade Guide*.

For high availability installation of these components, see [“Installing a Highly Available Content Engine” on page 34](#), [“Installing a Highly Available Process Engine” on page 53](#), and [“Installing a Highly Available Application Engine/Workplace XT” on page 76](#).

2. Review the Installation Planning Considerations in the *FileNet Rendition Engine Installation and Upgrade Guide*.

NOTE 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.

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

- (SQL Server only) Do task “Install Microsoft SQL Server for Rendition Engine.”

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

- (Oracle only) Do task “Install Oracle Server for Rendition Engine.”

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

Install and configure Rendition Engine

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

1. Install and configure the first Rendition Engine.

- a. Do task “Install Rendition Engine.”

In the “Add Server to an Existing Vista Domain” screen, select “Leave as a stand alone Server” to create a new Vista Domain.

- b. Do task “Configure Rendition Engine Connection 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. Do task “Set the site preferences on the Application Engine to use path-based URLs.”

NOTE You only need to do this once per HA environment.

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

NOTE Do this on each Rendition Engine server.

- e. Do task “Install Service Packs, Fix Packs and Interim Fixes for Rendition Engine.”

CAUTION Before you can verify your installation and use the rendering services you must install the required service pack to bring Rendition Engine to the minimally supported level for high availability (RE 4.0.1).

- f. Do 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.

2. Install additional Rendition Engines.

- a. Do task “Install Rendition Engine.”

In the “Add Server to an Existing Vista Domain” screen, select “Add to an Existing Vista domain” to add the new server to the Vista Domain you created with the first Rendition Engine.

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. Do task “Enable word processing macros for rendering.”
- NOTE** Do this on each Rendition Engine server.
- c. Do task “Install Service Packs, Fix Packs and Interim Fixes for Rendition Engine.”
- CAUTION** Before you can verify your installation and use the rendering services you must install the required service pack to bring Rendition Engine to the minimally supported level for high availability (RE 4.0.1).
- d. Do the optional task “(Optional) Verify Installation.”

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

Installing a Highly Available BPF

CAUTION This topic covers installation of Business Process Framework (BPF) 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

Business Process Framework is an optional component of IBM FileNet P8 Platform.

You can install BPF in highly available configurations using standard application server clustering software. The software is tightly 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 BPF Web Application only.
- When installing BPF Explorer in an environment with clustered databases, make sure you configure it to use the cluster’s virtual IP address for the ODBC data source.
- The following procedure covers the specific details required for installing BPF 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.

Procedure Overview for BPF

To install BPF in a highly available environment

High-Level Steps

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

1. Install Content Engine and Process Engine.
 Follow the instructions in the *FileNet P8 4.0.x Installation and Upgrade Guide*. For high availability installation of these components, see [“Installing a Highly Available Content Engine” on page 34](#) and [“Installing a Highly Available Process Engine” on page 53](#).
2. Install a highly available Application Engine. For information, see [“Installing a Highly Available Application Engine/Workplace XT” on page 76](#).
3. Configure your BPF environment.
 See [“Configure your BPF environment.” on page 106](#) and [“Create DataSources.” on page 106](#).
4. Install BPF on the first (AE) node.
 This configures the IBM FileNet P8 platform components for BPF. See [“Install and Configure the BPF Web Application on the first node.” on page 106](#).
5. Install BPF on the remaining (AE) nodes.
 See [“Install the BPF Web Application on additional nodes.” on page 108](#).
6. Deploy BPF in the cluster.
 See [“Manually Deploy the BPF Web Application.” on page 109](#).
7. Perform any necessary post-deployment steps.
 See [“Configure BPF” on page 111](#).
8. Make the DataSources highly available.
 See [“Make the DataSources highly available after failover.” on page 111](#).

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. (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.

2. Configure your BPF environment.

Follow the instructions in the “Configuration of the environment before BPF server installation” task in the *IBM FileNet Business Process Framework Installation Guide* to complete the following tasks:

- Prepare the Bootstrap Object Store.
- Configure Users and Groups.
- Edit Workplace Site Preferences.
- Create new Metastore database.
- On each AE node:
 - Install JDBC Drivers.
 - Copy CD contents to target machine.
 - Making staging folder accessible.

3. Create DataSources.

Create DataSources on the AE for the PE and Metastore databases. Follow the instructions in the “Create JDBC DataSource – BPF Metastore” and “Create JDBC DataSource – Process Engine” tasks in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*.

CAUTION When configuring the DataSources, make sure you use the virtual IP of the database cluster for the “Enter Host name: {meta data server}” value.

Installation Tasks

For more information about BPF installation, see the *IBM FileNet Business Process Framework Installation Guide*.

CAUTION Installation requires that the Application Engine and BPF be co-located on the same application server. For more information, see the “BPF Installation and Upgrade Overview” topic in the *IBM FileNet Business Process Framework Installation Guide*.

1. Install and Configure the BPF Web Application on the first node.

Follow the instructions in the “BPF server installation with automatic configuration” task in the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- In the Choose installation type screen, select “Typical.”
- In the Automatic Metastore creation screen, select “NO.”
- In the BPF Explorer XML Manifest screen, select “Manual import of XML Manifest.”

NOTE You will perform this import as part of [“Configure BPF” on page 111](#).

- WebSphere
 - In the WebSphere server access info screen, choose:
 - The Deployment Manager's profile
 - The Cell
 - The CellManager Node
 - The Deployment manager server instance
 - The Deployment Manager's SOAP port
 - The local node's HTTP Port
- WebLogic
 - In the WebLogic information screen, enter:
 - WebLogic Domain Name
 - Admin Server Name
 - Admin Server Port
 - Admin Login Name
 - Password
- JBoss
 - In the JBoss information screen, enter:
 - JBoss Home Directory
 - JBoss Server Type - all
 - HTTP Port
- In the Deployment to web app server screen choose "No" to manually deploy the application.
- In the Datasources screen choose "Use existing." for both PE Datasource and Metastore Datasource, and point to the existing PE DataSource and Metastore DataSource you created in ["Configure your BPF environment."](#) on page 106 above.
- In the Content Engine WSI access screen enter:
 - Content Engine Client Software URL: Point to a single CE server node using the correct host name and port number.
- In the Content Engine EJB access screen enter:
 - 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 the EJB Transport"](#) on page 174.
- In the Workplace Configuration Files screen:
 - Deployed Workplace Directory - Enter the Workplace installed location.

- Workplace Config Directory - The shared AE bootstrap location
- In the Workplace URL screen, specify:
 - Workplace URL: the proxy or load balanced URL for Workplace

2. Install the BPF Web Application on additional nodes.

Follow the instructions in the “BPF server installation with automatic configuration” task in the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- In the Choose installation type screen, select “Custom” and make sure only the following components are selected:
 - BPF Metastore
 - BPF Web Application
 - BPF Operations - Component Manager Component
- In the Automatic Metastore creation screen, select “NO.”
- In the BPF Explorer XML Manifest screen, select “Manual import of XML Manifest.”
- WebSphere

In the WebSphere server access info screen, configure the following:

- Profile - select the local application server profile (e.g., AppSrv01).
- Cell - select the same value as for the first node.
- Node - select the value for the local application server's node.
- Server - select the local application server instance (e.g., server1).
- SOAP Port - enter 8880 (the SOAP port for the local application server node).

CAUTION In a clustered environment the installer must be able to connect to a SOAP port during the installation process, or it might fail with an exception. By configuring the installation for a local SOAP port the installer will fail but will let you continue the installation.

- HTTP Port - enter the local application server HTTP port (e.g., 9080)
- WebLogic
 - In the WebLogic information screen, enter the same values you entered for the first node.
- JBoss
 - In the JBoss information screen, enter the same values you entered for the first node.
- In the Deployment to web app server screen choose “No” to manually deploy the application.
- In the P8 administrative account verify that the user name and password fields are blank.
- In the Content Engine WSI access screen verify that all fields are blank.

- In the Content Engine EJB access screen, enter the same value you entered for the first node.
- In the Workplace Configuration Files screen:
 - Deployed Workplace Directory - Enter the Workplace installed location.
 - Workplace Config Directory - The shared AE bootstrap location
- In the Workplace URL screen, enter the same values you entered for the first node.
- In the Process Engine access screen verify that all fields are blank.

CAUTION When running the installer with the above fields left blank a number of non-critical exceptions will display in the Installation Summary screen. These all relate to configuration entries that were not made and that you plan to perform manually or have already performed as part of the initial node installation, and can be ignored.

3. Manually Deploy the BPF Web Application.

Follow the instructions in the “Deploy BPF Web Application” task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- Application server cluster
 - WebSphere.

Configure and deploy the BPF Web Application from the administrative server instance.

NOTE When installing BPF through the administrative server make sure to deploy the BPF to the cluster of server instances.

Follow the instructions in the “Deploy BPF Web Application” task “IBM WebSphere <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- 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.

Configure and deploy the BPF Web Application from the administrative server instance.

NOTE When installing BPF through the administrative server make sure to deploy the BPF to the cluster of server instances.

Follow the instructions in the “Deploy BPF Web Application” task “BEA WebLogic <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM*

FileNet Business Process Framework Installation Guide, with the following modifications:

- Make sure the targets for the application include the WebLogic cluster and/or managed nodes that will be running the application.
- JBoss

Follow the instructions in the “Deploy BPF Web Application” task “JBoss 4.0.x” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - Copy the web application files from the first node to each additional node.
 - Copy the datasource XML files from the first node each additional node.
 - Restart JBoss on each node.
- Farm of independent application server instances
 - WebSphere.

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “IBM WebSphere <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - For WebSphere make sure you select the cluster and/or managed nodes that will be running the application when mapping modules to servers. The Deployment Manager 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 If you are using an IBM HTTP Server to proxy/load-balance application requests to WebSphere, make sure to select the webserver node as well as the application servers when mapping modules to the application.
 - WebLogic.

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “BEA WebLogic <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - Make sure the targets for the application include the WebLogic cluster and/or managed nodes that will be running the application.
 - JBoss

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “JBoss 4.0.x” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- Copy the web application files from the first node to each additional node.
- Copy the datasource XML files from the first node each additional node.
- Restart JBoss on each node.

4. Configure BPF

Follow the instructions in the “Post-install BPF server configuration” tasks in the *IBM FileNet Business Process Framework Installation Guide*.

NOTE This task includes importing the XML Manifest as well as updating the Workplace app_server.war/ear file and redeploying Workplace.

5. Make the DataSources highly available after failover.

Follow the instructions in “[Enable Application Server Connections after Database Failover](#)” on [page 48](#), substituting XA/non-XA DataSources with BPF Metastore and Process Engine DataSources as appropriate.

Installation Verification

To verify the installation follow the instructions in the “Installation Validation” task in the *IBM FileNet Business Process Framework Installation Guide*.

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

1. Using the application server administration tool, verify that the BPF application is running.
2. Open a web browser and type in the URL for the BPF application.
 - If using a load-balancer or proxy server the URL has the form:
`http://<virtual_name>:<port_number>/BPF`
 - If connecting to an individual cluster node the URL would have the form:
`http://<clustered_server_name>:<port_number>/BPF`
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.

Installing a Highly Available CFS-IICE

CAUTION This topic covers installation of Content Federation Services for Information Integrator Content Edition (CFS-IICE) 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

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

You can install CFS-IICE 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. For installation of IICE you will follow the instructions provided in *High availability deployment scenario for WebSphere Information Integrator Content Edition*. For information, see [“Install and configure IICE on each node in the cluster.” on page 114](#).
- 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.

Procedure Overview for CFS

The following high-level steps are necessary to configure CFS 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 [“Installing a Highly Available Content Engine” on page 34](#) and [“Installing a Highly Available Process Engine” on page 53](#).

2. Install a highly available Application Engine. For information, see [“Installing a Highly Available Application Engine/Workplace XT” on page 76](#).
3. Install and configure an application server cluster for CFS.
4. Install and configure IICE in the cluster.

Follow the instructions provided in the *High availability deployment scenario for WebSphere Information Integrator Content Edition* document you retrieved as part of [“Gather Auxiliary Documentation” on page 14](#).

5. Configure the IICE/CFS environment.

See [“Configure the IICE/CFS environment.” on page 114](#).

6. Install and Deploy CFS.

See [“Install and Deploy CFS.” on page 115](#).

7. Perform Post Installation Steps.

See [“Perform Post Installation Steps.” on page 116](#).

CFS Installation and Deployment Tasks

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-IICE.
2. Verify that the cluster nodes and load balancers are started.

Installation Tasks

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

1. Install and configure IICE on each node in the cluster.

Follow the instructions provided in the *High availability deployment scenario for WebSphere Information Integrator Content Edition* document you retrieved as part of "[Gather Auxiliary Documentation](#)" on page 14.

2. Configure the IICE/CFS environment.

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

- "Install Updated IICE Connectors" - Do this on all nodes.
- "Create Session Pools"
- "Configure HTTP Streaming" - Do this on all nodes.
- "Create Content Engine to IICE Connector"

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

WebSphere:

```
Remote Server URL: cemp:corbaloc::<node1 hostname>:<bootstrap port>,:<node2  
hostname>:<bootstrap port>/cell/clusters/<your websphere cluster name>/  
FileNet/Engine
```

```
Remote Server Upload URL: cemp:corbaloc::<node1 hostname>:<bootstrap  
port>,:<node2 hostname>:<bootstrap port>/cell/clusters/<your websphere  
cluster name>/FileNet/Engine
```

```
Remote Server Download URL: 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).

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
```

- “Configure Direct Mode for Content Engine (<application server>)” - Do this on all nodes.

WebSphere:

Use the WebSphere administrative console to configure the Generic JVM parameter for Content Engine for direct mode, using the following parameters:

```
-Dvbr.as.operationMode=direct
-Dvbr.as.direct.rmi.urls=rmi://<node 1 hostname>:<bootstrap port>/
ConfigurationServer,rmi:// <node 2 hostname>:<bootstrap port>/
ConfigurationServer
```

WebLogic:

Edit the startWebLogic script and add the following parameters:

```
JAVA_OPTIONS="{JAVA_OPTIONS} -Dvbr.as.operationMode=direct"
JAMA_OPTIONS="{JAVA_OPTIONS} -Dvbr.as.direct.rmi.urls=rmi:// <node1
hostname>:<bootstrap port>/ConfigurationServer,rmi:// <node2
hostname>:<bootstrap port>/ConfigurationServer"
```

- “Start RMI Bridges” - Do this on all nodes.
- “Create the CFS Federator Database and Database User Account (<db>)”

3. Install and Deploy CFS.

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

- “Install CFS”
 - Choose HA options where available.
 - In the IICE Configuration panel set the SSO Server URL and Configuration Server URL as follows:

```
IICE SSO Server URL: rmi:// <node 1 hostname>:<bootstrap port>/SSOserver,rmi://
/ <node2 hostname>:<bootstrap port>/SSOserver
```

```
IICE Configuration Server URL: rmi://<node 1 hostname>:<bootstrap port>/
ConfigurationServer,rmi:// <node 2 hostname>:<bootstrap port>/
ConfigurationServer
```

- “Configure CFS Federator Database (<db>)”

- “Deploy CFS Federator Administration application (WebSphere)”
 - Deploy to the iiceCluster - Do this on the administrative server.
 - NOTE** When deploying the CFS application, make sure you select the iiceCluster and webserver in the Map Modules to Servers screen.
 - Do step 2. Set the jdbc.drivers application custom property on all nodes.
 - Do step 3. Set the application classpath on all nodes.
 - “Deploy CFS Federator Administration application (WebLogic)”
 - Deploy to the iiceCluster - Do this on the administrative server.
 - Do step 2. Edit the following application server startup script for the IICE domain on all nodes.
 - “Start and Stop CFS Exporter” - Do this on all nodes.
 - “Create Data Source for CFS Federator Database“ - Do this on the administrative server.
 - “Configure Content Engine to Support CFS“
 - “Start and Stop CFS Importer“
4. Synchronize Connector, Data Maps, and Session Pools across the nodes.
- After configuring the IICE 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 IICE Admin Tool. Refer to the IICE documents for more details on the export / import options.
- NOTE** Make sure the JAR files in the IICE_HOME/lib directory are identical on both nodes.
- CAUTION** You must do this each time you make changes to your environment, such as “Set Up Additional External Repositories” below.
5. Perform Post Installation Steps.
- Perform the following tasks as outlined in the *IBM FileNet P8 Content Federation Services Installation and Upgrade Guide* with the following modifications:
- “Set Up Additional External Repositories”
 - NOTE** After making configuration changes like these you must redo [Step 4 “Synchronize Connector, Data Maps, and Session Pools across the nodes.”](#) on page 116.
 - “Start Up and Shut Down CFS”

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 IICE 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.

Installing a Highly Available Connector for SharePoint Web Parts 2.1

CAUTION This topic covers installation of Connector for SharePoint Web Parts 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

You can install Connector for SharePoint Web Parts 2.1 in a highly available SharePoint environment.

Before you begin installing the software, retrieve the following document:

- *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.1* (PDF)

To download this document from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

NOTES

- The following procedure covers the specific details required for installing Connector for SharePoint Web Parts 2.1 in an existing SharePoint farm rather than on a standalone server.
- You can configure Connector for SharePoint Web Parts 2.1 to connect to a highly available Content Engine or Process Engine, see [“Install IBM FileNet Connector for SharePoint Web Parts 2.1 on each farm node.” on page 119](#).
- IBM FileNet Connector for SharePoint Document Libraries 2.1

The following procedure covers installation of the Connector for SharePoint Web Parts 2.1 only. Connector for SharePoint Document Libraries doesn't have high availability support and can only be active on one server.

When installing Connector for SharePoint Document Libraries 2.1 in a highly available SharePoint environment you can optionally:

- Install the component on one node of your SharePoint farm only.

CAUTION This will introduce a single point of failure if that node goes offline.

- Install the component on all nodes of the SharePoint farm, configure them identically but only activate the component on one of the nodes. If the active node goes offline you can activate the Connector for SharePoint Document Libraries 2.1 on one of the other nodes.

Installation Tasks

To install SharePoint Web Parts 2.1 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.”

2. Install IBM FileNet Connector for SharePoint Web Parts 2.1 on each farm node.

For each node, follow the instructions in the *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.1* with the following modifications:

- a. Create a new User Profile Database for the first node you configure.
- b. For each consecutive node, make sure you enter the following User Profile Database Information exactly as you entered it for the first node:

- Database Instance
- Database Name
- Database Port
- Database User
- Database Password
- Encryption Key

CAUTION Each SharePoint node must use the same crypto key.

- c. For each consecutive node, make sure the information you enter for “Specify Web Part Properties” is identical.

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.

3. Verify that the encryption key specified in the web.config file is identical for all nodes.

- a. Open the **web.config** file on each 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 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.

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 Web Parts connects to FileNet P8, and that you can browse your object store through SharePoint Web Parts.
 Add a document to the object store using SharePoint.
 - c. Connector for SharePoint Web Parts
 Log in to Workplace or Enterprise Manager and verify that you can see existing SharePoint documents using those user interfaces.
2. Verify that 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.
 - c. Verify that SharePoint clients have access after a SharePoint failover.
 - i. Fail over your highly available SharePoint server.
 - ii. Verify that the SharePoint client can still connect to SharePoint and that SharePoint can connect to FileNet P8.

Installing 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 *FileNet P8 Hardware and Software Requirements* and the *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

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

Before you begin installing the software, retrieve *IBM FileNet Connector for SharePoint Web Parts Installation and Upgrade Guide, Version 2.2* (PDF) from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

NOTES

- The following procedure covers the specific details required for installing IBM FileNet Connector for SharePoint Web Parts in an existing SharePoint farm rather than on a standalone SharePoint server.
- You can configure IBM FileNet Connector for SharePoint Web Parts 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 123](#).
- For information on how to install IBM FileNet Connector for SharePoint Document Libraries in a highly available environment, see [“Installing a Highly Available IBM FileNet Connector for SharePoint Document Libraries 2.2” on page 126](#).

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 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

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 122.
 - 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 124.
- 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.

Installing 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 *FileNet P8 Hardware and Software Requirements* and the *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

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

Before you begin installing the software, retrieve *IBM FileNet Connector for SharePoint Document Libraries Installation and Upgrade Guide, Version 2.2* (PDF) from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

NOTE For information on how to install IBM FileNet Connector for SharePoint Web Parts in a highly available environment, see [“Installing a Highly Available IBM FileNet Connector for SharePoint Web Parts 2.2” on page 121](#).

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*.

The IBM FileNet Connector for SharePoint Document Libraries 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 128](#).
- 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 129](#).

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 128.

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

To install IBM FileNet Connector for SharePoint Document Libraries 2.2 in a highly available environment

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 127.
 - 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*.

To install IBM FileNet Connector for SharePoint Document Libraries 2.2 in a highly available environment

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 127](#), 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\UFIDatabase 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.

- 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
- Open Cluster Administrator for the Microsoft active/passive cluster.
- Right-click your group name and choose New > Resource.
- In the New Resource Box, give the new resource a name. For example: “Doc Lib Connector.”
- Under Resource Type, choose “Generic Service.”
- Click **Next**.
- Make sure both servers are listed under “Possible Owners.” Click **Next**.
- For Dependencies, move the drive that you installed the Doc Lib Connector Remote Admin Services on from “available resources” to “resource dependencies.”
- Under Service Name, enter:


```
ibm.ctms.taskrouteservice
```
- Leave start parameters blank. Click **Next**.
- On the registry replication page, click **Add** and enter:


```
SOFTWARE\IBM\EMRC\4.0
```
- 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 131.](#)
- [“To verify failover of the Document Libraries Connector Remote Admin Services on the active/passive cluster.” on page 132.](#)
- [“To verify the Document Libraries Connector Services on the SharePoint server farm nodes” on page 132.](#)
- [“To verify document transfer after Content Engine failover” on page 133.](#)

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 130](#) and bring it on-line.

- iii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 133.](#)

To install IBM FileNet Connector for SharePoint Document Libraries 2.2 in a highly available environment

- 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 130](#) and bring it on-line.
 - iii. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 133.](#)
 - 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 133.](#)
- 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 130](#) and bring it on-line.
 - iv. Do [“To verify transfer of documents from SharePoint to IBM fileNet P8” on page 133.](#)
 - 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 130](#) 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 130](#) and bring it on-line.

ix. Do ["To verify transfer of documents from SharePoint to IBM fileNet P8" on page 133.](#)

- 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 130](#) and bring it on-line.

- iii. Do ["To verify transfer of documents from SharePoint to IBM fileNet P8" on page 133.](#)

- 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 130](#) 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 130](#) and bring it on-line.

- viii. Do ["To verify transfer of documents from SharePoint to IBM fileNet P8" on page 133.](#)

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.

Installing a Highly Available IBM FileNet Services for Lotus Quickr

CAUTION This topic covers installation of FileNet Services for Lotus® Quickr™ 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

Overview

You can install FileNet Services for Lotus Quickr in highly available configurations using standard WebSphere application server clustering software.

A highly available installation of FileNet Services for Lotus Quickr should connect to a highly available deployment of Content Engine.

In highly available configurations using WebSphere clusters, a load balancing proxy web server is required to distribute load among the server nodes. The following procedures document using IBM HTTP Server (IHS) as the proxy server.

Procedure Overview for FileNet Quickr Services

The following high-level steps are necessary to configure IBM FileNet Services for Lotus Quickr for high availability:

1. Install and configure FileNet P8 platform.
2. Install and configure a WebSphere application server cluster for FileNet Services for Lotus Quickr.
3. Install and configure IBM HTTP Server (IHS) for proxy server and load balancing.
4. Install FileNet Services for Lotus Quickr on all nodes.
5. Create EAR file and deploy FileNet Services for Lotus Quickr to the cluster.
6. Generate IHS plug-ins for the deployed FNQuickrServices application

IBM FileNet Services for Lotus Quickr Installation and Deployment Tasks

Pre-deployment Tasks

1. Install and configure Content Engine, Process Engine, and Application Engine/Workplace XT.

Follow the instructions in the *IBM FileNet P8 Installation and Upgrade Guide*. For high availability installation of these components, see “Installing a Highly Available Content Engine” on page 34, “Installing a Highly Available Process Engine” on page 53, and see “Installing a Highly Available Application Engine/Workplace XT” on page 76.

2. Install IBM HTTP Server (IHS) and configure for both load balancing and proxy features.
3. Verify that the WebSphere cluster nodes and IHS proxy server/load balancer are started.

Deployment Tasks

For more information about FileNet Services for Lotus Quickr installation, see the “Install FileNet Services for Lotus Quickr” task in the *IBM FileNet Services for Lotus Quickr Installation and Upgrade Guide*.

To deploy FileNet Services for Lotus Quickr in an HA environment

1. Install FileNet Services for Lotus Quickr.

Follow the instructions in “Install FileNet Services for Lotus Quickr” of the *IBM FileNet Services for Lotus Quickr Installation and Upgrade Guide* with the following HA specific modifications:

- Install FileNet Services for Lotus Quickr on each node.

2. Generate the Quickr Services EAR file.

Follow the instructions in “Build and deploy the EAR file” of the *IBM FileNet Services for Lotus Quickr Installation and Upgrade Guide* with the following HA specific modifications:

- In a highly available environment you can deploy the EAR file either using the WebSphere administrative console or using the Configuration Manager. If you choose to deploy using the WebSphere administrative console you should to build the EAR file and configure the FileNet Services for Lotus Quickr bootstrap settings, and then deploy the EAR file.
- CE URI addressing format

When creating the XML configuration file with the bootstrap information, use the information below for the CE URI address:

```
ceURI=corbaloc::<node1_server_name>:<port#>,[<node#_server_name>:<port#>],:<node  
Z_server_name>:<port#>/cell/clusters/<CE_Cluster>/FileNet/Engine
```

Where:

<node#_server_name> represent the nodes in the cluster.

<nodeZ_server_name> represents the final node in the cluster.

Example:

If the cluster contains three server nodes, the CE URI addressing could be:

```
ceURI=corbaloc::node1_server_name:9811,
node2_server_name:9810,:node3_server_name:9809/cell/clusters/CECluster/FileNet/
Engine
```

3. Deploy the EAR file to the cluster.

Use the WebSphere administrative console or the Configuration Manager to deploy FileNet Services for Lotus Quickr to the cluster, making sure the web modules are mapped into all cluster node servers.

4. Use the WebSphere administrative console to generate IHS plug-ins for the deployed FNQuickrServices application.
5. Restart IHS and WebSphere cluster.

Deployment Verification

1. Make sure the WebSphere cluster and the IHS proxy server/load balancer are up and running.
2. Open a web browser and type in the following URL:

To verify deployment of the cluster

If using a proxy server the URL has the form:

```
http://<proxy_server_name>:<port_number>/<context_root>
```

Example:

```
http://appserver-proxy.server.com:9080/dm
```

To verify deployment of a single node

For an individual node the URL has the form:

```
http://<FNQuickr_Services_server>:<port_number>/<context_root>
```

Example:

```
http://fnquickr_services.server.com:9080/dm
```

3. Verify that the information page displays.

The following example shows some of the information displayed:

```
IBM FileNet Services for Lotus Quickr
Product name IBM FileNet Services for Lotus Quickr
Product version v1.0
Build number p8quickr100.026
P8 Content Engine URL corbaloc://node1_server_name:9811,
node2_server_name:9810,:node3_server_name:9809/cell/clusters/CECluster/FileNet/
Engine
```

P8 Application Engine URL `http://example.com:9080/WorkplaceXT`

P8 domain `example-domain`

P8 object store `example-OS`

4. Fail one of the nodes.
5. Verify that you can reload the information page.

Core Component Upgrade Tasks

This section provides the information needed to upgrade an existing FileNet P8 3.5.x highly available environment to FileNet P8 4.0.x.

A number of service packs, fix packs, and interim fixes have been released since the initial release of 4.0.0 software are required to reach a fully supported 4.0.x level. This is especially true for customers using DB2 as their database engine.

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 install and configure IBM FileNet P8 platform in an HA environment, follow the installation road map found in the “Upgrade Roadmaps” section of the *FileNet P8 4.0.x Platform Installation and Upgrade Guide*, with the following HA specific modifications:

1. Review IBM FileNet P8 platform requirements and other planning considerations.
In the *Upgrade Overview* section, see “General Requirements for all FileNet P8 Systems.”
2. Upgrade the IBM FileNet P8 platform documentation.
In the *Upgrade Core Components* section, do the task “Upgrade FileNet P8 Documentation.”
3. Ensure that you are running the minimum required release of IBM FileNet P8 platform software and service packs to perform an upgrade.
In the “Upgrade Core Components” section, do the task “Verify Current Release Level.”
4. Upgrade the highly available Content Engine Cluster/Farm.
For more information, see [“Upgrading a Highly Available Content Engine” on page 140](#).
5. Upgrade the highly available Process Engine Cluster.
For more information, see [“Upgrading a Highly Available Process Engine” on page 146](#).
6. Upgrade the highly available Application Engine Cluster/Farm.
For more information, see [“Upgrading a Highly Available Application Engine” on page 156](#).
7. Install service packs, fix packs, and interim fixes for core components.
In the “Upgrade Core Components” section, do the task “Apply Service Packs and Interim Fixes for Core Components.”

Upgrading a Highly Available Content Engine

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

Overview

- The 4.0.x 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 4.0.x 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 4.0.x Content Engine will take ownership of these, and they will no longer be available to the 3.5.x Content Engine.

Content Engine Upgrade Overview

The 4.0.x 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 [“Installing a Highly Available Content Engine” on page 34](#).

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 4.0.0 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 4.0.0 GCD database repository.
- b. Migrate and upgrade object stores and file stores to the FileNet P8 4.0.0 Content Engine.

Content Engine Upgrade

CAUTION If you install version 4.0.x 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

Prepare Content Engine for Upgrade involves performing the following tasks as described in “Prepare the Content Engine for upgrade” from the Upgrade Content Engine task in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

NOTE The bulleted items below correspond to procedures in the *IBM FileNet P8 Platform Installation and Upgrade Guide*. Some of the items are marked (*No HA specific steps needed*). You must still perform these steps, but they do not require any HA specific input.

- “To create the database for the GCD.”
For more information about making databases highly available with P8 components, see [“Database Services” on page 30](#).
- “To create or designate user accounts required for upgrade” (No HA specific steps needed).
- “To install and configure an application server for Content Engine Server.”
For pre-installation, verification, and configuration steps needed for a highly available Content Engine, see [“Installing a Highly Available Content Engine” on page 34](#). The steps listed are needed both for upgrades and new installations.
- “To verify that all in-progress event actions have finished” (No HA specific steps needed).
- “To verify that publishing requests have been completed” (No HA specific steps needed).
- “To delete custom subclasses of the PublishRequest object” (No HA specific steps needed).
- “To verify that all content-search indexing activity has completed” (No HA specific steps needed).
- “To stop and disable all Content-Engine-related services on all servers in the FileNet P8 domain.”

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 the shared folders using the same shared path that is stored in the GCD for the file store folder.

- “To verify that all in-progress Content Engine transactions involving content files have finished” (No HA specific steps needed).
- “To verify the directory permissions on file stores” (No HA specific steps needed).
- “To back up the database” (No HA specific steps needed).
- “To install and configure Autonomy K2” (No HA specific steps needed).
- “To upgrade Content Engine.”

- Installing Content Engine.

As part of the upgrade you must install version 4.0.x of Content Engine. The installation procedure is the same regardless of an upgrade or a new installation. For more information, see [“Installation and Deployment Tasks” on page 40.](#)

CAUTION When deploying Content Engine in an HA environment you must specify the cluster as the deployment target.

- Installing the Object Store 3.5 to 4.0 Upgrade Tool.

Install the upgrade tool on at least one Windows server by performing a “Custom” installation of the Content Engine and selecting “Object Store 3.5 to 4.0 Upgrade Tool.”

This Windows system must meet the requirements laid out in the *FileNet P8 Hardware and Software Requirements Guide*.

The Windows system does not need to be a part of a Windows Active Directory Domain.

The upgrade 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.

- 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 configure data sources and connection pools for existing object stores.”

For information on how to create data sources and connection pools for object stores, see [“Creating Object Stores in FileNet P8 4.0.x” on page 47.](#)

NOTE Create the data sources and connection pools for all object stores that will be upgraded.

- “To edit the upgrader utility file” (No HA specific steps needed).

To upgrade Content Engine Items

- Upgrading using the command line tool.

Perform the tasks in the “Command Line Interface to CE 3.5.2 to 4.0 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 Upgrader Tool using the CLI.”

- Upgrading using the Graphical User Interface

Run the upgrade tool perform the tasks in the “Graphical User Interface to CE 3.5.2 to 4.0 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 check Upgrader Tool prerequisites.”
- “To create the XML upgrade status file.”
- “To configure JAAS, P8 domain, CBR, and FCD 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.

Perform the tasks in the “Prepare Items for Upgrade” section

Do the following tasks from the “Prepare Items for Upgrade” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- “To prepare object store nodes for upgrade” (No HA specific steps needed).
- “To prepare file store and fixed file store nodes for upgrade.”

NOTE In order to complete the upgrade for file stores make sure that the highly available share is online and accessible from the node on which the upgrader tool is being run.

Perform the tasks in the “Perform the Upgrade” section

Do the following task from the “Perform the Upgrade” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- “To upgrade items” (No HA specific steps needed).

Perform the tasks in the “Clear the Read-Only Attribute for NTFS File Stores” section

Do the following tasks from the “Clear the Read-Only Attribute for NTFS File Stores” section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- “To clear the read-only attribute for NTFS file stores” (No HA specific steps needed).
- “To rebuild the spanword index” (No HA specific steps needed).

The 3.5.x/4.0.x Content Engine Farm/Cluster after Upgrade

After the 4.0.x upgrade is performed you can bring the Content Engine farm/cluster resources online if needed. The 3.5.x Content Engine and 4.0.x 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/4.0.x 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 4.0.x 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 4.0.x environment.

Upgrading a Highly Available Process Engine

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

Procedure Overview for Process Engine Upgrade

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 4.0.x.
3. Uninstall PE 3.5.x from the second node in the cluster.
4. Install PE 4.0.x 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.

The following procedures are based on these steps.

Migrating PE 3.5.x Clusters to PE 4.0.x Farms

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. 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 4.0.x.
 - i. Upgrade the first Process Engine in the cluster as described in the *IBM FileNet P8 Platform Installation and Upgrade Guide* task "Upgrade Process Engine (<platform>)."

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 4.0.x 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.
2. 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.

MS SQL

- 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.

3. 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.

4. (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**.
5. Install Process Engine 4.0.x 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 Process Engine (<platform>)” 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_r10
- 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 the settings listed in the “Complete <database> database screens” step are 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 - Do “To configure contiguous free memory for Process Engine (Windows only).”

d. Start Process Engine on the second node.

6. 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.

MS SQL

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.

7. Configure the hardware load balancer.

- a. Verify that your hardware load balancer is set up to allow ALL PROTOCOLs traffic, not just TCP.
- b. For each Process Engine server, in the hosts file, set the hardware load balancer virtual name to the IP address of the PE server itself.

Example:

<loadbalancer_virtual_name> (10.15.4.48) is the load balancer virtual name, and <PE_farm_1> (10.15.4.67), <PE_farm_2> (10.15.4.106), and <PE_farm_3> (10.15.4.111) are three PE servers.

The hosts file entries should look like this:

Server	Hosts file entry
<PE_farm_1>	10.15.4.67 <PE_farm_1> <loadbalancer_virtual_name>
<PE_farm_2>	10.15.4.106 <PE_farm_2> <loadbalancer_virtual_name>

CAUTION On UNIX verify that the `nsswitch.conf` file is configured to resolve names via the local hosts 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 (`vtaskman`), 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.

Upgrading UNIX Clusters

NOTE The bulleted items below correspond to procedures in the *IBM FileNet P8 Platform Installation and Upgrade Guide*. Some of the items are marked (*No HA specific steps needed*). You must still perform these steps, but they do not require any HA specific input.

Pre-Upgrade Tasks

1. Stop all Process Routers in the FileNet P8 environment.
2. Stop the clustered Process Engine resource.
3. Delete the FileNet Process Engine resource from the cluster.
4. Perform all other pre-upgrade tasks listed in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, the Upgrade Core Components section, Upgrade Process Engine (UNIX) task.
 - “To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x.”
 - “To verify that all Process Analyzer 3.5.x events have been published.”
 - “To purge event logs and statistics records.”
 - “To stop all Process Engine-related services and applications.”
Do this on each node prior to upgrade.
 - “To turn off Oracle Password Complexity Verification.”

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 “To update Process Engine (UNIX)” section in the “Upgrade Process Engine (Unix)” task of 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 the Web Service Transport” on page 175](#).
3. Perform post-update tasks as documented in the “Upgrade Process Engine (UNIX)” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - “To set file permissions.”
 - “To restore any custom modifications for root and fnsd users.”
 - “To enable ports.”**CAUTION** Do not re-enable Oracle Password Complexity Verification until after all nodes have been updated.

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. Perform the tasks listed in the “Complete Post-Upgrade Process Engine Configuration” task in the Upgrade Core Components section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - “To update the Process Engine database objects.”
Perform this step from the active cluster node.
 - “To update the Process Engine security.”
Perform this step 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.
 - “To update email notification.”
 - “To clean up Process Engine database views.”
6. Stop the Process Engine services prior to putting them under Veritas control.

Execute:

```
killfnsw -DAyS
```

7. Install the FileNet P8 4.0.0 Process Engine Agent.
 - a. Untar the FileNet Process Engine Agent package located under the Veritas folder on the **FileNet P8 Process Engine software** 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 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_r10 partition resource

8. Bring the cluster online.

Upgrading Windows Clusters

NOTE The bulleted items below correspond to procedures in the *IBM FileNet P8 Platform Installation and Upgrade Guide*. Some of the items are marked (*No HA specific steps needed*). You must still perform these steps, but they do not require any HA specific input.

Pre-Upgrade Tasks

1. Stop all Process Routers in the FileNet P8 environment.
2. Stop the Process Engine service and Process Engine registry replication resources.
3. Disable or remove all clustered resources for FileNet Process Engine services (IMSService and Process Services Manager) and registry replication in the cluster.
4. Perform all other pre-upgrade tasks listed in the *IBM FileNet P8 Platform Installation and Upgrade Guide*, the “Upgrade Core Components” section, “Upgrade Process Engine (Windows)” task.
 - “To verify that all Process Engine 3.x events have been transmitted to Process Analyzer 3.x.”
 - “To verify that all Process Analyzer 3.5.x events have been published.”
 - “To purge event logs and statistics records.”
 - “Create the Process Engine ODBC data source and test the connection (SQL Server only).”
Perform this step on every node that will run Process Engine with a SQL database backend.
Make sure to use the highly available virtual name for your database if needed.
 - “To stop all Process Engine-related services and applications.”
 - “To turn off Oracle Password Complexity Verification.”

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 “To update Process Engine (Windows)” section in the “Upgrade Process Engine (Windows)” task of 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 the Web Service Transport” on page 175](#).

CAUTION Do not re-enable Oracle Password Complexity Verification until after all nodes have been updated.

3. Move the cluster group to another node and repeat the previous step. Do this for all nodes in the cluster that will run Process Engine services.
4. Perform the tasks listed in the “Complete Post-Upgrade Process Engine Configuration” task in the Upgrade Core Components section of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.
 - “To update the Process Engine database objects.”
Perform this step from the active cluster node.”
 - “To update the Process Engine security.”
Perform this step 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.
 - “To update email notification.”
 - “To clean up Process Engine database views.”
 - “To run SQL Server post-upgrade scripts (SQL Server only).”
 - “To configure continuous free memory for Process Engine (Windows only).”
5. 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 66.](#)
 - For Microsoft Cluster Server, see [“Modify the registry to set the NCHBroadcast value.” on page 71.](#)
6. 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.

Upgrading a Highly Available Application Engine

This section covers the tasks and requirements for upgrading a highly available Application Engine in a FileNet P8 environment. The Application Engine upgrade process is straight forward for both single instance and farmed/clustered Application Engines and 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.

Procedure overview for Application Engine Upgrade

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.

Pre-Upgrade Tasks

Perform the Steps under “Before you upgrade Application Engine” in the Upgrade Application Engine task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

- Backup the shared bootstrap properties file in the shared network location.
- 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**.

Upgrade Tasks

Perform the steps listed under “To upgrade Application Engine” in the “Upgrade Application Engine” task of the *IBM FileNet P8 Platform Installation and Upgrade Guide*, with the following modifications.

- Stop Process Routers and Component Manager on all nodes. Make sure Process Task Manager is also stopped.
- In the Upgrade the Application Engine software step, specify the highly available Content Engine connection information. For information, see [“Content Engine in an Application Server Cluster using the EJB Transport” on page 174](#).
- In the Copy the retained configuration files step, copy the configuration files to the shared configuration directory, if you are using one.

Functional Expansion Components Upgrade Tasks

This section covers high availability upgrade and configuration of FileNet functional expansion components. FileNet functional expansions can be upgraded and configured for high availability in any order after you have upgraded your IBM FileNet P8 core components.

CAUTION Before attempting to Upgrade and configure these components, make sure your IBM FileNet P8 core environment (CE, PE, and AE) is up and running.

To upgrade FileNet functional expansions components in a highly available environment

- [“Upgrading a Highly Available Records Manager” on page 158.](#)
- [“Upgrading a Highly Available BPF” on page 161.](#)

Additional Documentation

NOTE For general instructions on how to navigate to this and other IBM FileNet product documentation on the IBM web site, see [“Access IBM FileNet Documentation” on page 13.](#)

Depending on your setup you will need additional documentation to set up your FileNet Functional Expansions for high availability:

- *FileNet Records Manager Installation and Upgrade Guide* (PDF).
- *IBM FileNet Business Process Framework Installation Guide* (PDF).

Upgrading 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 4.0
 - 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 P8 Records Manager Installation and Upgrade Guide*, version 4.0 or 4.5.

Procedure overview for Records Manager Upgrade

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 4.0 or 4.5, with modifications to specific tasks as outlined below.

Pre-upgrade Tasks

- (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.

Upgrade Tasks

1. 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 details:

- “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 [“Installing a Highly Available Records Manager” on page 95](#).

- “Complete the Setup screens as follows:”

When prompted, enter the HA specific information you identified above.

2. Configure report viewing if required.

If your setup is using reports, follow the instructions in the “Add Report Viewing Support” task with the following modifications:

You must copy the Crystal Reports JAR files to all server instances/nodes in the farm/cluster.

3. 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.

4. 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 or “Redeploy Application Engine or Workplace XT” task of the *IBM FileNet Records Manager Installation and Upgrade Guide*, version 4.0.

Perform these steps on all server instances/nodes in the farm/cluster.

5. Redeploy the RM Web Application.

Follow the instructions in the “Deploy the Records Manager Web Application” task.

- For WebSphere, refer to the “IBM WebSphere 5.x and 6.x” procedures Using these steps re-deploy Records Manager on the cluster and proxy server instances.
- For WebLogic, refer to the “BEA WebLogic 8.1.x and 9.x” procedures. Use these steps to deploy Records Manager on the cluster.

NOTE When deploying the Records Manager Web Application, specify the cluster as the deployment location instead of individual server instances.

6. 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 FileNet P8 Administration > Enterprise-wide Administration > Process Task Manager > Getting started > Start Process Task Manager.

Upgrading a Highly Available BPF

CAUTION This topic covers upgrade of Business Process Framework (BPF) in a supported FileNet P8 HA environment. Do NOT upgrade 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 *P8-4.x.x Fix Pack Dependency/Compatibility Matrix* documents. To download these documents from the IBM support page, see [“Access IBM FileNet Documentation” on page 13](#).

NOTES

- The following procedure covers upgrade of the BPF Web Application and BPF Operations components.
- When upgrading BPF Explorer in an environment with clustered databases, make sure you configure it to use the cluster’s virtual IP address for the ODBC data source.
- The following procedure covers the specific details required for upgrading BPF in an application server cluster rather than on a standalone server.

Procedure Overview for BPF

To upgrade BPF in a highly available environment

High-Level Steps

The following high-level steps are necessary to upgrade BPF for high availability:

1. Upgrade Content Engine and Process Engine.
Follow the instructions in the *FileNet P8 4.0.x Installation and Upgrade Guide*. For high availability upgrades of these components, see [“Upgrading a Highly Available Content Engine” on page 140](#) and [“Upgrading a Highly Available Process Engine” on page 146](#).
2. Upgrade your highly available Application Engine. For information, see [“Upgrading a Highly Available Application Engine” on page 156](#).
3. Configure your BPF environment.
See [“Configure your environment before BPF server upgrade.” on page 162](#).
4. Upgrade BPF on the first (AE) node.
This configures the FileNet P8 platform components for BPF. See [“Upgrade and Configure the BPF Web Application on the first node.” on page 163](#).
5. Upgrade BPF on the remaining (AE) nodes.
See [“Install the BPF Web Application on additional nodes.” on page 164](#).
6. Deploy BPF in the cluster.
See [“Manually Deploy the BPF Web Application.” on page 166](#).
7. Perform any necessary post-deployment steps.
See [“Configure BPF” on page 168](#).

Pre-upgrade Tasks

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

1. (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.
2. Configure your environment before BPF server upgrade.
Follow the instructions in the subtopic “Prerequisites for Upgrade” and “Configuration of the environment before BPF server upgrade” of the “Automatic Upgrade BPF to Release 4.1” task in the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

Perform the following “Configuration of the environment before BPF server upgrade” steps on all AE/BPF nodes:

- Back up your existing bp8ciops.jar used by BPF_Operations component.
- Back up the entire deployed folder used by the BPF Web Application for recovery purposes only.
- Back up the Workplace files used by BPF, for recovery purposes only:
- Create a custom code folder and copy any custom code modules to be deployed during the upgrade into this folder.
- If you have any custom web modules registered using Layout Designer, take a back up of the BPFModule.xml file located under the WEB-INF folder in the deployed BPF Web Application.
- Export the configuration of your {meta database} to an XML manifest using the Export feature in BPF Explorer and make a full database backup within your RDBMS.
- Ensure that the BPF metastore database server is running, so that the installer can update the metastore.

Upgrade Tasks

For more information about BPF installation, see the *IBM FileNet Business Process Framework Installation Guide*.

CAUTION Installation requires that the Application Engine and BPF be co-located on the same application server. For more information, see the “Automatic Upgrade BPF to Release 4.1” topic in the *IBM FileNet Business Process Framework Installation Guide*.

1. Upgrade and Configure the BPF Web Application on the first node.

Follow the instructions in the “BPF server upgrade with automatic configuration” task in the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- In the Installation vs. Upgrade screen, select “Upgrade.”
- In the Select Features screen, select BPF Metastore, BPF Web application, and BPF Operations only.
- In the Existing BPF App info screen, enter the correct BPF Application name.

CAUTION You must enter the correct BPF Application name for the upgrade application to be able to uninstall BPF.

- In the Automatic Metastore upgrade, select “NO.”
- WebSphere
 - In the WebSphere server access info screen, choose:
 - The Deployment Manager's profile
 - The Cell
 - The CellManager Node
 - The Deployment manager server instance

- The Deployment Manager's SOAP port
 - The local node's HTTP Port
 - WebLogic
 - In the WebLogic information screen, enter:
 - WebLogic Domain Name
 - Admin Server Name
 - Admin Server Port
 - Admin Login Name
 - Password
 - JBoss
 - In the JBoss information screen, enter:
 - JBoss Home Directory
 - JBoss Server Type - all
 - HTTP Port
 - In the Deployment to web application server screen choose “No” to manually deploy the application.
 - In the Content Engine WSI access screen enter the Content Engine Client Software URL.

NOTE Point to a single CE server node using the correct host name and port number. Do not use the CE cluster address.
 - In the Content Engine EJB access screen enter the Content Engine URL.

NOTE 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 the EJB Transport” on page 174.](#)
 - In the Workplace Configuration Files screen:
 - Deployed Workplace Directory - Enter the Workplace installed location.
 - Workplace Config Directory - The shared AE bootstrap location.
 - In the Workplace URL screen, specify:
 - Workplace URL: the proxy or load balanced URL for Workplace
2. Install the BPF Web Application on additional nodes.
- Follow the instructions in the “BPF server upgrade with automatic configuration” task in the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:
- In the Install vs. Upgrade screen, select “New Install.”
 - In the Choose installation type screen, select “Custom.”

- In the Select Features screen, make sure only the following components are selected:
 - BPF Metastore
 - BPF Web Application
 - BPF Operations - Component Manager Component
- In the Automatic Metastore creation screen, select “NO.”

NOTE This has already been done as part of upgrading the first node.
- In the BPF Explorer XML Manifest screen, select “Manual import of XML Manifest.”

NOTE This has already been done as part of upgrading the first node.
- WebSphere

In the WebSphere server access info screen, configure the following:

In the WebSphere server access info screen, configure the following:

 - Profile - select the local application server profile (e.g., AppSrv01).
 - Cell - select the same value as for the first node.
 - Node - select the value for the local application server's node.
 - Server - select the local application server instance (e.g., server1).
 - SOAP Port - enter 8880 (the SOAP port for the local application server node).

CAUTION During the upgrade the installer attempts to uninstall the application. In a clustered environment the installer must be able to connect to a SOAP port during the upgrade process, or it will fail with a critical exception. By configuring the upgrade for a local SOAP port the installer will fail with a non-critical exception, which will allow the upgrade to continue.

 - HTTP Port - enter the local application server HTTP port (e.g., 9080)
- WebLogic
 - In the WebLogic information screen, enter the same values you entered for the first node.
- JBoss
 - In the JBoss information screen, enter the same values you entered for the first node.
- In the Deployment to web application server screen choose “No” to manually deploy the application.
- In the P8 administrative account verify that the user name and password fields are blank.
- In the Content Engine WSI access screen verify that all fields are blank.
- In the Content Engine EJB access screen, enter the same value you entered for the first node.

- In the Workplace Configuration Files screen:
 - Deployed Workplace Directory - Enter the Workplace installed location.
 - Workplace Config Directory - The shared AE bootstrap location
- In the Workplace URL screen, enter the same values you entered for the first node.
- In the Process Engine access screen verify that all fields are blank.

CAUTION When running the installer with the above fields left blank a number of non-critical exceptions will display in the Installation Summary screen. These all relate to configuration entries that were not made and that you will either perform manually or have already performed as part of the initial node installation, and can be ignored.

3. Manually Deploy the BPF Web Application.

Follow the instructions in the “Deploy BPF Web Application” task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- Application server cluster
 - WebSphere.

Configure and deploy the BPF Web Application from the administrative server instance.

NOTE When installing BPF through the administrative server make sure to deploy the BPF to the cluster of server instances.

Follow the instructions in the “Deploy BPF Web Application” task “IBM WebSphere <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- 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.

Configure and deploy the BPF Web Application from the administrative server instance.

NOTE When installing BPF through the administrative server make sure to deploy the BPF to the cluster of server instances.

Follow the instructions in the “Deploy BPF Web Application” task “BEA WebLogic <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- Make sure the targets for the application include the WebLogic cluster and/or managed nodes that will be running the application.
- JBoss

Follow the instructions in the “Deploy BPF Web Application” task “JBoss 4.0.x” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - Copy the web application files from the first node to each additional node.
 - Copy the datasource XML files from the first node each additional node.
 - Restart JBoss on each node.
- Farm of independent application server instances
 - WebSphere.

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “IBM WebSphere <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - For WebSphere make sure you select the cluster and/or managed nodes that will be running the application when mapping modules to servers. The Deployment Manager 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 If you are using an IBM HTTP Server to proxy/load-balance application requests to WebSphere, make sure to select the webserver node as well as the application servers when mapping modules to the application.
 - WebLogic.

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “BEA WebLogic <version>” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

 - Make sure the targets for the application include the WebLogic cluster and/or managed nodes that will be running the application.
 - JBoss

Configure and deploy the BPF Web Application on each node.

Follow the instructions in the “Deploy BPF Web Application” task “JBoss 4.0.x” sub-task in the “Troubleshooting and Manual Procedures” section of the *IBM FileNet Business Process Framework Installation Guide*, with the following modifications:

- Copy the web application files from the first node to each additional node.
- Copy the datasource XML files from the first node each additional node.
- Restart JBoss on each node.

4. Configure BPF

Follow the instructions in the “Post-upgrade BPF server configuration” tasks in the *IBM FileNet Business Process Framework Installation Guide*.

Upgrade Verification

To verify the installation follow the instructions in the “Validate BPF server upgrade” task in the *IBM FileNet Business Process Framework Installation Guide*.

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

1. Using the application server administration tool, verify that the BPF application is running.
2. Open a web browser and type in the URL for the BPF application.
 - If using a load-balancer or proxy server the URL has the form:
`http://<virtual_name>:<port_number>/BPF`
 - If connecting to an individual cluster node the URL would have the form:
`http://<clustered_server_name>:<port_number>/BPF`
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.

Appendixes

The appendixes section contains the following major topics:

- [“Appendix A – Setting up a Secure Native Mode Domain Installation” on page 170.](#)
- [“Appendix B – Connecting to a Highly Available Content Engine” on page 173](#)
- [“Appendix D – Recompiling the FileNet Process Engine Veritas Agent” 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
<Installer> ^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 171](#) 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 36](#).

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 FileNet P8 Administration > Content Engine Administration > FileNet P8 Domain > How to... > Logon to a domain.

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 the EJB Transport” on page 174](#).
- Application using the web service (WSI) protocol to connect to Content Engine using load-balancers must use the virtual server name of the load-balancing device for the connection. See [“Connecting Using the Web Service Transport” on page 175](#).

Content Engine in an Application Server Cluster using the EJB Transport

For clients such as Application Engine that use an EJB transport to communicate with Content Engine in an application server cluster configuration a 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 the Web Service Transport

Applications like Process Engine use the web service (also 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/FNCEWS40DIME
```

Example:

```
http://testvirtual1:7001/wsi/FNCEWS40DIME
```

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 4.0.

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 1 on page 176](#) outlines the types of load balancing that can be used with Content Engine for each transport.

Table 1: 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 – Email Manager – Records Crawler 	Yes	No	Yes	Yes

Table 1: 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 177)	Yes	No	No
FileNet Enterprise Manager and other administrative applications ^a	No	No	No	No

a.For information, see [“Administrative applications” on page 173.](#)

CAUTION To successfully configure your environment for Content Engine high availability for clients that use the EJB Transport ([Table 1 on page 176](#)), 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 the EJB Transport” on page 174](#). If you use a hardware load-balancer between a client and Content Engine using EJB transport, the load-balancer will only perform 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 4.0 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 1 on page 176](#))
- 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 179](#).

Table 2: Process Engine

PE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	Network Load Balancing
Java API AE/Workplace Workplace XT	Yes (TCP for client access, UDP for intraserver communication)	No	No	No
Web Services clients (WSI 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 plugin load balancing in combination with an HTTP server.

Table 3: Application Engine

AE Clients	Hardware Load Balancing (all protocols)	Application Server Load Balancing	HTTP-plugin Load Balancing	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 via 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 will force 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> should be 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/VRTSvcs/bin/FN_ProcessEngine

8. (If no agent is installed) Copy the newly created agent file to:

/opt/VRTSvcs/bin/FN_ProcessEngine

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo 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

GC31-5487-09

