IBM FileNet P8 Platform

Version 3.5

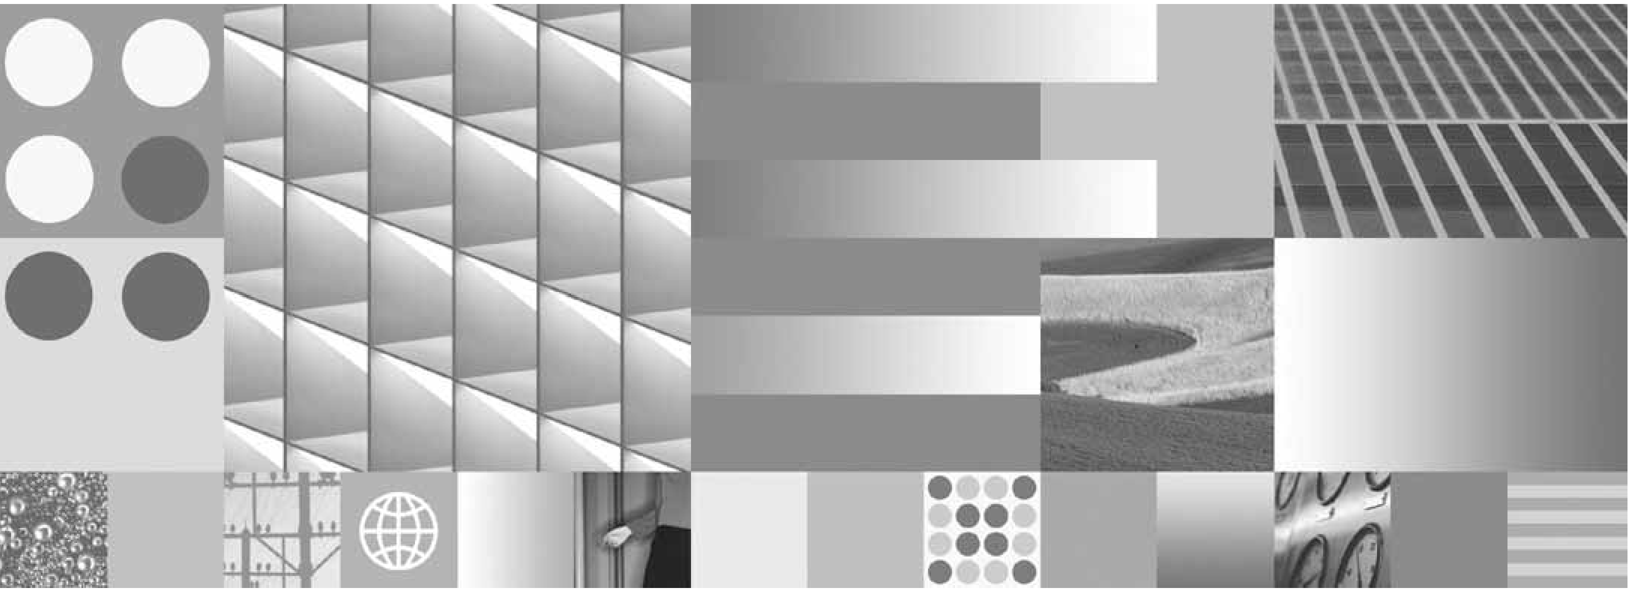

High Availability Technical Notice

GC31-5491-01

IBM FileNet P8 Platform

Version 3.5

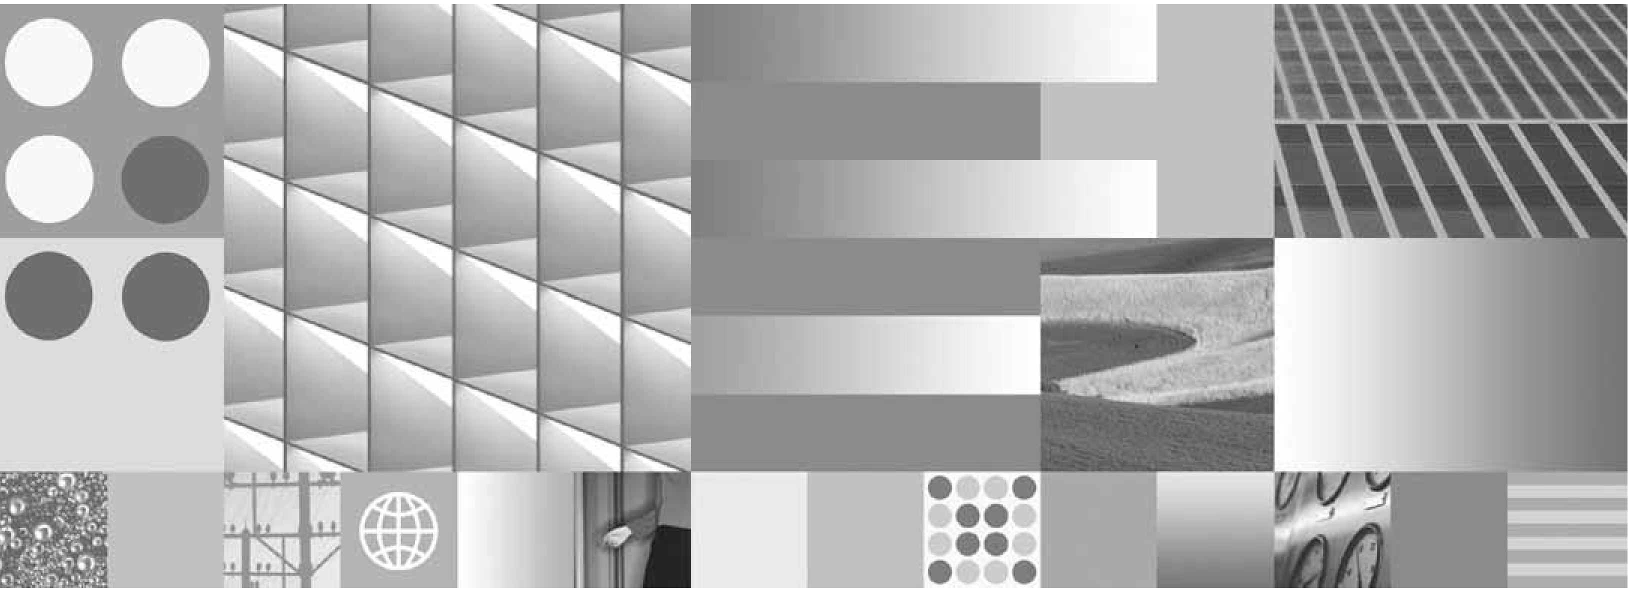

High Availability Technical Notice

GC31-5491-01

Note

Before using this information and the product it supports, read the information in "Notices" on page 121.

This edition applies to version 3.5.2 of IBM FileNet P8 Business Process Manager (product number 5724-R76), version 3.5.2 of IBM FileNet P8 Content Manager (product number 5724-R81), versions 3.5.1 and 3.7.0 of IBM FileNet Records Manager (product number 5724-S19), version 3.5.0 of IBM FileNet Team Collaboration Manager (product number 5724-S00), and to all subsequent releases and modifications until otherwise indicated in new editions.

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.
Italic	User-supplied variables and new terms introduced in text, names of additional documents (such as *IBM FileNet® P8 Platform Installation and Upgrade Guide*).
<*italic*>	User-supplied variables that replace everything between and including the angle bracket delimiters (< and >).
`Monospace`	Code samples, examples, display text, and error messages.

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

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

Revision Log

The following table identifies changes made to this document.

	Revision
06/08	Updated section "Access IBM FileNet Documentation" on page 12. Added section "Rendition Engine 3.5.0" on page 93.
01/08	Added section "Access IBM FileNet Documentation" on page 12. Added section "Team Collaboration Manager 3.5.0" on page 97.
02/07	Updated the section "Oracle Real Application Clusters" on page 22 with new requirements when using Oracle RAC servers with IBM FileNet P8: "Oracle RAC CE 3.5.x sites must have sequences set to use ORDER" on page 23 and "Oracle 9i RAC must use Broadcast on Commit (BOC)" on page 23. Updated section "Object Store Service" on page 29 with two new topics: • "Configure Content Engine COM Client stations to not start the object store service on a passive node" on page 30. • "Configure COM Clients in a farmed environment to use a specific Object Store service" on page 30. Added requirement to add a DTC resource in "Object Store Service - VERITAS Cluster Server" on page 33. Updated section "File Store Service HA Requirements" on page 39 with content indexing information in "(Windows 2000 only) Set security on the cluster temp directory to allow content search indexing." on page 39. Added section "Upgrade a Records Manager Cluster" on page 112.
12/06	Added section "Records Manager 3.7.0" on page 89 and "Appendix B – Redeploying Workplace after Records Manager is Installed" on page 119.
07/06	Added section "PE Database Reconnection" on page 64. Added section "Managing Workplace settings in a load balanced environment" on page 70. Added "Records Manager 3.5.0" on page 81 and "Install the Records Manager 3.5.x service pack" on page 85. Added information about using load balancers or proxy servers with AE and RM on page 67, page 78, page 84, and page 87.
12/20/05 (3.5.x)	Initial revision.

Table of Contents

IBM FileNet P8 Platform High Availability

This document describes the technical issues related to configuring the IBM 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 IBM FileNet P8 3.5.0 Content Engine, Process Engine, and Application Engine, and the application of IBM FileNet P8 3.5.x service packs to your HA environment. The final section provides information about how to upgrade an existing IBM FileNet P8 3.0.0 highly available environment.

CAUTION Under no circumstances are these instructions provided as a step-by-step guide detailing the use of the many third-party software and tools required to make highly available clusters.

NOTES

This technical notice is for experienced System and Database Administrators who will be using this documentation to implement their intended environments and requires:

- A working administrative knowledge with the IBM FileNet P8 system.
- 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 software, such as Microsoft NLB to set up and configure software load balancing.
- Experience with BEA WebLogic Server, IBM® WebSphere®, or JBoss/Tomcat for setting up application server clusters.
- Experience with supported operating systems and hardware platforms, volume management software, database platforms, and application servers.

Access IBM FileNet Documentation

To access documentation for IBM FileNet products

1. Access the Product Documentation for FileNet P8 Platform support page (http://www-1.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 support page (http://www-1.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 12.

Before you begin installing the software, retrieve the 3.5 versions of the following documents:

- *IBM FileNet P8 Platform Installation and Upgrade Guide* (PDF)
- *IBM FileNet P8 Hardware and Software Requirements* (PDF)
- *IBM FileNet P8 Release Notes* (PDF)
- *IBM FileNet Records Manager Installation and Upgrade Guide, version 3.5 or 3.7* (PDF)
- *IBM FileNet Rendition Engine Installation and Upgrade Guide* (PDF)
- *IBM FileNet Team Collaboration Manager Installation 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 such as META Group 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. The Gartner Group notes 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 people 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 note addresses just the hardware and software issues, but IBM 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.

How does High Availability relate to 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 the 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 is covered in the *IBM FileNet P8 Platform Disaster Recovery Technical Notice*. 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 12.

Building blocks for High Availability

Enterprise content management (ECM) systems typically consist 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 offers a software-based load balancing capability in Windows 2000 and Windows Server 2003 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, both are good candidates for server farms.

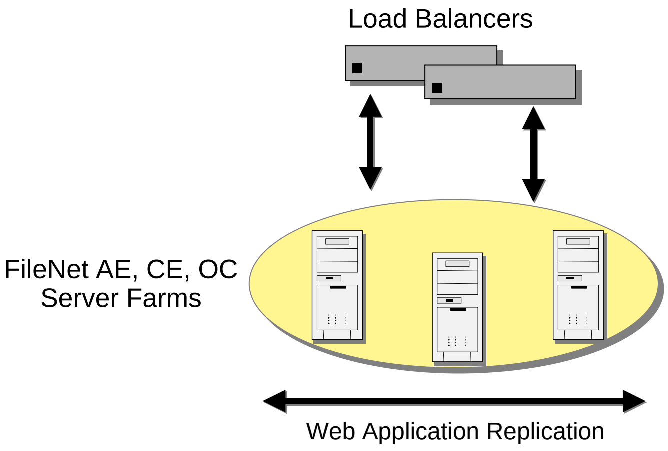

Figure 1: The basic server farm setup.

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.

Server Clusters

Server clusters are based on the concept of shared data storage, in contrast to server farms, which feature no sharing of data storage 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 high availability 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.

- 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.

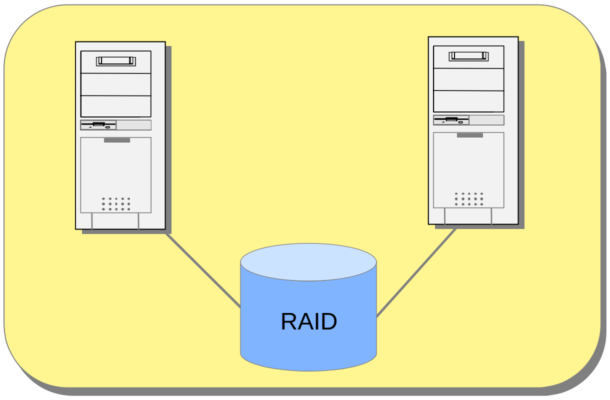

Figure 2: Basic server cluster using RAID storage.

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. VERITAS 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.

VERITAS 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.

Cluster Configurations

Cluster software vendors support several distinct types of cluster configurations. Using VERITAS's terminology, the notable configurations supported by the IBM FileNet P8 3.5.x 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.

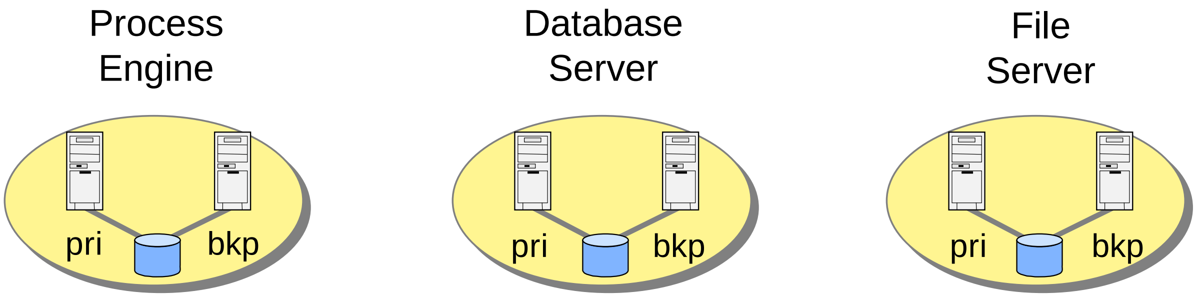

Figure 3: Asymmetric cluster.

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.

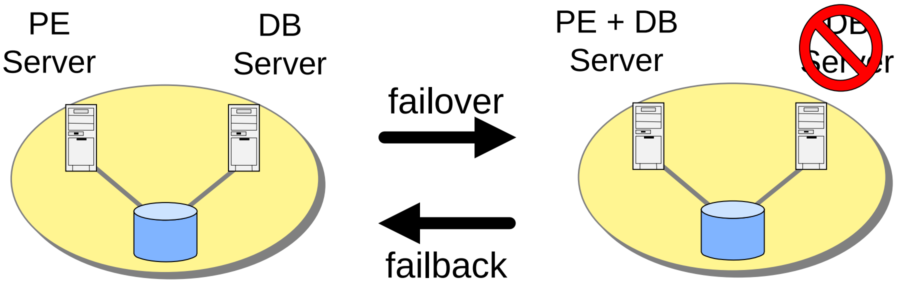

Figure 4: Symmetric cluster.

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 IBM FileNet P8 3.5.x 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.

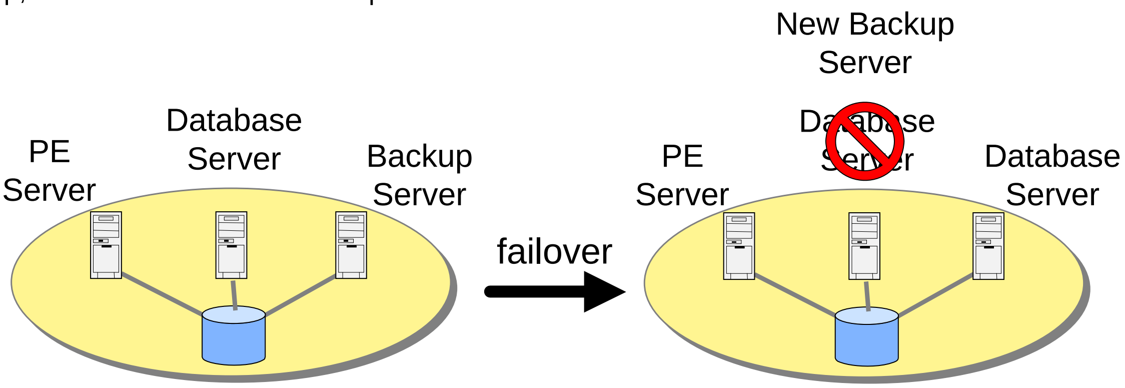

Figure 5: N+1 cluster.

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 IBM FileNet P8 N-to-M node could for example have four nodes:

1. Content Engine File Store Service node
2. Content Engine Object Store Service node
3. Process Engine server node

4. Remote database server node.

If the CE File Store Service node fails, it could be configured to fail over to the 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 CE Object Store Service node.

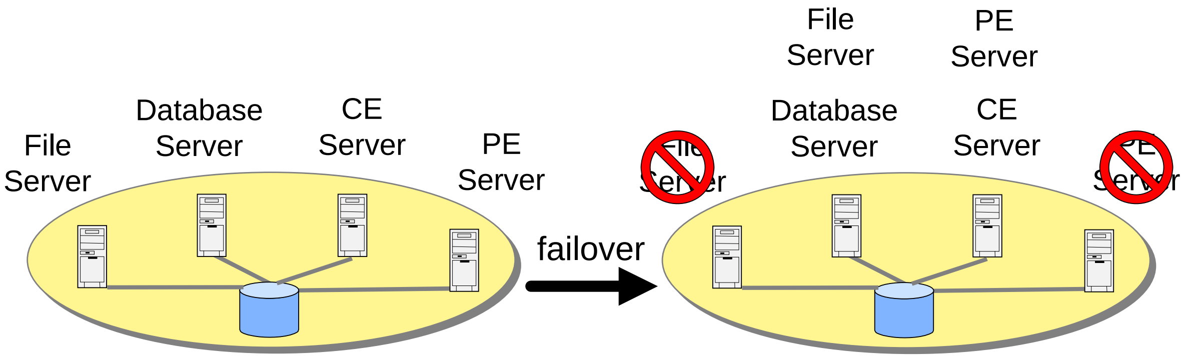

Figure 6: N to M cluster.

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 map from 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
VERITAS	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 with the IBM FileNet P8 Platform

IBM has tested the IBM FileNet P8 Platform with high availability products from BEA Systems, IBM, Microsoft, and VERITAS Software. 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. 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.

NOTE For a detailed list of the third-party high availability products tested and supported with IBM FileNet P8 Platform, please see the *IBM FileNet P8 3.5.x Hardware and Software Requirements* document located on the IBM Information Management support

www.ibm.com/software/data/support

In addition to the Content Engine and 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 53 for more details on the VERITAS high availability agent for the Process Engine.

Load Balancing/Farming

There are a number of hardware and software load-balancing products available for server farm configurations. IBM has tested the IBM FileNet P8 platform components with the following software load balancers to verify that the IBM FileNet P8 components function correctly:

- Microsoft NLB (Network Load Balancing)
- BEA WebLogic clusters
- IBM WebSphere server group of clones (called Cluster in later releases)

Microsoft NLB

Microsoft NLB is an example of a software load balancer for the Windows 2000 and Windows Server 2003 platforms.

BEA WebLogic and IBM WebSphere

The IBM FileNet P8 Application Engine is hosted on a Java™ application server. Both BEA and IBM's Java application servers have built-in capabilities for providing highly available web services. While these application servers use different terminologies for their solutions, the basic concepts are the same. 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. BEA WebLogic calls this collection of server instances a Cluster. IBM's WebSphere terms the collection of instances a Server Group of Clones, or Cluster. Both products function like a server farm.

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 Server Groups (Farms) of Clones

IBM's WebSphere application server provides high availability through server groups of clones. Server groups are a collection of servers that can be managed as a single server. You can deploy, start, and stop all servers in the server group as a single entity.

Clones are copies of an original server. They function as separate individual servers, yet their configuration is based on the same template configuration from an original server. Server groups can be used to group together cloned servers. Server groups 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.

Server Clusters

VERITAS Cluster Server and Microsoft Cluster Server are two examples of server cluster products. The VERITAS product runs on a number of hardware and operating systems platforms, including:

- RS/6000® AIX
- HP-UX
- Sun Solaris
- Microsoft Windows Server

VERITAS Cluster Server

VERITAS 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. VERITAS 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:

- Windows 2000 Advanced Server supports two nodes.
- Windows 2000 Datacenter supports up to four nodes.
- Windows Server 2003 Enterprise Edition supports up to eight nodes.
- Windows Server 2003 Datacenter Edition supports up to eight nodes.
- The Microsoft Hardware Compatibility List contains the hardware configurations that have been tested and certified for MSCS.

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 to use highly available database instances for any IBM FileNet P8 components (with the exception of Oracle Real Application Clusters as noted below).

IBM has tested the IBM FileNet P8 platform components with the following databases to verify that the IBM FileNet P8 components function correctly with a highly available database instance and that they can reconnect to the database instance after a failover:

- SQL Server on Microsoft Cluster Server (MSCS)
- Oracle on VERITAS Cluster Server
- IBM DB2® on VERITAS Cluster Server

Oracle Real Application Clusters

In addition to the traditional database server clusters IBM 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.

Oracle RAC CE 3.5.x sites must have sequences set to use ORDER

When performing Content Indexing with Oracle RAC, unless the sequences are in order, on rare occasions there could be a one-time lack of order in the sequences which would manifest itself in a few documents not being indexed.

Existing CE installations with Oracle RAC must issue “alter sequence CISEQUENCE CACHE 1000 ORDER,” where the CACHE 1000 command is added for performance and the ORDER command is required to ensure that the sequences are in order.

To set sequences to use ORDER

Perform the following on each object store:

1. Shut down all Content Engine services accessing the object store
2. Login via sqlplus as the object store schema owner, and issue the following SQL command:

   ```
   alter sequence CISEQUENCE CACHE 1000 ORDER
   ```

3. Restart the Content Engine services

NOTE You can perform the above steps at any time. For new object stores in a RAC environment, set the sequence order immediately after creating the object store.

Oracle 9i RAC must use Broadcast on Commit (BOC)

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

Core Components Installation Tasks

High Availability Installation Order/Priority

FileNet P8 components can be installed/configured for high availability in any order. To make it easier to verify functionality of as many components as possible, the recommended order for installing a full environment is as follows:

1. Database Services

   This typically involves third-party software and does not involve any IBM software. Usually the database server instance and necessary tablespaces are created before any IBM software is installed. This makes it easier to troubleshoot problems that are not related to IBM software.

2. Global Configuration Data

   The Content Engine installation must create information for the FileNet P8 Domain (explained in the Content Engine section) in a highly available location; therefore, this should be set up and running before installing any Content Engine components. For more information, see "Content Engine Global Configuration Data (GCD)" on page 26.

3. Content Engine Object Store Service Farm/Cluster

   This is the front end service needed to gain access to content and is likely the first Content Engine Service installed. This installation will create the FileNet P8 Domain using the highly available Global Configuration Data location. Application Engine requires that the Content Engine Object Store Service is installed and running to enable log in through Workplace. Object stores cannot be created until an Object Store Service is online. For more information, see "Object Store Service" on page 29.

4. Content Engine File Store Service Cluster

   This would most likely follow the installation and configuration of the Content Engine Object Store Service in order to add File Stores to existing or new object stores. The Object Store Service must already be installed in order to create File Stores. For more information, see "File Store Service" on page 39.

5. Process Engine

   The Process Routers on the Application Engine server must have a Process Engine up and running to establish a connection. Installing and configuring a highly available Process Engine at this point makes it possible to verify Process Engine functionality after Application Engine comes online. For more information, see "Process Engine" on page 52.

6. Application Engine Server Farm/Cluster

   After installation, Application Engine needs all other components online in order to log in, browse content, and perform any activities with the Process Engine. This is the most common component to be installed or configured last. For more information, see "Application Engine" on page 67.

7. FileNet functional expansions

   You can also configure FileNet functional expansions, such as Records Manager, for high availability. For more information, see “Functional Expansion Components Installation Tasks” on page 80.

Content Engine

The Content Engine provides access and storage facilities for content. When a critical service 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, making the health and stability of the Content Engine a key factor in creating a highly available IBM FileNet P8 environment.

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

- Content Engine Global Configuration Data
- Object Store Service
- File Store Service

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

Content Engine Global Configuration Data (GCD)

A Content Engine installation can involve a number of components including: IBM FileNet Enterprise Manager (EM), Object Store Service, and File Store Service. Also, more than one server may be running the same components, such as two object store servers, two EM clients, etc. The IBM FileNet P8 domain is the collection of Content Engine Services and servers on which services are hosted. For an overview of the IBM FileNet P8 Domain architecture, go to the *IBM FileNet P8 Platform* help and navigate to **IBM FileNet P8 Administration > Content Engine Administration > IBM FileNet P8 Domain**.

The Content Engine Global Configuration Data (GCD) contains information for the IBM FileNet P8 domain, such as:

- Location of Content Engine Services
- Object stores in use in the domain
- File Stores in use in the domain
- Authentication provider, such as the type and location of a directory server from which to authenticate

When the GCD is inaccessible, Content Engine servers and object stores can not be added to a IBM FileNet P8 domain. The GCD can be made highly available by placing GCD data on a clustered file share.

Procedure overview for GCD

The following high-level steps are required to make the GCD highly available:

1. Add a fileshare resource to a cluster.

2. Ensure that the clustered share can be accessed from a remote machine, or the machine on which you will be installing Content Engine.
3. Set security on the fileshare.

The following topics use these high-level steps to make the GCD highly available for a particular cluster package.

GCD - VERITAS Cluster Server

VERITAS uses “Service Groups” of resources to provide highly available services to users. The following procedure specifies how to add resources to a pre-existing service group to provide a highly available clustered resource that will store the GCD.

Pre-Deployment Tasks

1. Install the clustering software prior to installing or configuring any IBM FileNet P8 services for high availability. This includes VERITAS Cluster Server, VERITAS Volume Manager, data replication software, and backup software.
2. 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)

     NOTE The mount should have at least 20 MB free space.
   - A clustered IP resource
   - A Lanman name resource
3. Verify that each node in the cluster has access to the network or clustered storage used to store the files and folders that will comprise the GCD.
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.
5. Create the directory where the GCD will reside on clustered storage.

   Set the following permissions for the directory:
   - Full Control for the group defined as Content Engine Servers.
   - Full Control for all Windows administrative users/groups that will directly maintain files (i.e., do backups and restores).

   Give the share itself one of the following permissions:
   - Full Control for the Authenticated Users group (restrictive permissions on the folder will keep out unauthorized users).
   - Full Control for Content Engine Servers and Windows administrative users/groups.

   For more information, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 4: Create groups and users required for IBM FileNet P8 installation.

Deployment Tasks

This procedure expands the cluster group to include resources needed for the GCD.

1. Add a Lanman resource to ensure that clients can connect to the server using the cluster name.
   - Set IPResName to the name of the IP resource already in the cluster group.
   - Set VirtualName to the DNS name assigned to the clustered IP.
   - Make the Lanman resource dependent on the IP resource.
2. Create a new “FileShare” VERITAS resource for this directory.
   - Configure the resource parameters as appropriate (PathName, ShareName, and MountResName should be set based on the environment requirements and use the directory created in the previous step. At a minimum the share’s security permissions must grant FULL_CONTROL to Administrators (IBM FileNet P8 Administrators and Content Engine Administrators) and to the Content Engine Servers group).
   - Make the FileShare resource dependent on the Lanman and clustered storage resources.
3. Bring the cluster group online.

Deployment Verification

CAUTION You must use this share as the SYSINIT location when running the Content Engine Global Configuration Data (GCD) Wizard as part of the Content Engine installation.

a. To verify that the GCD share is working correctly, verify that you can access the share using:

```
\\<DNS name of clustered IP>\<fileshare name>
```

b. Fail the group to another node, and verify that you can still access the share.

If this brings up the shared folder, the share is working and ready to accept the GCD data.

GCD - Microsoft Cluster Server

Microsoft Cluster Server provides a framework to create “Groups” of resources that collectively provide a highly available service. The following procedure details how to add to a Group of resources to provide a highly available clustered disk to store the GCD.

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, see http://www.microsoft.com/whdc/hcl/.
2. Install the clustering software prior to installing or configuring any IBM FileNet P8 services for high availability.
3. Verify that a cluster group already exists with the following resources (at a minimum):
   - A clustered disk resource

     NOTE The mount should have at least 20 MB free space.
   - A clustered IP resource

- A Network Name resource

4. Verify that each node in the cluster has access to the network storage or clustered storage used to store the files and folders that will comprise the GCD.
5. Verify that the cluster group can failover to all nodes in the cluster, and that the clustered storage can be accessed from all nodes.
6. Create the directory where the GCD will reside on clustered storage.

   Set the following permissions for the directory:

   - Full Control for the group defined as Content Engine Servers.
   - Full Control for all Windows administrative users/groups that will directly maintain files (i.e., do backups and restores).

   Give the share itself one of the following permissions:

   - Full Control for the Authenticated Users group (restrictive permissions on the folder will keep out unauthorized users).
   - Full Control for Content Engine Servers and Windows administrative users/groups.

   For more information, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 4: Create groups and users required for IBM FileNet P8 installation.

Deployment Tasks

1. Create a new “File Share” resource for this directory.
2. Configure the resource parameters as appropriate (Share Name and Path should be set based on the environment requirements. At a minimum the share’s security permissions must grant FULL_CONTROL to Administrators (IBM FileNet P8 Administrators and Content Engine Administrators) and to the Content Engine Servers group).
3. Make the File Share resource dependent on the Network Name and clustered disk resource.
4. Bring the cluster group online.

Deployment Verification

CAUTION You must use this share as the SYSINIT location when running the Content Engine Global Configuration Data (GCD) Wizard as part of the Content Engine installation.

a. To verify that the GCD share is working correctly, verify that you can access the share using:

```
\\<DNS name of clustered IP>\<fileshare name>
```

b. Fail the group to another node, and verify that you can still access the share.

   If this brings up the shared folder, the share is working and ready to accept the GCD data.

Object Store Service

The Object Store Service manages content for database object stores and is the main service required for authenticating with an LDAP repository. If the Object Store Service goes offline users will not be able to login from Application Engine or IBM FileNet Enterprise Manager, and content cannot be accessed (including FileStore content) even if all other services are up and running elsewhere. The Object Store Service can be started, stopped, and controlled like most other Windows Services.

This section covers the tasks and requirements needed to make the Object Store Service highly available in a cluster or server farm environment. The goal of making the Object Store Service highly available is to continue providing access to Content Engine data (both database and filestore data) for all applications.

Object Store Service HA Requirements

The Object Store Service can be made highly available using either cluster (active-passive) or server farm (active-active) technologies.

Installing remote IBM FileNet Enterprise Manager clients

When installing a remote IBM FileNet Enterprise Manager client, join the IBM FileNet P8 domain by entering the physical machine name (not a virtual IP or name) of another server with Content Engine Object Store Service installed and running.

Configure Content Engine COM Client stations to not start the object store service on a passive node

In a Content Engine active/passive cluster environment the object store service should only be running on the active node. If the service is launched on a passive node unexpected behavior can occur.

CAUTION Do not manually launch IBM FileNet Enterprise Manager on a passive node, as this will launch the object store service on that node.

Example: In a cluster environment, if the object store service is running on both the active and passive nodes, subscription workflows could queue on the passive node, effectively locking these subscriptions from being processed.

To configure COM client stations in an active/passive Content Engine cluster environment to not start the object store service on passive nodes

Point the nodes to the virtual IP of the Content Engine cluster group on the CE COM clients by adding the virtual IP address of the Content Engine cluster group to the etc\hosts file for each node.

1. Open the hosts file for editing (typical location):

   ```
   WINDOWS\System32\drivers\etc\hosts
   ```

2. Add/Edit the host file entry for each CE node to point to the virtual IP of the Content Engine cluster group instead of pointing to the physical IP of the CE nodes.

Configure COM Clients in a farmed environment to use a specific Object Store service

By default, COM Clients in a CE Farm environment will dynamically select an Object Store service in the farm environment. Depending on your farm setup, certain farm servers may not be ideally situated for this behavior.

Example:

A COM Client may select a farmed server at a remote site where communication lag with the remote server could lower the performance of the client. The CE COM clients do not have any information about network performance or the proximity of the servers, and will sometimes dynamically select the Object Store services at the remote site, instead of using one that is locally available.

To avoid this possibility you can configure your farmed environment to only use a certain CE server for the Object Store service by adding the CE server name to the registry of the farmed servers.

To add the CE server name to the registry entry

1. Verify that you have applied CE 3.5.2 on all the CE COM client stations.
2. Add the Content Engine server name to the registry entry:

   Under the key:

   `HKLM\Software\FileNet\ECM\Content Engine\Client`

   add a new string value `UseServer` and set this value to the (Netbios) name of the Content Engine server.

NOTE This solution has also been known to work for the "Configure Content Engine COM Client stations to not start the object store service on a passive node" on page 30 issue described above, but has not been fully qualified by the IBM testing team. In this case, set the value of the registry key to the virtual network name of the CE cluster group.

Procedure overview for Object Store Service

The following high-level steps are required to make the Object Store Service highly available.

Making the Object Store Service highly available in a *server farm* involves the following:

- Install Content Engine on all nodes.

Making the Object Store Service highly available in a *cluster* involves the following:

1. Install Content Engine on all nodes in the cluster.
2. Add cluster resources for the following:
   - Content Engine Object Store Service
   - API Listener (Apache2 service)
   - Distributed Transaction Coordinator Service
   - Oracle MTS Service (if using Oracle RDBMS)

   Optional cluster resources that rely upon the Object Store Service:

   - Process Services Manager
   - FileNet P8 CFS Server for Image Services

The following topics use the preceding high-level steps to make the Object Store Service highly available.

Object Store Service - Server Farm

A server farm is a collection of servers that functions as a normal individual server with a front-end layer providing load-balancing capabilities that distributes requests to all servers in the farm. Should one server go down, traffic is automatically directed to another server in the farm. The load balancing layer can be either hardware (Big IP, Cisco Local Director) or software (Microsoft NLB).

NOTE An Object Store Service Farm cannot be collocated on the same machines as a File Store Service Cluster.

Pre-Deployment Tasks

1. Verify that each server in the farm has the same configuration, including OS levels, fix packs, and software (both IBM and third-party). Identically configured servers make automatic re-direction within the server farm transparent to the end user.

2. Set up and configure load-balancing software/hardware prior to Content Engine installation.

   - Installation procedures for Microsoft NLB can be found in Microsoft Knowledge Base articles #240997 (Windows 2000) and #323431 (Windows 2003)

   - If you are using Microsoft NLB, make sure each node in the farm has at least two network interface cards with one dedicated to direct connection (private network) within the nodes, and the other for the public network.

Deployment Tasks

1. Install Content Engine on all nodes in the server farm. See *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8a: Install Content Engine (Typical) or Task 8b: Install Content Engine (Custom).

   NOTE If this is the first installation of any Content Engine services in the configuration, then the GCD wizard will help you create the IBM FileNet P8 domain. For all subsequent installations of Content Engine services, the GCD wizard will help you join the IBM FileNet P8 domain. When joining the IBM FileNet P8 domain, enter the physical machine name (not a virtual IP or virtual name) of another server with Content Engine Object Store Service installed and running.

   You can perform either a Typical or a Custom installation. At a minimum the following components must be installed:

   - Content Engine Object Store Service

   - Client connectivity

   CAUTION If you perform a Typical installation, make sure to disable the File Store Service after installation.

2. Set the ObjectStore and Apache2 services to automatic startup.

3. Start the Process Routers.

   In order to provide access to the same Process Engine regions for all users, routers must be started on all machines in the server farm. Each router should maintain the same name, region, and Process Engine Server connection information across all servers in the farm.

Deployment Verification

NOTE To complete this verification, a database has to be available and configured.

1. Bring up the IBM FileNet Enterprise Manager and login as a Content Engine Administrator.
2. Create an object store.
3. Add a folder and document to the object store.
4. If an Application Engine has been setup and configured for this server farm then sign-in to verify that a connection is being made to the API Listener.
5. Fail one of the nodes and verify that remote IBM FileNet Enterprise Manager nodes and Application Engines can still login and view content.

Object Store Service - VERITAS Cluster Server

VERITAS 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 Object Store Service.

Pre-Deployment Tasks

1. Verify that a service group exists with the following resources:
   - Clustered storage resources (can include: VERITAS Volume Group (VMDg) and MountV, or Mount for a basic disk)

     NOTE This is only necessary if this environment will support subscriptions.
   - A clustered IP resource
   - A Lanman name resource
   - Clustered Distributed Transaction Coordinator (DTC) service
2. Verify that the service group can failover to all nodes in the cluster prior to installing Content Engine.
3. Decide the service account policy for the Content Engine services in your VERITAS Cluster Server.

   When you configure the resources you must enter the service account information. For clusters you have two options when selecting the service account:
   - Each node in the cluster utilizes its own (local) domain service account to run Content Engine services.
   - All nodes in the cluster uses the same (global) domain service account to run Content Engine services.

NOTE For more information, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 4: Create groups and users required for IBM FileNet P8 installation.

Deployment Tasks

1. Install Content Engine on all nodes in the cluster. See *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8a: Install Content Engine (Typical) or Task 8b: Install Content Engine (Custom).

   NOTE If this is the first installation of any Content Engine services in the configuration, then the GCD wizard will help you create the IBM FileNet P8 domain. For all subsequent installations of Content Engine services, the GCD wizard will help you join the IBM FileNet P8 domain. When joining the IBM FileNet P8 domain, enter the physical machine name (not a virtual IP or virtual name) of another server with Content Engine Object Store Service installed and running.

   - Ensure that the service group is up and running on the node you are installing on.
   - You can perform either a Typical or a Custom installation. At a minimum the following components must be installed:
     - Content Engine Object Store Service
     - Client connectivity
   - If you perform a Typical installation and will not be using File Stores on the cluster, make sure to disable the File Store Service after installation.
   - If you plan to use Fixed Content Devices utilizing Image Services, perform a Custom install and be sure to also select FileNet P8 Content Federation Service Server for Image Services.

   CAUTION After the Content Engine installation finishes you will be prompted to reboot the machine. This reboot will most likely fail the cluster group to another node. Before you login to the machine after the reboot to run the GCD Wizard, 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 configure the GCD using the GCD Wizard.

2. After installation set the startup type for Content Engine Object Store Service to Manual on all nodes.
3. Stop the Content Engine Object Store service.
4. Add a Lanman resource to ensure clients can connect to the server using the cluster name.
   - Set IPResName to the name of the IP resource already in the cluster group.
   - Set VirtualName to the DNS name assigned to the clustered IP.
   - Make the Lanman resource dependent on the IP resource.
   - Bring the Lanman resource online and try to access the cluster again using:

     ```
     \\<DNS name of clustered IP>
     ```

5. Add the following resources to the cluster:
   - API Listener resource
     - Resource type “GenericService.”
     - Set ServiceName parameter to “Apache2.”
   - Distributed Transaction Coordinator resource
     - Resource type “GenericService.”
     - Set the ServiceName parameter to “Distributed Transaction Coordinator.”
   - OracleMTSRecoveryService (only required if using Oracle rdbms)
     - Resource type “GenericService.”
     - Set the ServiceName parameter to “OracleMTSRecoveryService.”
   - Content Engine Object Store Service
     - Resource type “GenericService.”
     - Set the ServiceName parameter to “Content Engine Object Store Service.”
     - Make this resource dependent on the Lanman, OracleMTSRecoveryService, and Distributed Transaction Coordinator resources.
     - Set the Domain parameter to the fully qualified domain name of Active Directory.
     - Set the UserAccount parameter using one of the two following setups:
       - Only one service account runs Content Engine Services on all nodes.
         - Enter the account name for the UserAccount parameter.
         - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account.
       - Separate service accounts exist for each node running Content Engine Service.
         - Make the UserAccount and Password resource attributes local. This makes the attributes unique to each node.
         - Enter the account name for the UserAccount parameter for each node.
         - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account for each node.
6. Make the clustered IP resource dependent on the API Listener resource.
7. (Optional) Enable subscriptions.

   a. Modify the registry.

   Modify the registry on each node to specify the cluster name that gets assigned ownership of subscriptions.

Perform the following on every server in the object store cluster if you plan to enable subscriptions in this environment.

i. Open the registry and navigate to:

HKEY_LOCAL_MACHINE\SOFTWARE\FileNet\ECM\Content Engine

ii. Create a new string called veritas_virtual_name, whose value is your service group's cluster name (Lanman VirtualName).

b. Add cluster resources.

Add the following cluster resources to your Object Store Service group and perform the following configuration steps if you plan to enable subscriptions.

i. Create a cluster resource for registry replication.

- Resource type "RegRep."
- Set MountResName to the mount resource for the clustered disk.
- Leave the ReplicationDirectory location on the clustered disk as default or set as desired.
- Set Keys to:

```
HKLM\\SYSTEM\\CurrentControlSet\\Services\\VWServices
```

- Make this resource dependent on the clustered MountV resource.

ii. Create a cluster resource to control Process Services Manager.

- Resource type "GenericService."
- Set the ServiceName parameter to "VWServices."
- Make this resource dependent on the RegRep resource.
- Set the Domain parameter to the fully qualified domain name of Active Directory.

iii. Configure your routers to start automatically

- From the Start menu, open **FileNet P8 Platform > Process Services Administrator**.
- Double-click each router entry to bring up the properties dialog.
- Check **Automatically started** and **Automatically restarted** for each router.

8. Bring all services in the group online.

Deployment Verification

NOTE To complete this verification, a database has to be available and configured.

1. Bring up the IBM FileNet Enterprise Manager and login as a Content Engine Administrator.
2. Create an object store.
3. Add a folder and a document to the object store.

4. If an Application Engine has been setup and configured for this server farm sign in to verify that a connection is being made to the API Listener.

5. Fail one of the nodes and verify that remote IBM FileNet Enterprise Manager nodes and Application Engines can still login and view content.

Object Store Service - Microsoft Cluster Server

Microsoft Cluster Server provides a framework to create “Groups” of resources that collectively provide a highly available service. The following procedures describe how to add to a Group of resources in order to provide a highly available Object Store Service.

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, see http://www.microsoft.com/whdc/hcl/.

2. Install the clustering software prior to installing or configuring any IBM FileNet services for high availability.

3. Verify that a cluster group already exists with the following resources (at a minimum):

   - A clustered disk resource

     NOTE The MSDTC resource is dependent on a clustered storage resource.

   - A clustered IP resource

   - A Network Name

   - A Distributed Transaction Coordinator Resource

     NOTE On Windows 2000 this resource is created using the comclust utility.

4. Verify that the cluster group can failover to all nodes in the cluster prior to installing Content Engine.

Deployment Tasks

1. Install Content Engine on all nodes in the cluster. See *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8a: Install Content Engine (Typical) or Task 8b: Install Content Engine (Custom).

   NOTE If this is the first installation of any Content Engine services in the configuration, then the GCD wizard will help you create the IBM FileNet P8 domain. For all subsequent installations of Content Engine services, the GCD wizard will help you join the IBM FileNet P8 domain. When joining the IBM FileNet P8 domain, enter the physical machine name (not a virtual IP or virtual name) of another server with Content Engine Object Store Service installed and running.

   - Ensure that the cluster group is up and running on the node you are installing on.

   - You can perform either a Typical or a Custom installation. At a minimum the following components must be installed:

     – Content Engine Object Store Service

    - Client connectivity

- If you perform a Typical installation and will not be using File Stores on the cluster, make sure to disable the File Store Service after installation.
- If you plan to use Fixed Content Devices utilizing Image Services, perform a Custom install and be sure to also select FileNet P8 Content Federation Service Server for Image Services.

CAUTION After the Content Engine installation finishes you will be prompted to reboot the machine. This reboot will most likely fail the cluster group to another node. Before you login to the machine after the reboot to run the GCD Wizard, 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 configure the GCD using the GCD Wizard.

2. Add the *Content Engine Servers* group to the cluster security from one node in the cluster.

    a. Open the MSCS Cluster Administrator.

    b. Right-click the root node of the cluster and select **Properties**.

    c. Select the **Security** tab.

    d. Add the Content Engine Servers group to the cluster security, and give it “Full Control.”

3. Add the following resources to the cluster group:

    - API Listener resource
        - Resource type “Generic Service.”
        - Set Service name parameter “Apache2.”
    - OracleMTSRecoveryService (only required if using Oracle rdbms)
        - Resource type “Generic Service.”
        - Set Service name to “OracleMTSRecoveryService.”
    - Content Engine Object Store Service
        - Resource type “Generic Service.”
        - Set Service name to “Content Engine Object Store Service.”
        - Make this resource dependent on the OracleMTSRecoveryService and Distributed Transaction Coordinator resources.
        - Check **Use Network Name for computer name**.

4. (Windows 2000 only) Make the clustered IP resource dependent on the API Listener resource.

5. (Optional) Add the following cluster resources to your object store cluster group and perform the following configuration steps if you plan to enable subscriptions.

    a. Create a cluster resource to control Process Services Manager.

        - Resource type “Generic Service.”
        - Set the Service name to “VWServices.”

- Add values to the Registry Replication tab:

  `SYSTEM\CurrentControlSet\Services\VWServices`

b. Configure your routers to start automatically.

   i. From the Start menu, open **FileNet P8 Platform > Process Services Administrator**.

   ii. Double-click each router entry to bring up the properties dialog.

   iii. Check **Automatically started** and **Automatically restarted** for each router.

6. Bring all services in the group online.

Deployment Verification

NOTE To complete this verification, a database must be available and configured.

1. Bring up the IBM FileNet Enterprise Manager and login as a Content Engine Administrator.
2. Create an object store.
3. Add a folder and a document to the object store.
4. If an Application Engine has been setup and configured for this server farm, sign in to verify that a connection is being made to the API Listener.
5. Fail one of the nodes and verify that remote Enterprise Manager nodes and Application Engines can still login and view content.

File Store Service

The File Store Service manages and controls file store content, and can be made highly available using cluster technology to monitor and control the service. The File Store Service can be started, stopped, and controlled like most other Windows Services.

This section covers the tasks and requirements needed to make the File Store Service highly available in a cluster environment. The goal of this configuration is to continue providing access to file store content in the event that the active node in the cluster fails.

File Store Service HA Requirements

- Cluster Configuration.

  The File Store Service must run in an active-passive type of cluster configuration, with only one cluster node running the File Store Services at any one time.

- (Windows 2000 only) Set security on the cluster temp directory to allow content search indexing.

  To allow indexing of certain file types for content search on the object store, the Content Engine Servers group (which contains the user running the search engine) must have write access to the temp directory of the cluster user.

Procedure overview for File Store Service

The following high-level steps are necessary to make the File Store Service highly available:

1. Install Content Engine File Store Services on all nodes in the cluster.
2. Add the following resources to the cluster:
   - Content Engine File Store Service
   - Clustered FileShare for filestore data
   - Clustered Folder for filestore logs
3. Modify the registry to store File Store Service transaction logs on clustered storage.
4. Run the P8Cluster.exe tool to add the clustered network name into the IBM FileNet P8 domain and have it appear as a selectable server when creating File Stores.

The following topics use the preceding steps to make the File Store Service highly available for a particular cluster package.

File Store Service - VERITAS Cluster Server

VERITAS 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 File Store Service.

A File Store can be managed by only one File Store Service at any given time. This type of deployment constitutes an active-passive type of configuration. The active node in the cluster runs the File Store Service while the passive node remains without any services running until the active node fails.

Pre-Deployment Tasks

1. Install the clustering software prior to installing or configuring any IBM FileNet services for high availability. This includes VERITAS Cluster Server, VERITAS Volume Manager, data replication software, and backup software.
2. 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)
   - Clustered IP resource
   - Lanman name resource
   - Clustered Distributed Transaction Coordinator (DTC) service
3. Create two directories on the clustered disk: one to hold file store data (this directory will be shared), the other to hold the File Store Service transaction logs (this will not be shared).
4. Verify that the service group can failover to all nodes in the cluster, and that the clustered storage can be accessed from all nodes.

5. Decide the service account policy for the Content Engine services in your VERITAS Cluster Server.

   When you configure the resources you must enter the service account information. For clusters you have two options when selecting the service account:

   - Each node in the cluster utilizes its own (local) service account to run Content Engine services.
   - All nodes in the cluster uses the same (global) service account to run Content Engine services.

   NOTE For more information, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 4: Create groups and users required for IBM FileNet P8 installation.

Deployment Tasks

1. Install Content Engine on all nodes in the cluster. See *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8a: Install Content Engine (Typical) or Task 8b: Install Content Engine (Custom).

   NOTE If this is the first installation of any Content Engine services in the configuration, then the GCD wizard will help you create the IBM FileNet P8 domain. For all subsequent installations of Content Engine services, the GCD wizard will help you join the IBM FileNet P8 domain. When joining the IBM FileNet P8 domain, enter the physical machine name (not a virtual IP or virtual name) of another server with Content Engine Object Store Service installed and running.

   - Ensure that the service group is up and running on the node you are installing on.
   - You can perform either a Typical or a Custom installation. At a minimum the following components must be installed:
     - Content Engine File Store Service
     - Client connectivity
   - If you are installing on a dedicated File Store cluster, and will be running the Object Store Service on another cluster/server, make sure you disable the Object Store Service on this server.

   CAUTION After the Content Engine installation finishes you will be prompted to reboot the machine. This reboot will most likely fail the cluster group to another node. Before you login to the machine after the reboot to run the GCD Wizard, 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 configure the GCD using the GCD Wizard.

2. After installation, set the startup type for the Content Engine File Store Service to Manual on all nodes.
3. Stop the File Store service.
4. Add a Lanman resource to ensure clients can connect to the server through the cluster name.
   - Set IPResName to the name of the IP resource already in the cluster group.
   - Set VirtualName to the DNS name assigned to the clustered IP.
   - Make the Lanman resource dependent on the IP resource.

- Bring the Lanman resource online and try to access the cluster again using:

  `\\<DNS name of clustered IP>`

5. Create and configure security for two directories on the clustered storage.

   - One directory to hold File Store data.

     Set security for this directory to give full access to the following groups/users:

     – Content Engine Servers group

     – P8Apache user

     – The domain user that logged on to the server running the IBM FileNet Enterprise Manager client

   - One directory to hold the File Store Service transaction logs.

     Set security for this directory to give full access to the following groups/users:

     – Content Engine Servers group

     – P8Apache user

6. From the active node, modify the registry to specify the location for the storage directory where transaction logs will be created.

   a. Navigate to:

   HKEY_LOCAL_MACHINE\SOFTWARE\FileNet\ECM

   b. Create a new key named Content RM.

   c. Under the Content RM key, create a new string value named `Log File Directory`, whose value is the path to the directory for transaction logs you created in Step 5 on page 42.

7. Add a resource for the Registry Replication of the Content RM key.

   a. Set the Resource Type: RegRep

   b. Set MountResName to the mount resource for the clustered disk.

   c. Leave the ReplicationDirectory location on the clustered disk as default or set as desired.

   d. Set Keys to:

   ```
   HKLM\\SOFTWARE\\FileNet\\ECM\\Content RM
   ```

   e. Make the resource dependent on the following resource:

   - Clustered disk resource

8. Add the following resources to the cluster:

   - FileShare resource for FileStore data.

     – Resource type: FileShare.

     – Set PathName, ShareName, MountResName, and User_Permissions as appropriate for the environment and use the File Store data directory created in Step 5 on page 42.

  - The FileShare resource should be made dependent on the Lanman resource and the clustered storage resources.

- Content Engine File Store Service.
  - Resource type: GenericService.
  - Set ServiceName parameter to “Content Engine File Store Service.”
  - Make the resource dependent on the FileShare, registry replication, and the Content Engine Object Store Service resource (only if both resources are being collocated in the same service group), and a Distributed Transaction Coordinator resource.
  - Set the UserAccount parameter using one of the two following setups:
    - Only one service account runs Content Engine Services on all nodes.
      - Enter the account name for the UserAccount parameter.
      - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account.
    - Separate service accounts exist for each node running Content Engine Service.
      - Make the UserAccount and Password resource attributes local. This makes the attributes unique to each node.
      - Enter the account name for the UserAccount parameter for each node.
      - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account.

9. Run the P8Cluster.exe tool from one of the nodes.

   The P8Cluster.exe tool is located in the Content Engine installation folder; the default location is:

   C:\Program Files\FileNet\Content Engine

   The P8Cluster tool allows you to add a virtual cluster name to the list of available IBM FileNet P8 platform servers. When you create a File Store you can select the virtual cluster name for the location of the File Store Service instead of an individual node.

To add a virtual cluster name

i. For the Cluster Node Name, enter the actual machine name of the cluster node from which you are running the tool.

ii. For the Cluster Name enter the cluster virtual machine name, that is, the name (DNS name or IP address) that clients will use to connect to the cluster resources.

iii. Click **Add**.

To delete a cluster name (i.e. retiring a server, name change, etc.)

i. For the Cluster Node Name, enter the actual machine name of the cluster node from which you are running the tool.

  ii. For the Cluster Name enter the cluster virtual machine name, that is, the name (DNS name or IP address) that clients will use to connect to the cluster resources.

  iii. Click **Delete**.

NOTE The **Add** button is only enabled when the Cluster Node Name is an existing server with the File Store Service installed. The **Delete** button is only enabled when the Cluster Name is a cluster name previously added by the tool.

10. Bring up the services in the cluster group.

Deployment Verification

Verifying this configuration covers use cases for retrieving and adding content to the File Store. You should ensure that usage of the File Store is consistent even after a failover. The user sessions from IBM FileNet Enterprise Manager and Workplace should not expire or request new credentials.

1. Verify that the File Store Service is running.
2. Use IBM FileNet Enterprise Manager to Create a File Store.
   - Specify the virtual cluster name as the location of the File Store Service.
   - Specify the clustered share location for the location of the File Store data.
3. Add documents using IBM FileNet Enterprise Manager.
4. Retrieve documents using IBM FileNet Enterprise Manager.
5. Check out documents using IBM FileNet Enterprise Manager.
6. Add documents using Application Engine.
7. Retrieve documents using Application Engine.
8. Check out documents using Application Engine.
9. Failover the File Store Service to another node in the cluster.
10. Verify that the File Store is still usable after failover.
11. Add documents using IBM FileNet Enterprise Manager.
12. Retrieve documents using IBM FileNet Enterprise Manager.
13. Check out documents using IBM FileNet Enterprise Manager.
14. Add documents using Application Engine.
15. Retrieve documents using Application Engine.
16. Check out documents using Application Engine.

File Store Service - Microsoft Cluster Server

Microsoft Cluster Server provides a framework to create “Groups” of resources that collectively provide a highly available service. The following procedure describes how to add to a Group of resources in order to provide a highly available File Store Service.

NOTE A File Store can be managed by only one File Store Service at any given time. This type of deployment constitutes an active-passive type of configuration. The active node in the cluster runs the File Store Service while the passive node remains without any services running until the active node fails.

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, see http://www.microsoft.com/whdc/hcl/.
2. Install the clustering software prior to installing or configuring any IBM FileNet services for high availability.
3. Verify that a cluster group already exists with the following resources (at a minimum):
   - A clustered disk resource.
   - A clustered IP resource.
   - A Network Name resource.
   - A Distributed Transaction Coordinator Resource.

     NOTE On Windows 2000 this resource is created using the comclust utility.

   CAUTION The length of your virtual server's cluster name must not be longer than the network name resource for your quorum cluster group. If it is longer, truncation will occur when you try to use the P8cluster.exe tool.
4. Create two directories on the clustered disk: one to hold file store data (this directory will be shared), another to hold the File Store Service transaction logs (this will not be shared).
5. 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. Install Content Engine on all nodes in the cluster. See *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8a: Install Content Engine (Typical) or Task 8b: Install Content Engine (Custom).

   NOTE If this is the first installation of any Content Engine services in the configuration, then the GCD wizard will help you create the IBM FileNet P8 domain. For all subsequent installations of Content Engine services, the GCD wizard will help you join the IBM FileNet P8 domain. When joining the IBM FileNet P8 domain, enter the physical machine name (not a virtual IP or virtual name) of another server with Content Engine Object Store Service installed and running.

   - Ensure that the cluster group is up and running on the node you are installing on.
   - You can perform either a Typical or a Custom installation. At a minimum the following components must be installed:
     - Content Engine File Store Service
     - Client connectivity

NOTE If you are installing on a dedicated File Store cluster, and will be running the Object Store Service on another cluster/server, make sure you disable the Object Store Service on this server.

CAUTION After the Content Engine installation finishes you will be prompted to reboot the machine. This reboot will most likely fail the cluster group to another node. Before you login to the machine after the reboot to run the GCD Wizard, 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 configure the GCD using the GCD Wizard.

2. Add the *Content Engine Servers* group to the cluster security from one node in the cluster.

   a. Open the MSCS Cluster Administrator.

   b. Right-click the root node of the cluster and select **Properties**.

   c. Select the **Security** tab.

   d. Add the Content Engine Servers group to the cluster security, and give it “Full Control”.

3. Create and configure security for two directories on the clustered storage:

   - A directory to hold File Store data.

     Set security for this directory to give full access to the following groups/users:

     – Content Engine Servers group

     – P8Apache user

     – The domain user that logged on to the server running the IBM FileNet Enterprise Manager client

   - A directory to hold the File Store Service transaction logs.

     Set security for this directory to give full access to the following groups/users:

     – Content Engine Servers group

     – P8Apache user

4. From the active node, modify the registry to specify the location for the storage directory where transaction logs will be created.

   a. Navigate to:

      HKEY_LOCAL_MACHINE\SOFTWARE\FileNet\ECM

   b. Create a new key named Content RM.

   c. Under the Content RM key, create a new string value named `Log File Directory`, whose value is the path to the directory for transaction logs.

5. Add the following resources to the cluster:

   - A file share resource for FileStore data.

     – Resource type: FileShare.

- – Set Path and Share name as appropriate for the environment and use the File Store data directory.
- – Make the File Share resource dependent on the network name and shared disk resources.
- – Set permissions for the file store.

  For Windows Active Directory authentication, IBM recommends adding the following permissions on the shares containing local and remote file stores:

  - P8Apache user = Full Control
  - #AUTHENTICATED-USERS = Full Control
  - Other accounts as required for your specific environment

- A Content Engine File Store Service.
  - – Resource type: Generic Service
  - – Set the Service name parameter to “Content Engine File Store Service.”
  - – Check **Use Network Name for computer name**.
  - – Make the resource dependent on the File Share resource, the Distributed Transaction Coordinator resource, and the Content Engine Object Store Service resource (only if the Content Engine Object Store Service and File Store Service resources are being collocated in the same cluster group).
  - – Add the following value to the **Registry Replication** tab:

    ```
    software\FileNet\ECM\Content RM
    ```

6. Add a virtual cluster name.

   The P8Cluster tool allows you to add a virtual cluster name to the list of available IBM FileNet P8 platform servers. When you create a File Store you can select the virtual cluster name for the location of the File Store Service instead of an individual node.

To add a virtual cluster name

i. From one of the nodes, run the P8Cluster.exe tool.

ii. For the Cluster Node Name, enter the actual machine name of the cluster node from which you are running the tool.

iii. For the Cluster Name enter the cluster virtual machine name, that is, the name (DNS name or IP address) that clients will use to connect to the cluster resources.

iv. Click **Add**.

To delete a cluster name (i.e. retiring a server, name change, etc.)

i. From one of the nodes, run the P8Cluster.exe tool.

ii. For the Cluster Node Name, enter the actual machine name of the cluster node from which you are running the tool.

iii. For the Cluster Name enter the cluster virtual machine name, that is, the name (DNS name or IP address) that clients will use to connect to the cluster resources.

iv. Click **Delete**.

NOTE The **Add** button is only enabled when the Cluster Node Name is an existing server with the File Store Service installed. The **Delete** button is only enabled when the Cluster Name is a cluster name previously added by the tool.

7. Bring up the services in the cluster group.

Deployment Verification

Verifying this configuration covers use cases for retrieving and adding content to the File Store. You should ensure that usage of the File Store is consistent even after a failover. The user sessions from IBM FileNet Enterprise Manager and Workplace should not expire or request new credentials.

1. Verify that the File Store Service is running.
2. Use IBM FileNet Enterprise Manager to Create a File Store.
   - Specify the virtual cluster name as the location of the File Store Service.
   - Specify the clustered share location for the location of the File Store data.
3. Add documents using IBM FileNet Enterprise Manager.
4. Retrieve documents using IBM FileNet Enterprise Manager.
5. Check out documents using IBM FileNet Enterprise Manager.
6. Add documents using Application Engine.
7. Retrieve documents using Application Engine.
8. Check out documents using Application Engine.
9. Failover the File Store Service to another node in the cluster.
10. Verify that the File Store is still usable after failover.
11. Add documents using IBM FileNet Enterprise Manager.
12. Retrieve documents using IBM FileNet Enterprise Manager.
13. Check out documents using IBM FileNet Enterprise Manager.
14. Add documents using Application Engine.
15. Retrieve documents using Application Engine.
16. Check out documents using Application Engine.

Content Federation Services for Image Services (CFS-IS)

CFS-IS integrates Content Engine and Image Services based systems with a deep native integration. CFS-IS enables Content Manager applications and servers to communicate and work together with Image Manager applications and servers. You may initiate work through either system. For example, you can capture documents through IBM FileNet P8 Platform applications, store them in Image Services, and view them with other P8 applications. Or you may capture documents in Image Services, send the cataloging information to Content Engine, and view the documents through an Image Manager viewer such as IDM Desktop and IBM FileNet P8 Platform applications such as Workplace.

CFS-IS is an optional component installed with Content Engine that can be selected by performing a Custom installation of Content Engine. CFS-IS needs to be installed on every server where the Content Engine Object Store Service or File Store Service is installed. In general it is recommended to cluster the CFS-IS service within the Content Engine File Store Service cluster group. On servers running Content Engine Object Store Service it is recommended to use the high availability solution already in place (for example, farm the CFS-IS service if using Object Store Service farms, or cluster the CFS-IS service if clustering the Object Store Service).

One exception is when using a Symmetric cluster with Object Store Services in one cluster group and File Store Services in another cluster group. If both groups are capable of running on the same servers at the same time then the CFS-IS service must be left un-clustered (since we cannot control the same service from two different cluster groups). In this configuration CFS-IS should be configured as an automatically started Windows service.

The following sections provide information on how to cluster CFS-IS in Microsoft and VERITAS cluster.

Content Federation Services for Image Services (CFS-IS) - VERITAS Clusters

VERITAS 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 in order to make CFS-IS highly available.

Pre-Deployment Tasks

1. Verify that IBM FileNet Content Federation Services for Image Services (CFS-IS) has been installed on all Content Engine servers running Object Store Services or File Store Services.

2. Verify that each server that has IBM FileNet Content Federation Services for Image Services (CFS-IS) installed also has entries in its host file for the Image Server systems to be used in the IBM FileNet P8 domain.

Deployment Tasks

Add a new resource to the cluster to control the CFS-IS service if your environment will be utilizing Image Services as Fixed Content Devices.

CAUTION Do not put the “FileNet P8 CFS Server for Image Services” under cluster control until you have successfully created at least one fixed file store that is using Image Services as its Fixed

Content Device and mapped at least one set of properties. The CFS-IS service will not start until such a fixed file store has been created and properties have been mapped.

Configure the resource for CFS-IS as follows:

- Resource type “GenericService.”
- Set the ServiceName parameter to “FileNet P8 CFS Server for Image Services.”
- Set the Domain parameter to the fully qualified domain name of Active Directory.
- Set the UserAccount parameter using one of the two following setups:
  - Only one service account runs Content Engine Services on all nodes.
    - Enter the account name for the UserAccount parameter.
    - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account.
  - Separate service accounts exist for each node running Content Engine Service.
    - Make the UserAccount and Password resource attributes local. This makes the attributes unique to each node.
    - Enter the account name for the UserAccount parameter for each node.
    - Enter the vcs encrypted password in the cluster configuration instead of a clear text password for the service account for each node.

Deployment Verification

Verifying this configuration involves executing use cases for capturing documents and retrieving content. You should ensure that functionality of CFS-IS is consistent before and after a failover.

Test the system by capturing a small number of documents, checking the IS connection Monitor tool on the CE Server to make sure the documents are committed on the IS system, and retrieving them from either the CE or IS system or both depending on your configuration.

Content Federation Services for Image Services (CFS-IS) - Microsoft Clusters

Microsoft Cluster Server provides a framework to create “Groups” of resources that collectively provide a highly available service. The following procedure describes how to add to a Group of resources in order to make CFS-IS highly available.

Pre-Deployment Tasks

1. Verify that IBM FileNet Content Federation Services for Image Services (CFS-IS) has been installed on all Content Engine servers running Object Store Services or File Store Services.
2. Verify that each server that has IBM FileNet Content Federation Services for Image Services (CFS-IS) installed also has entries in its host file for the Image Server systems to be used in the IBM FileNet P8 domain.

Deployment Tasks

Add a new resource to the cluster to control the CFS-IS service if your environment will be utilizing Image Services as Fixed Content Devices.

CAUTION Do not put the “FileNet P8 CFS Server for Image Services” under cluster control until you have successfully created at least one fixed file store that is using Image Services as its Fixed Content Device and mapped at least one set of properties. The CFS-IS service will not start until such a fixed file store has been created and properties have been mapped.

Configure the resource for CFS-IS as follows:

- Resource type “GenericService.”
- Set the Service name parameter to “FileNet P8 CFS Server for Image Services.”

Deployment Verification

Verifying this configuration involves executing use cases for capturing documents and retrieving content. You should ensure that functionality of CFS-IS is consistent before and after a failover.

Test the system by capturing a small number of documents, checking the IS connection Monitor tool on the CE Server to make sure the documents are committed on the IS system, and retrieving them from either the CE or IS system or both depending on your configuration

Process Engine

This section covers the tasks and requirements for making the Process Engine highly available in a cluster environment.The Process Engine allows users to create, manage, control, and participate in workflows. When the Process Engine is unavailable, Application Engine or Content Engine cannot connect to the Process Engine to 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 continue running Process Engine Services and provide service to Application Engine and Content Engine in the event that the active node in the cluster 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.

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

- Process Service
- Pooled Process Manager (PPM)

Process Engine HA Requirements

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

Procedure Overview

The procedure for making the Process Engine highly available differs by platform.

UNIX Platforms

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. Install the Process Engine on all nodes.
3. Configure the Process Engine for the cluster.
4. 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 *IBM 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)
- OCOR_Listen (required)

  NOTE There may be many of these running; monitoring any one is sufficient.
- SEC_Daemon (optional)
- NCH Daemon (optional)
- MKF Write (optional)
- MKF Clean (optional)

Windows Platforms

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

1. Install the Process Engine on all nodes.
2. Configure the Process Engine for the cluster.
3. Add resources to the cluster to control the Process Engine.
   - IMS® Control Service
   - Process Engine Services Manager
   - Registry replication

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.

VERITAS 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.

VERITAS provides a framework to create custom agents that can control and monitor particular resources. IBM 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 VERITAS 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.

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. Verify that a cluster group already exists with the following resources (at a minimum):
    - Clustered storage resources (can include VERITAS Volume Group and/or Mount)
    - 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. Create cluster resources for the Process Engine partitions on the clustered storage, using the permission and size settings as documented in the *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide*.
    a. To ensure that each node uses consistent information and provides the same data to all users, create the following partitions to store configuration data on the clustered storage:
        - local
        - fn_sec_db0
        - fn_sec_rl0

    b. IBM recommends that you keep the fnsw directory on a local (non-clustered) storage. Create the following partition on local storage on each node:
        - fnsw
2. Install the Process Engine on all nodes in the cluster.
    a. On the first node in the cluster do the following:
        i. Install Process Engine. For detailed instructions, see the *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 11*x*: Install Process Engine (*Operating System*).
        ii. Once the installation completes, and you have completed all post-installation Process Engine specific tasks, stop the Process Engine services manually, then failover the cluster to another node.

    b. On all other nodes do the following:
        i. Once the cluster group is up and running on the node, remove all files and directories underneath **/fnsw/local**.
        ii. Run the Process Engine installation, specifying the same information used when installing the first node.
        iii. Once the installation completes, and you have completed all post-installation Process Engine specific tasks, stop the Process Engine services manually, then failover the cluster to another node.

3. Configure the Process Engine for the cluster.

   a. On one of the cluster nodes, 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.

   e. At this point, configure the LDAP connection. For more information, see *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 13: Configure the Process Engine LDAP Connection.

4. Install the FileNet P8 3.5.0 Process Engine Agent.

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

   a. Untar the FileNet Process Engine Agent located under the VERITAS folder on the FileNet P8 Process Engine software package.

   b. Copy the files to all nodes in the cluster.

   c. On all nodes perform the following steps:

      i. Create the 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.

  i. Set the Process Engine resource attributes, as required by your environment:

  - Oracle_SID (required for Oracle) - SID of the Process Engine's Oracle database.
  - Oracle_Home (required for Oracle) - Path to Oracle binaries and configuration files.
  - Oracle_UID (required for Oracle) - The owner of the Oracle files and directories.
  - DB2Instance (required for DB2) - The user/owner of the DB2 client instance (usually fnsw).
  - DB2Home (required for DB2) - The location of the home directory for your DB2 Instance owner (usually fnsw).
  - PPM_Port (optional) - The port on which the PPM for the Process Engine will be started. If no value is specified the default port 32771 is used.
  - LocalHost (optional) - The DNS name or IP Address of the clustered IP resource.

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

  iii. Verify that the status of the resource is *not* "UNKNOWN".

  NOTE The status "UNKNOWN" usually means that one or more of the attributes are not set correctly; double check the resource attributes.

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.

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 any session connection or be prompted for login credentials.

1. Start a router on the local Process Engine node. Make sure that it connects without error.
   - Specify the DNS name of the clustered IP resource for the Process Engine server.
2. Start the Process Router on the Application Engine server.
   - Specify the DNS name of the clustered IP resource for the Process Engine server.

3. Log on to Workplace.
4. Create and launch a workflow.
5. Initiate a failover of Process Engine in the cluster.

   NOTE Always reboot the failed node to ensure services get killed properly.

6. Access the user's inbox and retrieve workflow information during/after the failover.

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

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

VERITAS Cluster Server (Windows)

VERITAS 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 the Process Engine on Windows, IBM 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 where configuration data will reside.

   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 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.

NOTE Before starting your installation of Process Engine, please follow the configuration steps outlined in "Appendix A – Setting up a Secure Native Mode Domain Installation" on page 115.

1. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 11*x*: Install Process Engine (*Operating System*).

   NOTE Run the Process Engine installer as a domain user created for installing Process Engine software as detailed in "Create IBM FileNet Users" on page 115. 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 the Process Service has 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**.
         - Start the Process Service from task manager to verify that the service starts correctly.

      iii. If this is the first installation in the cluster, configure the LDAP connection at this point. For more information, see *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 13: Configure the Process Engine LDAP Connection.

      iv. Stop the Process Service from the Process Task Manager.

v. 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.

vi. 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. Configure the 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. In the Process Task Manager, check the box to automatically start the PPM.

iii. Modify the registry to set the NCHBroadcast value

- Navigate to:

  HKEY_LOCAL_MACHINE\Software\FileNET\IMS\CurrentVersion

- Add a new DWORD Value:

  – Name: NCHBroadcast

  – Value: 0

3. 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.

4. Bring the cluster group online.

Deployment Verification

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

1. Start a router on the local Process Engine node, and make sure it connects without error.

   NOTE Specify the Lanman name for the Process Engine server.

2. Start the Process Router on the Application Engine server.

   NOTE Specify the Lanman name for the Process Engine server.

3. Log on to Workplace.
4. Create and launch a workflow.
5. Initiate a failover of the Process Engine in the cluster.

   NOTE Always reboot the failed node to ensure services get killed properly.

6. Access the user's inbox and retrieve workflow information during/after a failover.

   NOTE There might be a slight delay, but eventually the Process Router automatically connects to the Process Engine after the fail-over. Make sure to reboot the inactive node after fail-over to ensure a clean startup environment on each successive fail-over.

7. Complete a step in a workflow.
8. Create and launch another workflow.
9. 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 instructions 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, see http://www.microsoft.com/whdc/hcl/.
2. Verify that your cluster nodes are members of a Windows domain running in either Windows 2000 Native mode or Windows Server 2003 mode.

   CAUTION The Process Engine installation does not support Windows 2000 Mixed 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.

   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 IBM 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 failover 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. Please refer to the configuration steps outlined in "Appendix A – Setting up a Secure Native Mode Domain Installation" on page 115.

1. Install Process Engine using the following steps. For detailed instructions, see the *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 11*x*: 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 IBM FileNet Users" on page 115.

   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. If this is the first installation in the cluster, configure the LDAP connection at this point. For more information, see *IBM FileNet P8 Platform 3.5.x Installation and Upgrade Guide,* Task 13: Configure the Process Engine LDAP Connection.

iv. 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.

v. 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. 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. In the Process Task Manager, check the box to automatically start the PPM.

c. 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

3. Add resources to the cluster to control the Process Engine.

a. Add a resource for the IMSService.

i. Set the Resource Type: Generic Service

ii. Set the Service name to: IMSService

iii. Check the **Use Network Name for computer name** box.

iv. Make the resource dependent on the following resources:

- Clustered network name
- Clustered disk resources

v. Add values to the Registry Replication tab:

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

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

i. Set the Resource Type: Generic Service

ii. Set the Service name to: VWservicesPE

iii. Check the **Use Network Name for computer name** box.

iv. Make the resource dependent on the following resource:

- IMSService

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

```
system\CurrentControlSet\Services\VWServicesPE
```

c. Bring the cluster group online.

Deployment Verification

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

1. Start a router on the local Process Engine node, and make sure it connects without error.
   - Specify the Network Name resource for the Process Engine server.
2. Start the Process Router on the Application Engine server.
   - Specify the Network Name resource for the Process Engine server.
3. Log on to Workplace.
4. Create and launch a workflow.
5. Initiate a failover of the Process Engine in the cluster.

   NOTE Always reboot the failed node to ensure services get killed properly.
6. Access the user's inbox and retrieve workflow information during/after a failover.

   NOTE There might be a slight delay, but eventually the Process Router automatically connects to the Process Engine after the fail-over. Make sure to reboot the inactive node after fail-over to ensure a clean startup environment on each successive fail-over.
7. Complete a step in a workflow.
8. Create and launch another workflow.
9. Access the user's inbox and retrieve workflow information.

PE Database Reconnection

Each PE background process (VWRs, VWJs, VWSs, VWdone, VWtime, and vwaemsg) 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 may 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 elog-files, showing the progress of re-connecting 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.

Application Engine

Application Engine provides convenient front-end access to content throughout the IBM FileNet P8 platform for users who want to check in documents, perform searches, and publish content. By making Application Engine highly available you are ensuring that users will always have access to a running instance of Application Engine. The other instances of a running Application Engine will provide service to users while the original Application Engine instance is offline.

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

- Workplace application
- bootstrap.properties file
- Process Routers

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

Application Engine HA Requirements

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.

CAUTION If a load-balancer or proxy server is used in this configuration you must use the load-balancer or proxy URL when performing installation steps that require an application URL. If you use the URL of a specific AE server when logging in to Workplace for the first time, the Base URL setting, which must be set to the load-balancer or proxy URL, will be incorrectly configured.

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://<proxyurl>:<proxyport>/Workplace
```

Use this URL when completing the installation steps for Application Engine.

Procedure overview for Application Engine

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

1. Create a setup location for the bootstrap.properties file.
2. Install Application Engine on all nodes.
3. Configure the application server cluster for Application Engine.

4. Deploy Workplace.
5. Perform any necessary post-deployment steps, and restart Application Engine.
6. (If you enable User Tokens on your IBM FileNet P8 system) Make sure each Application Engine server has the same UTCryptoKeyFile.properties file.
7. Start Process Routers on all nodes in the server farm.

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

Pre-deployment tasks

The following procedure assumes that a cluster of server instances has been created and is running prior to deploying the IBM FileNet P8 Application Engine. Aside from the configuration data (bootstrap.properties), the rest of the Application Engine software does not need to be placed on any network or shared storage systems. The Application Engine software may be installed onto local storage for each machine in the cluster.

1. Verify that your cluster of servers is running.
2. Create a setup location for the bootstrap.properties file.

   To ensure that each Application Engine instance functions identically in a cluster, it is important that each instance access the same Application Engine configuration data. The Application Engine installer allows you to select the location of the bootstrap.properties file during installation.

   The bootstrap.properties file should ideally be placed on a highly available Windows share or NFS exported directory that is accessible by all systems in the cluster.

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

   NOTE Do not use one of the cluster servers for the file location as that 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 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 instances will be unable to find the bootstrap.properties file and will return error messages.
3. Mount the bootstrap.properties file location from all node servers in the cluster.

   Once the location of the Application Engine configuration data is decided, that directory must be mounted and accessible from all servers in the cluster. Also, ensure that the same directory path is used on all systems (i.e. if the path on the first system is **/home/clusterdata**, make sure this same path is available on all other systems).

Deployment Tasks

You must install Application Engine on all nodes in the cluster farm. Even though web applications get deployed to the farm through a single node, Application Engine also uses functionality outside of the web application such as the router and component manager. These are only available through a full Application Engine installation.

1. Install Application Engine on all nodes.

   Follow the instructions in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 14: Install Application Engine, with the following additional information:

   - (On the Content Engine Java API Configuration screen) If the Content Engine has been configured for high availability, make sure to specify the clustered/farmed Content Engine name when you specify the machine name of the Content Engine.
   - (On the Folder Location of bootstrap.properties screen) Specify the location created to hold Application Engine configuration data.

   NOTE JBoss is designed to cluster deployed applications easily. Most application servers communicate state over multicast so it is critical that JBoss clusters be configured to segregate Application Engine multicast traffic from other JBoss clusters. See JBoss documentation for instructions regarding JBoss cluster configuration.

2. Configure and deploy the application server cluster for Application Engine.

   Follow the instructions in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 15*x*: Configure and start Application Engine (*Application server*), for each individual node server, with the following additional information.

   - When deploying Workplace, specify the cluster as the deployment location instead of individual server instances.

3. Perform any post-deployment configuration steps such as configuring symmetric encryption or setting up SSL security.

4. (If you enable User Tokens on your IBM FileNet P8 system) Verify that each Application Engine server has the same UTCryptoKeyFile.properties file.

   CAUTION For multiple applications to pass user tokens between one another, each participating application must use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all other Application Servers in the cluster.

   For general information about user tokens, go to the *Documentation for IBM FileNet P8 Platform* help and navigate to Developer Help > Workplace Integration and Customization > User Tokens > Configuring Applications to Use Tokens.

5. Restart the Application Engine application in the cluster after performing post-deployment configuration.

6. Configure and start routers on all machines in the cluster, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 18: Start Process Routers.

   In order to provide the same Process Engine data to each clustered Application Engine instance, routers will need to be started on all machines in the cluster. Use the same router configurations (name, region, and Process Engine Server connection information) across all Application Engine cluster machines.

NOTES

- Changes made to the Site Preferences > Process Router - Host:port/name field do not get automatically propagated to other Application Engine cluster machines in the server farm. You must make this change manually on each Application Engine, by either changing the value in Site Preferences or by manually editing the router information on each machine in the cluster.

  The router information is stored in the following file:

  – UNIX: **/opt/FileNet/Router/taskman.login.config**

  – WINDOWS: **C:\Program Files\FileNet\Router\taskman.login.config**

- If a proxy server is used in the cluster configuration, and the connection to the Workplace Process Configuration Console will be made via the proxy server (i.e. a user does not access the Configuration Console from an individual cluster node) then a Process Engine router must be started on the proxy server for each region that is to be administered and configured.

  For information, go to the *Documentation for IBM FileNet P8 Platform* help and navigate to FileNet P8 Documentation > Workplace > Admin tools > Site Preferences > General Preferences > General Settings - Tasks > Router.

Deployment Verification

1. Verify that the Workplace application is running on the cluster.
2. Open a web browser and type in the URL for the Workplace application
   - If using a proxy server the URL has the form:

     ```
     http://<proxy_server_name>:<port_number>/Workplace
     ```
   - Otherwise, if connecting to an individual cluster node the URL would have the form:

     ```
     http://<clustered_server_name>:<port_number>/Workplace/
     ```
3. Verify that the sign-in page displays.
4. Fail one of the nodes.
5. Verify that you can reload the sign-in page.

Managing Workplace settings in a load balanced environment

This topic contains documentation for managing workplace settings in a load balanced Application Engine environment consisting of a hardware load balancer, several HTTP servers and several application server instances.

Some Workplace Site Preference 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 server hosting the current user session, and the modified file(s) will not be distributed throughout the cluster or farm.

Currently, after making a configuration change, you must synchronize the configuration file(s) across your cluster or farm nodes, and then reload the Workplace configuration files individually on each node.

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 settings in a load balanced environment

1. From the Workplace Admin user interface, make your configuration changes.

   Depending on the type of configuration change made, one or more of the following files will be updated:

   - Actions.xml
   - PrimaryViews.xml
   - SystemPropertiesView.xml

   The default location for these files is the deployed directory for the Workplace application on your application server.

   NOTE The bootstrap.properties file is normally already set up to be shared between the servers in a highly available environment. This file does not need to be manually copied.

2. Back up the existing configuration files on the other node servers.

3. Copy the updated configuration file(s) to the cluster or farm nodes.

   Synchronize the files across the servers, overwriting existing files with the newly modified files.

4. Reload the Workplace configuration files

   After synchronizing the files you must reload the Workplace 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.

To reload settings

a. Sign into Workplace as a user who is assigned the Application Engine Administrator access role.

   Use the following URL to sign in to Workplace:

   ```
   http://<node_server_name>:<port#>/Workplace
   ```

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:

Developer Help > Workplace Integration and Customization > Appendixes > Reloading Workplace Configuration Files.

Service Pack Installation Tasks

This section provides the information needed to install IBM FileNet P8 3.5.x Services Packs in an existing IBM FileNet P8 3.5.0 highly available environment.

Additional Documentation

In addition to this technical notice you will need the following documents which provide instructions on how to apply the IBM FileNet P8 Platform Service Packs:

- *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide* - available on the IBM Information Management support www.ibm.com/software/data/support
- *Content Engine 3.5.x Service Pack Readme* - included on the Content Engine Service Pack software package.
- *Process Engine 3.5.x Service Pack Readme* - included on the Process Engine Service Pack software package.
- *Application Engine 3.5.x Service Pack Readme* - included on the Application Engine Service Pack software package.

High Availability Service Pack Installation Order/Priority

You must apply the service packs in a highly available environment in the following order:

5. Apply the service pack to the Content Engine Cluster/Farm. For more information, see "Content Engine 3.5.x service pack" on page 74.
6. Apply the service pack to the Process Engine Cluster. For more information, see "Process Engine 3.5.x service pack" on page 76.
7. Apply the service pack to the Application Engine Cluster/Farm. For more information, see "Application Engine 3.5.x service pack" on page 78.
8. Apply the service pack to available functional expansions.

   For more information, see "Functional Expansion Components Installation Tasks" on page 80.

Content Engine 3.5.x service pack

This section covers the installation of the Content Engine service pack in a highly available environment.

Service Pack Installation in an Object Store Service Farm

The load balancing software or hardware does not need to be specially configured or modified to install the service pack. Install the Content Engine service pack into the Object Store Service farm just like you would multiple standalone Content Engine Servers.

Service Pack Installation in a Content Engine Cluster

The service pack installation process for a Content Engine cluster is the same regardless of the services being clustered, cluster configuration in use, and cluster software being used (for example there is no difference between the service pack installation on a system using MSCS versus a system using VERITAS). Install the service pack on each node that has Content Engine software installed.

Service Pack Installation Tasks

1. Ensure that the GCD share, if clustered, is online when installing the Content Engine service pack.
2. Suspend cluster operations on the active cluster node so that the cluster groups do not failover when services are stopped/started during the service pack installation process. The specific procedure for this varies by cluster vendor (MSCS can "Pause" a node, while VERITAS can "freeze" a node).
3. Install the Content Engine service pack on the active cluster node.

   Follow the instructions listed in the "Installation Instructions" section of the *IBM FileNet P8 3.5.x Content Engine Service Pack Readme*.

   NOTE The log files for the File Store are most likely placed on shared storage. Perform the verification actions for Step 2 of "Allow all Content Engine-related transactions to complete" on the shared storage location.

   CAUTION The service pack installation process may involve a reboot of the active node. After the reboot, logon to another node in the cluster and fail the cluster group back to the node on which the Content Engine service pack is still being installed, then logon to this node to finish the installation of the Content Engine service pack.
4. After the service pack installation is finished, resume cluster operations on the current node. The specific procedure for this varies by cluster vendor (MSCS can "Resume" a node, while VERITAS can "unfreeze" a node).
5. Move the cluster group to another node.

6. Repeat steps 1 - 4 on all Content Engine servers in your cluster/farm. .

Process Engine 3.5.x service pack

This section covers the installation of the Process Engine service pack in a highly available environment.

Service Pack 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 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. See the "Stop all Process Engine-related services and applications" section of the "Installation Instructions" in the *Process Engine P8 3.5.x Service Pack Readme*.
3. Install the Process Engine service pack in the cluster by following the steps under "Update Process Engine (UNIX)" in the "Installation instructions" section of the *Process Engine P8 3.5.x Service Pack Readme*.
4. Once the service pack installation is finished on the node, 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 1-5 for all nodes in the cluster that run Process Engine services until the service pack has been completely installed on all nodes in the cluster.
7. Start Process Engine services.

Service Pack 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 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. See the "Stop all Process Engine-related services and applications" section of the "Installation Instructions" in the *Process Engine P8 3.5.x Service Pack Readme*.
3. Install the Process Engine service pack in the cluster by following the steps under "Update Process Engine (UNIX)" in the "Installation instructions" section of the *Process Engine P8 3.5.x Service Pack Readme*.
4. Once the service pack installation is finished on the node, resume cluster operations on the current node. The specific procedure for this varies by cluster vendor (for example MSCS can "Resume" a node).
5. Move the cluster group to another node.
6. Repeat steps 1-5 for all nodes in the cluster that run Process Engine services.

7. Start Process Engine services.

Application Engine 3.5.x service pack

This section covers the installation of the Application Engine service pack in a highly available environment.

The Application Engine service pack installation 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.

Service Pack Installation Tasks

1. Stop all Process Routers on all nodes in the Application Engine farm. See Steps 1-3 of "To install the Application Engine Service Pack 3.5.1" in the *Application Engine 3.5.x Service Pack Readme*.
2. Back up, undeploy, and remove the Workplace web application from the J2EE application server. See step 4 of the *Application Engine 3.5.x Service Pack Readme*. Perform this step for all nodes in the cluster/farm.

   NOTES

   - Make a backup copy of one of the deployed Workplace directories. Since the same copy of Workplace is deployed to all nodes in the cluster, you only need to make a backup copy of one of the deployed applications.
   - Step 4 of the *Application Engine 3.5.x Service Pack Readme* includes instructions to delete a temporary Workplace directory on some application servers. If this step applies for your application server then you must make sure to delete the temporary Workplace directories on every node in the cluster/farm.
3. Copy existing configuration files.

   See step 5 of the *Application Engine 3.5.x Service Pack Readme*. Copy these files to the installation location on each node.
4. Install the service pack on all nodes in the cluster/farm.

   See step 6 of the *Application Engine 3.5.x Service Pack Readme*.
5. Perform steps 2-6 under Deployment Tasks for Application Engine installation and deployment in a cluster/farm.
6. Perform the remaining service pack installation steps 8-12 in the *Application Engine 3.5.x Service Pack Readme* as applicable for your environment.

   CAUTION If a load-balancer or proxy server is used in this configuration you must use the load-balancer or proxy URL when performing 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://<proxyurl>:<proxyport>/Workplace
   ```

   Use this URL when completing the installation steps for Application Engine.

NOTE When starting Process Routers in step 11, make sure to start routers on every node in the cluster/farm.

Functional Expansion Components Installation Tasks

This section covers high availability installation and configuration of FileNet functional expansion components and their service packs.

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

- “Records Manager 3.5.0” on page 81.
- “Install the Records Manager 3.5.x service pack” on page 85.
- “Records Manager 3.7.0” on page 89.
- “Rendition Engine 3.5.0” on page 93.
- “Team Collaboration Manager 3.5.0” on page 97.

  CAUTION P8TCM-3.5.0-001 or later is required for Team Collaboration Manager to be used with a P8 3.5.1 Platform installation.

Additional Documentation

In addition to this technical notice you will need the following documents to set up the FileNet Functional Expansions for high availability:

- *IBM FileNet Rendition Engine Installation and Upgrade Guide* (PDF)
- *IBM FileNet Records Manager Installation and Upgrade Guide, version 3.5 or 3.7* (PDF)
- *IBM FileNet P8 Hardware and Software Requirements* (PDF)
- *IBM FileNet Team Collaboration Manager Installation Guide* (PDF)

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 12.

Records Manager 3.5.0

IBM FileNet P8 Records Manager 3.5.0 can be installed 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.

The following procedure covers the differences required for installing Records Manager in an application server cluster. The remainder of the IBM FileNet P8 components required to support Records Manager are installed as per the IBM FileNet P8 High Availability Technical Notice found on the IBM Information Management support site. 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 12

NOTES

- The following procedure will install IBM FileNet Records Manager 3.5.0 in a high availability environment. After installing version 3.5.0, to install IBM FileNet Records Manager 3.5.x, follow the instructions in "Install the Records Manager 3.5.x service pack" on page 85.
- The following procedure covers installation of the RM Web Application only.
- For instructions on installing RM components such as the RM Sweep Application or Crystal Reports, see the *IBM FileNet P8 Records Manager 3.5.x Installation and Upgrade Guide*.

Procedure Overview for Records Manager 3.5.0

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 3.5.x Installation and Upgrade Guide.* For high availability installation of these components, see "Content Engine" on page 26 and "Process Engine" on page 52.
2. Install a highly available Application Engine and note the location of the bootstrap.properties file (a highly available share is recommended). For information, see "Application Engine" on page 67.

   NOTE Do not deploy Workplace. This step will be completed in tandem when deploying the Records Manager component.
3. Install Records Manager on each node in the cluster where Application Engine is installed.
4. Configure the application server cluster for Records Manager.
5. Deploy Workplace and then Records Manager.
6. Perform any necessary post-deployment steps.
7. Start Process Routers on all nodes in the server farm.

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

CAUTION Redeploying Workplace after RM is installed.

When you install RM, the files **rmapi.jar** and **rmapiresources.jar** are copied to the **WEB-INF/lib** folder of the *deployed* Workplace application. If you ever redeploy Workplace, you must perform the following steps after it has been redeployed. For more information, see “Appendix B – Redeploying Workplace after Records Manager is Installed” on page 119.

The RM Web Application

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.

Records Manager 3.5.0 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 Application Engine software may be installed onto local storage for each machine in the cluster.

1. If the application server used to cluster Records Manager is WebLogic or WebSphere, verify that the cluster nodes and load balancers are started. If the application server is JBoss, just start the load balancer. JBoss deploys applications on startup so there is no need to start the application server until the installation process is complete.

   NOTE JBoss is designed to cluster deployed applications easily. Most application servers communicate state over multicast so it is critical that JBoss clusters be configured to segregate Application Engine and Records Manager multicast traffic from other JBoss clusters. See JBoss documentation for instructions regarding JBoss cluster configuration.

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

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.

Records Manager 3.5.0 Deployment Tasks

For more information about Records Manager installation, see Task 7a in the *IBM FileNet P8 Records Manager 3.5.0 Installation Guide.*

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

   Refer to Task 7a, steps 1-6 in the *Install RM (Typical)* procedure.

   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 Step 6 of the *Install RM (Typical)* procedure 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.

   Step 6 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).

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

2. Configure the application server cluster and deploy Application Engine according to the *IBM FileNet P8 Platform 3.5.0 Installation and Upgrade Guide*, then deploy Records Manager.

   - For WebSphere, refer to Task 9 of the *IBM WebSphere 5.x and 6.0* procedures in the IBM FileNet *Records Manager Installation and Upgrade Guide Release 3.5.0* 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 Task 9 of the *BEA WebLogic 7.x and 8.1* procedures in the *Records Manager Installation and Upgrade Guide Release 3.5.0* after the cluster and administrative domains have been created. Use these steps to deploy Records Manager on the cluster.
   - For JBoss, refer to Task 9 of the *JBoss* procedure in the *Records Manager Installation and Upgrade Guide Release 3.5.0*.

     NOTE Clustered JBoss instances are deployed from the ‘all’ directory and not the ‘default’ directory. See JBoss documentation for instructions regarding JBoss cluster configurations.

3. 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://<proxyurl>:<proxyport>/Workplace`

   Records Manager would have a URL of the form:

   `http://<proxyurl>:<proxyport>/RecordsManager`.

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

   NOTE Installation and configuration steps that utilize the Records Manager or Application Engine applications do not need to be duplicated for every node in the farm. Installation and configuration steps that utilize the Process Task Manager (to make configuration changes to Routers or the Component Manager) or the Records Manager Sweep applications must be performed on each node in the farm of application servers.

Records Manager 3.5.0 Deployment Verification

1. Make sure Application Engine 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.

Install the Records Manager 3.5.x service pack

This section covers the installation of the Records manager service pack in a highly available environment.

The Records Manager service pack installation process is straight-forward for both single instance and farmed/clustered Records Manager and is similar to a fresh installation in terms of configuration and deployment.

This topic identifies the tasks required to update your RM 3.5.0 installation to the latest RM 3.5.x Service Pack.

CAUTION Before you proceed, verify that you have successfully installed and configured your RM 3.5.0 software, including selected optional components, see "Records Manager 3.5.0" on page 81.

CAUTION Redeploying Workplace after RM is installed.

When you install RM, the files **rmapi.jar** and **rmapiresources.jar** are copied to the **WEB-INF/lib** folder of the *deployed* Workplace application. If you ever redeploy Workplace, you must perform the following steps after it has been redeployed. For more information, see "Appendix B – Redeploying Workplace after Records Manager is Installed" on page 119.

Records Manager 3.5.x Service Pack Installation Tasks

1. Identify any IBM FileNet P8 Platform Service Packs and fix packs required before you install the RM 3.5.x Service Pack. For details, contact you service representative, or access the IBM FileNet P8 3.5.x Patch Compatibility Matrix document on the IBM Information Management support www.ibm.com/software/data/support

1. Verify that Content Engine and Process Engine have been updated to the required service pack levels.

2. Verify that your highly available Application Engine has been updated to the required service pack and fix pack level.

3. Verify that your highly available Records Manager has been updated to the required fix pack level.

4. Install and deploy the RM 3.5.x Service Pack on each node in the cluster where Application Engine is installed.

   To install the Records Manager 3.5.x service pack, follow the instructions provided in the *Records Manager 3.5.x Service Pack Readme*, which covers a non-clustered installation procedure. For a highly available installation these step differ somewhat.

   Modify the readme steps as follows:

   a. Step 2: Copy the required configuration files.

      Where the instructions tell you to copy to the Content Engine server, copy to the server on which you will be running Enterprise Manager. In a clustered Content Engine environment, copy the files to the Active node.

   b. Step 4: Update the Records Manager application.

Do this step on all nodes.

In step A. Identify the complete path for the deployed location of the WEB-INF folder of the Workplace application.

The complete path for the deployed location of the WEB-INF folder of the Workplace application is different in an HA installation.

Examples of default WEB-INF deployment locations:

- WebSphere 5.x:

  ***<WAS_HOME>*/installedApps/*<node_name>*/app_engine_war.ear/app_engine.war/WEB-INF/**

- WebSphere 6.x:

  ***<WAS_HOME>*/profiles/*<profile_name>*/installedApps/*<cell_name>*/app_engine_war.ear/app_engine.war/WEB-INF/**

- WebLogic:

This path differs depending on how you deployed the Workplace application in your cluster.

– Deploy locally on each node (“from this location”)

  ***<AE_install_path>*/FileNet/Workplace/WEB-INF/**

– Deploy from one node (“Make copy to all nodes in cluster”)

  ***<WL_HOME>*/user_projects/domains/*<Domain_Name>*/*<Managed_Server>*/stage/Workplace/Workplace/WEB-INF**

- JBoss:

***<JBoss_HOME>*/server/all/deploy/Workplace.war/WEB-INF/**

In step C. Backup of the existing RMSweepConfiguration.xml file

Do this on all nodes from which you will be running the Records Manager Sweep Application.

In step J. Copy the backed up RMSweepConfiguration.xml file to its original installed location.

Do this on all nodes from which you will be running the Records Manager Sweep Application.

c. Step 5. (WebLogic and Tomcat only) Adjust the JVM memory settings.

   (WebLogic) Modify startManagedServer.cmd (Windows) or startManagedServer.sh (UNIX)

d. Step 8. Redeploy the Records Manager application.

   During redeployment 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).

Configure the application server cluster and deploy Application Engine according to the *IBM FileNet P8 Platform 3.5.0 Installation and Upgrade Guide*, then deploy Records Manager.

- For WebSphere, refer to Task 9 of the *IBM WebSphere 5.x and 6.0* procedures in the IBM FileNet *Records Manager Installation and Upgrade Guide Release 3.5.0* 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 Task 9 of the *BEA WebLogic 7.x and 8.1* procedures in the *Records Manager Installation and Upgrade Guide Release 3.5.0* after the cluster and administrative domains have been created. Use these steps to deploy Records Manager on the cluster.
- For JBoss, refer to Task 9 of the *JBoss* procedure in the *Records Manager Installation and Upgrade Guide Release 3.5.0*.

  NOTE Clustered JBoss instances are deployed from the ‘all’ directory and not the ‘default’ directory. See JBoss documentation regarding JBoss cluster configurations.

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://<proxyurl>:<proxyport>/Workplace`

Records Manager would have a URL of the form:

`http://<proxyurl>:<proxyport>/RecordsManager`.

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

NOTE Installation and configuration steps that utilize the Records Manager or Application Engine applications do not need to be duplicated for every node in the farm. Installation and configuration steps that utilize the Process Task Manager (to make configuration changes to Routers or the Component Manager) or the Records Manager Sweep applications must be performed on each node in the farm of application servers.

e. Deployment Verification

i. Make sure Application Engine and Records Manager applications are running on the cluster.

ii. 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`

iii. Verify that the Sign-in page displays.

iv. Fail one of the nodes.

v. Verify that you can reload the sign-in page.

f. Start Process Routers on all nodes in the server farm.

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

g. Step 9. Update the process region.

Do this only one time.

h. Step 10. Update and Configure Existing object stores.

In a clustered Content Engine environment, do this on the Active node.

i. Step 12. Verify the existence of the IS NULL operator in your FPOS(s) disposition schedules.

Step 12 is no longer required. The IS_NULL verification now takes place in a required RM 3.5.0 fix pack.

j. Step 13. Update and/or Reconfigure the Records Manager Sweep Application.

i. Do step a of this step on all nodes from which you will be running the Records Manager Sweep Application.

ii. Do steps b and c for each sweep configuration you want to set up.

iii. (Optional) Copy the sweep configuration files to other cluster nodes.

If you plan to configure two or more servers with identical sweep configurations you can copy the reconfigured RMSweepConfiguration.xml file and the newly created RMHoldSweepConfiguration.xml file to all cluster nodes from which you will be running the same sweeps.

Both of these files are located in:

***<RM_install_path>*/RecordsManagerSweep/lib/config**

CAUTION Having identical sweep configuration files on multiple servers will limit your sweeps to be run against either one FPOS only, or all FPOS connected to the CE.

Records Manager 3.7.0

IBM FileNet P8 Records Manager 3.7.0 can be installed 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.

The following procedure covers the differences required for installing Records Manager in an application server cluster. The remainder of the IBM FileNet P8 components required to support Records Manager are installed as per the IBM FileNet P8 High Availability Technical Notice.

NOTES

The following procedure will install IBM FileNet Records Manager 3.7.0 in a high availability environment.

The following procedure covers installation of the RM Web Application only.

For instructions on installing RM components such as the RM Sweep Application or Crystal Reports, see the *IBM FileNet P8 Records Manager 3.7.0 Installation and Upgrade Guide*.

Procedure Overview for Records Manager 3.7.0

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 3.5.x Installation and Upgrade Guide.* For high availability installation of these components, see "Content Engine" on page 26 and "Process Engine" on page 52.
2. Install a highly available Application Engine and note the location of the bootstrap.properties file (a highly available share is recommended). For information, see "Application Engine" on page 67.

   NOTE Do not deploy Workplace. This step will be completed in tandem when deploying the Records Manager component.
3. Install Records Manager on each node in the cluster where Application Engine is installed.
4. Configure the application server cluster for Records Manager.
5. Deploy Workplace and then Records Manager.
6. Perform any necessary post-deployment steps.
7. Start Process Routers on all nodes in the server farm.

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

CAUTION Redeploying Workplace after RM is installed.

When you install RM, the files **rmapi.jar** and **rmapiresources.jar** are copied to the **WEB-INF/lib** folder of the *deployed* Workplace application. If you ever redeploy Workplace, you must perform the following steps after it has been redeployed. For more information, see "Appendix B – Redeploying Workplace after Records Manager is Installed" on page 119.

The RM Web Application

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.

Records Manager 3.7.0 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 Application Engine software may be installed onto local storage for each machine in the cluster.

1. If the application server used to cluster Records Manager is WebLogic or WebSphere, verify that the cluster nodes and load balancers are started. If the application server is JBoss, just start the load balancer. JBoss deploys applications on startup so there is no need to start the application server until the installation process is complete.

   NOTE JBoss is designed to cluster deployed applications easily. Most application servers communicate state over multicast so it is critical that JBoss clusters be configured to segregate Application Engine and Records Manager multicast traffic from other JBoss clusters. See JBoss documentation for instructions regarding JBoss cluster configuration.

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

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.

Records Manager 3.7.0 Deployment Tasks

For more information about Records Manager installation, see the Install Records Manager task in the *IBM FileNet P8 Records Manager 3.5.0 Installation Guide.*

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

   Refer to steps1-10 in 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 step 10 of 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.

   Step 10 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).

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

2. Configure the application server cluster and deploy Application Engine according to the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, then deploy Records Manager.

   - For WebSphere, refer to the *Deploy the RM Web Application* task, *IBM WebSphere 5.x and 6.0* procedures in the IBM FileNet *Records Manager Installation and Upgrade Guide Release 3.7.0* 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 RM Web Application* task, *BEA WebLogic 7.x and 8.1* procedures in the *Records Manager Installation and Upgrade Guide Release 3.7.0* after the cluster and administrative domains have been created. Use these steps to deploy Records Manager on the cluster.

   - For JBoss, refer to the *Deploy the RM Web Application* task, *JBoss* procedure in the *Records Manager Installation and Upgrade Guide Release 3.7.0*.

     NOTE Clustered JBoss instances are deployed from the “all” directory and not the “default” directory. See JBoss documentation for instructions regarding JBoss cluster configuration.

3. 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://<proxyurl>:<proxyport>/Workplace`

   Records Manager would have a URL of the form:

   `http://<proxyurl>:<proxyport>/RecordsManager`.

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

   NOTES

   - Installation and configuration steps that utilize the Records Manager or Application Engine web applications need only be done on one node per HA environment, and not duplicated for every node in the farm.
   - The Transfer Workflow Definitions task is done from the command prompt using WorkflowTransfer.bat/.sh, and need only be done on one node per HA environment, and not duplicated for every node in the farm.
   - Installation and configuration steps that utilize the Process Task Manager (to make configuration changes to Routers or the Component Manager) or the Records Manager Sweep applications must be performed on each node in the farm of application servers.

Records Manager 3.7.0 Deployment Verification

1. Make sure Application Engine 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.

Rendition Engine 3.5.0

Rendition Engine high availability is achieved without farming or clustering of the Rendition Engine software. In fact, collocated Content Engine/Rendition Engine servers *cannot* be made highly available through deployment in a farm or cluster as:

- Active-passive clusters are not supported for Rendition Engine as the proprietary protocol used by the Rendition Engine software is not interoperable with clustering.
- Load-balanced server farms for the Content Engine Object Store Service are not supported when collocated with the Rendition Engine services.

Instead, to configure Rendition Engine for high availability, you configure a number of standalone Rendition Engine servers to each pull rendition work from a common publishing queue rather than having work pushed to them. The latter solution would require pushing via an IP address which in turn would have to be virtualized by farming or clustering of the servers.

This following topic is a high-level discussion of the configuration recommended by IBM to use multiple IBM FileNet Rendition Engine (RE) servers in a highly available configuration.

NOTES

- This topic does not cover detailed installation or configuration instructions of Rendition Engine in the configurations discussed below. For detailed installation-related questions, refer to the *IBM FileNet Rendition Engine Installation and Upgrade Guide*, see "Access IBM FileNet Documentation" on page 12.
- For general information about Rendition Engine functionality and administration, go to the *IBM FileNet P8 Platform* help and navigate to FileNet P8 Administration > Enterprise-wide Administration > Publishing.
- For general information on how to configure a highly available Content Engine, see "Content Engine" on page 26.
- For general information on how to configure a highly available database for Rendition Engine, see "Database Services" on page 22.
- The high availability configuration in Figure 7 on page 94 has an added performance benefit.

  In installations with heavy publishing workloads, IBM recommends having the Content Engine clients go through Content Engine servers that are not collocated with Rendition Engines. Otherwise, the publishing workload on the Rendition Engine component might degrade Content Engine performance and responsiveness to client requests.

A Highly Available Rendition Engine Solution

In the configuration discussed below, customer applications point to a Content Engine cluster/farm for Content Engine high availability. Separate from the cluster/farm, a number of Content Engine/Rendition Engine servers, each with its own database are set up solely for the purpose of processing publishing requests from the publishing queue.

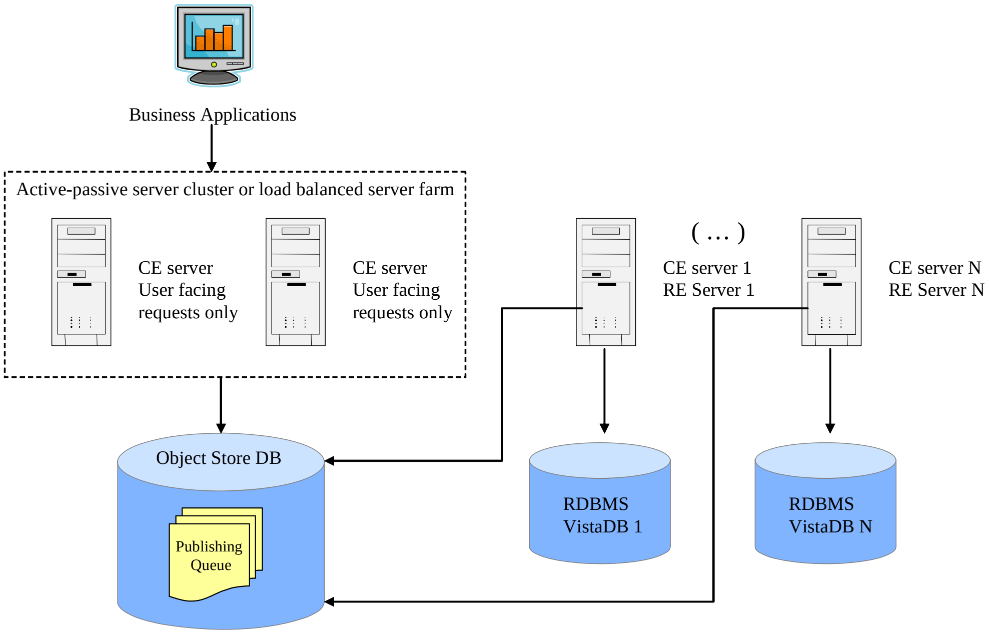

Figure 7: An example of a highly available Rendition Engine configuration.

To set up a highly available Content Engine cluster/farm with Rendition Engine support

1. Set up a Content Engine active-passive server cluster or load-balanced server farm with FileNet Publishing Plug-in Services disabled.

   These servers (Active passive server cluster or load balanced server farm in Figure 7) are necessary to provide a highly available Content Engine to serve Content Engine client applications and handle user-facing requests. They do not perform any rendition services; they do, however, submit publishing requests to the publishing queue.

For information on highly available Content Engines, see "Content Engine" on page 26.

NOTE Make sure you disable the FileNet Publishing Plug-in Services on the cluster/farm servers.

2. Set up two or more dedicated Content Engine/Rendition Engine servers with FileNet Publishing Plug-in Services enabled.

   These standalone servers (CEServer1 and CEServerN in Figure 7) each contain a Content Engine collocated with a fully configured Rendition Engine.

   For more information on installing and configuring Rendition Engine, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

   In this setup, each server contains the following components:

   - Content Engine
   - Publishing Plug-In services
   - Rendition Engine services
   - A unique Rendition Engine Domain containing:
     - Rendition Engine server
     - Rendition Engine database

   NOTES

   - Make sure you enable the FileNet Publishing Plug-in Services on the Content Engine/Rendition Engine servers.
   - Each Rendition Engine server must be paired with its own Content Engine. The Publishing Plug-In service running on the Rendition Engine/Content Engine server only sends one PDF and one HTML request at a time and waits for completion before sending another request of the same type.
   - To provide high availability, you must set up at least two Rendition Engine servers, but more than two can be deployed as the publishing workload demands.
   - If either of the Content Engine/Rendition Engine servers fail, the other server will continue processing requests from the publishing queue on the object store.
   - If a Content Engine/Rendition Engine server fails while performing a rendition request, the current job will eventually fail and an administrator must manually restart that request. For more detailed information, see "Completing publishing requests after a Rendition Engine server failure" on page 96.

Completing publishing requests after a Rendition Engine server failure

If one of the Rendition Engine servers fails, publishing requests that were in progress on the failed server will be set to a state of “In Error” and will not complete without manual intervention.

Job tickets sent to the Rendition Engine and the entries in the publishing queue will be in an “In Work” state. If these requests do not complete within a 45-minute period, they will go into an “In Error” state.

The “In Error” requests are stored in the system but will not be processed again unless an administrator manually sets them to be retried using the Publishing Queue interface of FileNet Enterprise Manager.

For more information, go to the *IBM FileNet P8 Platform* help and navigate to FileNet P8 Administration > Enterprise-wide Administration > Publishing > Publishing Queue Monitor > Accessing the Queue Monitor.

NOTE If a Rendition Engine server fails, a maximum of two requests will require this type of manual intervention. The Publishing Plug-In service running on Content Engine will only send two requests at a time to a Rendition Engine: one HTML and one PDF.

Team Collaboration Manager 3.5.0

CAUTION This topic covers installation of Team Collaboration Manager (TCM) 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 *P8-3.5.x / P8-3.6.x / P8-3.7.x Fix Pack Dependency/Compatibility Matrix* document.

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 12.

Team Collaboration Manager is composed of a number of components that can be configured for high availability. This section covers the tasks and requirements for making the following components highly available:

1. TCM Content Engine Integration
2. TCM Application
3. TCM Workflow Integration
4. Collaboration Engine
5. Collaboration Mail Server

CAUTION P8TCM-3.5.0-001 or later is required for Team Collaboration Manager to be used with a FileNet P8 3.5.x Platform installation, see "Team Collaboration Manager Fix Pack" on page 102.

NOTE The TCM components should be installed and configured for high availability in the order listed above.

TCM Web Application

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 TCM Web Application

1. Install the TCM web components on all nodes in the web farm.
2. Configure the application server cluster for the TCM web application.
3. Deploy the TCM web application to the cluster.
4. Perform any necessary post-deployment steps, and restart the TCM web application.

Pre-deployment tasks

The following procedure assumes that a cluster of server instances has been created and is running prior to deploying the FileNet Team Collaboration Manager. TCM does not require configuration data on network or clustered storage systems.

1. The TCMCryptoKeyFile.properties file gets created during the first installation of the TCM web components. Make sure to copy this file to the equivalent location on all other nodes in your cluster before proceeding with the TCM installation on subsequent nodes.
2. (If you enable User Tokens on your FileNet P8 system) Make sure all nodes in your TCM web farm are using the same UTCryptoKeyFile.properties file from the Application Engine. Copy this file from your Application Engine to the equivalent location on all your TCM cluster nodes.

   CAUTION For multiple applications to pass user tokens between one another, each participating application must use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all other Application Servers in the cluster. For more information, see Task 3: Configure Symmetric Encryption in the *FileNet Team Collaboration Manager 3.5.0 Installation Guide*.

Deployment Tasks

1. Install the TCM web components on all nodes. For more information about the TCM application install, see Task 5: in the *FileNet Team Collaboration Manager Installation Guide*.
2. Deploy the TCM Application, specifying the cluster as the deployment location instead of individual server instances. in the *FileNet Team Collaboration Manager Installation Guide*.
   - For JBoss/Tomcat, refer to Task 6a.
   - For WebLogic, refer to Task 6b.
   - For WebSphere, refer to Task 6c.

   NOTE You must install the TCM Application on all nodes in the cluster farm. The deployed web applications use functionality only available through a full TCM Application installation.

   Perform the following steps on each individual server.
3. Perform any post-deployment configuration steps such as configuring symmetric encryption or setting up SSL security.
4. Restart the TCM Application in the cluster after performing post-deployment configuration.

Deployment Verification

1. Make sure the TCM application is running on the cluster.
2. Open a web browser and type in the URL for the TCM application
   - If using a proxy server the URL has the form:

     ```
     http://proxy_server_name:port_number/TCM
     ```

- Otherwise, if connecting to an individual cluster node the URL would have the form:

  ```
  http://clustered_server_name:port_number/TCM
  ```

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

TCM Content Engine Integration

No special configuration is required to support the TCM Content Engine Integration components in a highly available environment. This component gets installed on a properly configured Content Engine cluster.

Procedure Overview for TCM Content Engine Integration

1. Install the component on the active node.
2. Failover the services to the next node in the cluster and install the TCM Content Engine Integration component on this, and each subsequent active node, in your cluster.

For more information, see Task 4 of the *FileNet Team Collaboration Manager 3.5.0 Installation Guide*.

TCM Workflow Integration

No special configuration is required to support TCM Workflow Integration components in a highly available environment. Install the TCM Workflow components on the servers where your environment's routers are configured to run (most likely the Application Engine).

For more information, see Task 7 of the *FileNet Team Collaboration Manager 3.5.0 Installation Guide*.

Collaboration Engine/Collaboration Mail Server

The Collaboration Engine and Collaboration Mail Server can be made highly available using active/passive clustering technology to monitor and control the services. The applications get installed as Generic Windows Services on Windows and standard daemons on UNIX. They can be started, stopped, and controlled using the clustering software, like most other Windows services and UNIX daemons.

This section covers the tasks and requirements needed to make the Collaboration Engine and Collaboration Mail Server highly available in a cluster environment. In such an environment, the Collaboration Engine Service and Collaboration Mail Server must run in an active-passive type of cluster configuration, with only one cluster node running these services at any one time.

NOTE The following procedure assumes you are running the Collaboration Engine and Collaboration Mail Server in the same cluster group of resources and specifies resource dependencies accordingly. If these services are not being co-located into the same cluster group then there will be no resource dependency between the Collaboration Engine and Collaboration Mail Server, nor will one be necessary.

Procedure overview for Collaboration Engine / Collaboration Mail Server

The following high-level steps are necessary to make the Collaboration Engine Service highly available:

1. Install Collaboration Engine Services on all nodes in the cluster.
2. Add the following resources to the cluster:
   - Collaboration Engine Service
   - Collaboration Mail Service
   - Clustered IP Address
   - Clustered Network Name

Collaboration Engine / Collaboration Mail Server - Microsoft Cluster Server

Microsoft Cluster Server provides a framework to create “Groups” of resources that collectively provide a highly available service. The following procedure describes how to add to a Group of resources in order to provide a highly available Collaboration Engine and Collaboration Mail Server Service.

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 Microsoft Hardware Compatibility list (HCL).
2. Install the clustering software prior to installing or configuring any FileNet services for high availability.
3. Verify that a cluster group already exists with the following resources (at a minimum):
   - Clustered IP resource
   - Network Name resource
4. Verify that the cluster group can failover to all nodes in the cluster.

Deployment Tasks

1. Install the Collaboration Engine and Collaboration Mail Server components on all nodes in the cluster. Refer to the FileNet Team Collaboration Manager 3.5.0 Installation Guide for each respective section.

   Task 8 - Install the Collaboration Engine

   Task 9 - Install the Collaboration Mail Server

   CAUTION You must specify a value for *Mail Domain* during the Collaboration Mail Server installation. Specify the cluster name from your cluster group and not the local hostnames of each node.

2. Add the following resources to the cluster:

   - Collaboration Mail Server Service
     - Resource type: Generic Service.
     - Set Service name parameter to “Collaboration Mail Server.”
     - Make the resource dependent on the Network Name resource.
   - Collaboration Engine Service
     - Resource type: Generic Service.
     - Set Service name parameter to “Collaboration Engine Service.”
     - Make the resource dependent on the Collaboration Mail Server resource (only if both resources are being collocated in the same cluster group).

       If the Collaboration Mail Server and Collaboration Engine are not co-located in the same cluster group then it will be dependent on the Network Name resource.

3. Bring up the services in the cluster group.

Deployment Verification

The following use cases can help ensure proper configuration of your cluster and verify consistent behavior amongst all nodes in your cluster before and after a failover.

1. Verify that the cluster resources you just created are running.
2. Start the rest of your TCM environment. See Task 12: Start Team Collaboration Manager of the *FileNet Team Collaboration Manager 3.5.0 Installation Guide*.
3. Login to your TCM web application and create a new teamspace, verify that a teamspace user receives an email notifying him of the new teamspace. If no email is received check your Collaboration Engine and Collaboration Mail Server logs.
4. From the new teamspace, select the send email function. On the send email dialog box, click the CC:Teamspace checkbox and click submit.

5. Check the email inbox from within your teamspace, you should be able to see the email that gets sent directly to the teamspace. If you do not, check your Collaboration Engine and Collaboration Mail Server logs.
6. After verification of email functionality is completed you should failover services to the passive node. Verify that emails get sent to users and to teamspaces from your other active node by performing Step 3 through Step 5 again. If problems exist, check your Collaboration Engine and Collaboration Mail Server logs to troubleshoot the problem.

Team Collaboration Manager Fix Pack

This section covers the installation of the Team Collaboration Manager fix pack in a highly available environment.

CAUTION

- P8TCM-3.5.0-001 or later is required for Team Collaboration Manager to be used with a FileNet P8 3.5.x Platform installation.
- If you install the FileNet TCM fix pack on a FileNet P8 3.5.0 Platform and subsequently upgrade to version 3.5.1 (or later), you must re-run the TCM fix pack installer to enable TCM on that version of FileNet P8. For more information, see the fix pack readme.

NOTES

- The fix pack must be installed on top of an existing TCM 3.5.0 installation.
- If you have installed TCM 3.5.0 components on different machines, you must install the fix pack on each machine.

Fix Pack Installation Tasks

On each node, follow the installation instructions in the fix pack readme to install the fix pack in your cluster. 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 12.

Upgrade Tasks

This section provides the information needed to upgrade an existing FileNet P8 3.0.0 highly available environment to IBM FileNet P8 3.5.x. It also provides instructions on how to upgrade a highly available RM 3.5.0 to RM 3.7.0

Additional Documentation

In addition to this technical notice you will need the following document to upgrade the IBM FileNet P8 platform for high availability:

- *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide* (PDF)

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 12.

High Availability Upgrade Order/Priority

The order of upgrade in a highly available environment is the same as detailed in the *Upgrade Core Components* section of the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*. This order, as applied to highly available environments, is as follows:

1. Review IBM FileNet P8 Platform requirements and other planning considerations. In the *Upgrade Overview* section, see "General Requirements for all IBM FileNet P8 Systems".
2. Upgrade the IBM FileNet P8 Platform documentation. In the *Upgrade Core Components* section, do Task 1: Upgrade IBM FileNet P8 Platform Documentation.
3. Ensure that you are running the minimum required release of IBM FileNet P8 Platform software and fix packs for doing an upgrade. In the *Upgrade Core Components* section, do Task 2: Verify Current Release Level.
4. If you're using an Oracle database engine, in the *Upgrade Core Components* section, do Task 3: Upgrade Oracle for Content Engine and Process Engine.

   If the Oracle database is remote, do Task 4: Upgrade Oracle Client for Content Engine and Process Engine as well. (If you're using a SQL Server database engine, no upgrade is required).
5. Upgrade the Content Engine Cluster/Farm. For more information, see "Upgrade a Content Engine Cluster/Farm" on page 105.
6. Upgrade the Process Engine Cluster. For more information, see "Upgrade a Process Engine Cluster" on page 107.
7. Upgrade the Application Engine Cluster/Farm. For more information, see "Upgrade an Application Engine Cluster/Farm" on page 110.
8. Install fix packs for core components. In the *Upgrade Core Components* section, do Task 9: Apply Service Packs and fix packs for core components.

NOTE For instructions on how to install IBM FileNet P8 Platform service packs in a high availability environment, see "Service Pack Installation Tasks" on page 73.

Step 5 to Step 7 are covered in the upgrade section of this technical notice, however, they provide only information specific to highly available environments, and rely on the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide* to detail all other information not related to high availability.

Upgrade a Content Engine Cluster/Farm

Object Store Service Farm Upgrade

The load balancing software or hardware does not need to be specially configured or modified to perform the upgrade. Upgrade the Object Store Service farm just like you would upgrade multiple Content Engine Servers.

The upgrade procedures are covered in detail in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, the *Upgrade Core Components* section, Task 5: Upgrade Content Engine. Use these procedures to upgrade the Content Engine Object Store services, as all procedures are applicable for a Farm setup.

Content Engine Cluster Upgrade

The upgrade process for a Content Engine cluster is the same regardless of the services being clustered, cluster configuration in use, and cluster software being used (for example there is no difference between the upgrade of a system using MSCS versus a system using VERITAS). The upgrade is performed on each node that has Content Engine software installed. The steps in this procedure will remove the Content Engine services from cluster control prior to upgrade, and add them back into the cluster after the upgrade is completed for all nodes.

Pre-Upgrade Tasks

1. Disable all cluster resources for Content Engine services (this includes any combination of the following if installed: Apache2, Object Store Service, Content Cache Service, File Store Service, and Process Services Manager).
2. Perform all other pre-upgrade tasks listed in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 5: Upgrade Content Engine.
3. Ensure that the GCD share, if clustered, is online when performing the upgrade.

Upgrade Tasks

1. Upgrade the Content Engine services on the active cluster node.

   Use the upgrade procedures in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, the *Upgrade Core Components* section, Task 5: Upgrade Content Engine.

   CAUTION The upgrade process may involve a reboot of the node. After the reboot, logon to another node in the cluster and fail the cluster group back to the node on which the Content Engine services are still being upgraded, then logon to the node that is being upgraded to continue the upgrade process.
2. After the upgrade is finished, fail the cluster group to another node.
3. Repeat Step 1 and Step 2 until all Content Engine servers are upgraded.

4. Perform the following steps from the active node of each cluster where File Store Service is installed.

   a. Modify the registry to specify the shared storage directory where transaction logs will be created.

      i. Navigate to:

         HKEY_LOCAL_MACHINE\SOFTWARE\FileNet\ECM

      ii. Create a new key named Content RM.

         Under the Content RM key, create a new string value named Log File Directory, whose value is the path to the directory for transaction logs.

   b. (VERITAS clusters) Create a registry replication resource and modify File Store service dependencies accordingly.

      i. Add a resource for the Registry Replication of the Content RM key

         - Set the Resource Type: RegRep
         - Set MountResName to the mount resource for the clustered disk.
         - Leave the ReplicationDirectory location on the clustered disk as default or set as desired.
         - Set Keys to:

           ```
           HKLM\\SOFTWARE\\FileNet\\ECM\\Content RM
           ```

         - Make the resource dependent on the following resource:
           - Clustered disk resource

      ii. Modify the dependency for the Content Engine File Store Service cluster resource so that it is dependent on the Content RM registry replication resource.

   c. (MSCS clusters) Add registry replication for the Content RM key:

      For the Content Engine File Store Service cluster resource add the following value to the Registry Replication tab:

      ```
      software\FileNet\ECM\Content RM
      ```

5. Re-enable the cluster resources and bring the cluster online.

Upgrade a Process Engine Cluster

Procedure overview for Process Engine Upgrade

Upgrading the software is performed on all nodes in the cluster, while upgrading the Process Engine database objects is performed only once from the node running Process Engine services.

The following high-level steps are necessary to upgrade the Process Engine for high availability:

1. Stop all Process Routers in the IBM FileNet P8 environment.
2. Stop all Process Engine services in the cluster.
3. Disable (on Windows clusters) or remove (on UNIX VERITAS clusters) any cluster resources monitoring and controlling the Process Engine.
4. Upgrade the Process Engine software on each node.
5. Recreate cluster resources to monitor and control the Process Engine.
6. Start Process Engine services.

The following procedures are based on these steps.

UNIX Clusters

Pre-Upgrade Tasks

1. Stop all Process Routers in the IBM FileNet P8 environment.
2. Stop the 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 3.5 Installation and Upgrade Guide*, the *Upgrade Core Components* section, Task 6: Upgrade Process Engine (Unix).

Upgrade Tasks

1. Make sure the cluster group is running on the node being upgraded.
2. Start the upgrade; follow the steps in Task 6 of the Upgrade Core Components section in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*.
3. Once the upgrade is finished on the node, fail 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. Install the FileNet P8 3.5.0 Process Engine Agent.
   a. Untar the FileNet Process Engine Agent package located under the VERITAS folder on the IBM FileNet P8 Process Engine software package.

b. Copy the files to **all nodes** in the cluster.

c. **On all nodes** perform the following steps:

i. Create the directory **/opt/VRTSvcs/bin/FN_ProcessEngine**.

ii. Copy the FN_ProcessEngineAgent agent binary to the following directory (overwriting the existing file):

/opt/VRTSvcs/bin/FN_ProcessEngine

iii. Change the permissions on the directory as follows:

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

iv. Copy the file FN_Types.cf 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 new resource of the type FN_ProcessEngine and set the attributes as described in Step g and Step h on page 56.

5. Bring the cluster online.

Windows Clusters

Pre-Upgrade Tasks

1. Stop all Process Routers in the IBM FileNet P8 environment.
2. Stop the Process Engine service and Process Engine registry replication resources.
3. Disable 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 3.5 Installation and Upgrade Guide*, the *Upgrade Core Components* section, Task 7: Upgrade Process Engine (Windows).

Upgrade Tasks

1. Make sure the cluster group is running on the node being upgraded.

2. Start the upgrade; follow the steps in Task 7 of the Upgrade Core Components section in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*.
3. Once the upgrade is finished on the node, fail 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. After upgrade is complete 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 59.
   - For Microsoft Cluster Server, see "Modify the registry to set the NCHBroadcast value." on page 63.
5. Modify the cluster resource controlling the Process Services Manager. This service has been renamed "Process Engine Services Manager."
   a. Set the "Service name" parameter for (MSCS) or "ServiceName" parameter for VERITAS to "VWServicesPE."
   b. (MSCS only) Modify the cluster resource controlling the IMSService so that it is dependent on the Network Name resource of the cluster group.
6. Re-enable the cluster resources and bring the cluster online.

Upgrade an Application Engine Cluster/Farm

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 IBM 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.
7. Start all Process Routers.

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

Pre-Upgrade Tasks

Perform the Steps under “Before you upgrade Application Engine” in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*, Task 8: Upgrade Application Engine.

Upgrade Tasks

1. Stop all Process Routers and IBM FileNet Java processes in the IBM FileNet P8 environment as documented in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 8: Upgrade Application Engine, the “To upgrade Application Engine” procedure steps 2 and 3.
2. Undeploy Application Engine as described in Step 4 of the “Upgrade Application Engine” section in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*.
3. Perform Application Engine upgrade tasks from Step 5 through 9 of the “Upgrade Application Engine” section in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide*. Perform these steps **on all server instances/nodes** in the farm/cluster.
   - When specifying the machine name of the Content Engine make sure to specify the clustered/farmed name if the Content Engine has been configured for high availability.
   - For the bootstrap.properties file location screen, specify the location created to hold Application Engine configuration data.

4. Configure and deploy the application server cluster for Application Engine.

   Follow the instructions in the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 15*x*: Configure and start Application Engine (*Application server*), for each individual node server, with the following additional information.

   - When deploying Workplace, specify the cluster as the deployment location instead of individual server instances.

5. Perform any post-deployment configuration steps such as configuring symmetric encryption or setting up SSL security.

6. (If you enable User Tokens on your IBM FileNet P8 system) Make sure each Application Engine server has the same UTCryptoKeyFile.properties file.

   CAUTION For multiple applications to pass user tokens between one another, each participating application must use the same encryption key file. Copy the UTCryptoKeyFile.properties file installed with Application Engine to all other Application Servers in the cluster.

   For general information about user tokens, go to the *Documentation for IBM FileNet P8 Platform* help and navigate to Developer Help > Workplace Integration and Customization > User Tokens > Configuring Applications to Use Tokens.

7. Restart the Application Engine application in the cluster after performing post-deployment configuration.

8. Configure and start routers on all machines in the cluster, see the *IBM FileNet P8 Platform 3.5 Installation and Upgrade Guide,* Task 18: Start Process Routers.

   In order to provide the same Process Engine data to each clustered Application Engine instance, routers will need to be started on all machines in the cluster. Use the same router configurations (name, region, and Process Engine Server connection information) across all Application Engine cluster machines.

Upgrade a Records Manager Cluster

The Records Manager 3.5.x to 3.7.0 upgrade process is straight forward for both single instance and clustered Records Managers and is similar to a fresh installation in terms of configuration and deployment.

NOTES

The following procedure covers upgrades of the RM Web Application only.

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 3.7.0 Installation and Upgrade Guide*.

This document covers upgrades from RM 3.5.x to RM 3.7.0, and does not include upgrade instructions for RM 3.0.0 to RM 3.5.x.

Procedure overview for Records Manager Upgrade

An upgrade of a Records Manager cluster/server farm includes the following steps:

1. If necessary, import fileplans for upgrade and verify custom changes to RM workflows.
2. Upgrade Records Manager Software.
3. If you will be using DoD Chapter 4 support, create and assign users to the Classified Records Security Personnel Access Role.
4. Configure report viewing if required.
5. Synchronize JAR files.
6. Redeploy the RM Web Application.
7. Start a Process Router.
8. Update the Process Region.
9. Upgrade RM Data, Data Model, and Security.
10. Restart Component Manager on the running Process Router.
11. Start Process Routers and Component Managers on the remaining nodes.

CAUTION Redeploying Workplace after RM is installed.

When you install RM, the files **rmapi.jar** and **rmapiresources.jar** are copied to the **WEB-INF/lib** folder of the *deployed* Workplace application. If you ever redeploy Workplace, you must perform the following steps after it has been redeployed. For more information, see “Appendix B – Redeploying Workplace after Records Manager is Installed” on page 119.

Pre-Upgrade Tasks

Perform the Steps under “To prepare to upgrade RM” in the *IBM FileNet Records Manager Installation and Upgrade Guide Software Release 3.7.0* task “Prepare to Upgrade Records Manager.”

Upgrade Tasks

To upgrade the RM Web application from version 3.5.x to 3.5.7, follow the tasks in the “Upgrade Tasks” section of the *IBM FileNet Records Manager 3.7.0 Installation and Upgrade Guide,* with the following modifications:

1. If you will be using DoD Chapter 4 support, upgrade the P8 domain's Marking Sets.

   Follow the instructions in the ”(DoD Chapter 4 only) Upgrade Marking sets on the FileNet P8 domain” task.

2. Upgrade Records Manager Software.

   To upgrade the Records Manager software, on each node, follow the instructions in the “To Upgrade RM” procedure 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 RM 3.5.x. See “Records Manager 3.5.0 Pre-deployment Tasks” on page 82.

   - “Undeploy the RM Web Application.”

   - Follow the instructions in the “Upgrade Records Manager Software” task, “To Upgrade RM” procedure, step 4.“Complete the Setup screens as follows:”

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

3. If you will be using DoD Chapter 4 support, create and assign users to the Classified Records Security Personnel access role.

   Follow the instructions in the ”(DoD Chapter 4 only) Create the Classified Records Security Personnel Access Role” and the “(DoD Chapter 4 only) Assign Users to the Classified Records Security Personnel Access Role” tasks.

4. 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.

5. Synchronize Jar files.

   Follow the instructions in the task “Synchronize Jar Files.”

   NOTES

- Perform these steps on all server instances/nodes in the farm/cluster.
- If you have installed the RM Sweep application on your nodes, follow the instructions in the “Upgrade Records Manager Software” task, “To Upgrade RM” procedure, step 13.

6. Redeploy the RM Web Application.

   Follow the instructions in the “Redeploy the RM Web Application” task.

   - For WebSphere, refer to the “IBM WebSphere 5.x and 6.0” procedures Using these steps re-deploy Records Manager on the cluster and proxy server instances.
   - For WebLogic, refer to the “BEA WebLogic 7.x and 8.1” procedures. Use these steps to deploy Records Manager on the cluster.
   - For JBoss, refer to the “JBoss” procedure.

     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.

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

7. Start a Process Router and update the Process Region.

   Follow the instructions in the “Update the Process Region” task.

8. Upgrade RM Data, Data Model, and Security on your object stores.

   Follow the instructions in the following tasks:

   - “Configure RM Data for Upgrade.”
   - “Upgrade RM Data Model.”
   - “Upgrade RM Security.”

9. Restart Component Manager on the running Process Router.

   Follow the instructions in the “Restart the Component Manager” task.

10. Start Process Routers and Component Managers on the remaining server instances/nodes in the farm/cluster.

    Follow the instructions in the “Update the Process Region” task.

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 IBM FileNet users and groups, and configure the node 1 and node 2 cluster servers.

Configure the Domain Controller

Create IBM FileNet Groups

Create the following domain local security groups on the domain controller:

Table 1: IBM FileNet Groups

Group Name	Group Description
FNADMIN	Members have all privileges on IBM FileNet files and databases
FNOP	Members can start/stop and execute IBM FileNet software
FNUSR	Members have normal privileges on IBM FileNet files and databases

Create IBM FileNet Users

On the Domain Controller, create the following users:

Table 2: IBM FileNet Users

User Name	**User Description**	**Member of**
<Installer>[a]	User that will install the IBM FileNet P8 Process Engine software.	FNADMIN, FNOP, and FNUSR
fnsw	Primary IBM 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.

b.The oracle username 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.

Add Nodes to Pre-Windows 2000 Compatible Access Properties

1. On the Active Directory Users and Computers window, click the **Builtin** folder.
2. Right-click the Pre-Windows 2000 Compatible Access object and select **Properties**, as shown above.
3. Click on the **Members** tab, and click the **Add** button.
4. Select **Node 1** and click the **Add** button.
5. Select **Node 2** and click the **Add** button.
6. 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.
7. Click **Apply** on the Pre-Windows 2000 Compatible Access Properties dialog box to apply the changes.
8. Click **OK** to close the Pre-Windows 2000 Compatible Access Properties window.
9. 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.

Create the LocalAdminInstall File

Use this procedure to create the `LocalAdminInstall` file in the **c:\temp** directory.

1. Turn-on the node 1 and node 2 servers and logon as Domain Administrator on each server.
2. Open a Command Prompt window.

3. 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.

4. At the **\temp** directory, type the following command and press Enter:

   ```
   copy con LocalAdminInstall
   ```

5. Press and hold **Ctrl** key, and press the **Z** key.
6. Press **Enter**.
7. Verify that the `LocalAdminInstall` file was successfully created in the **c:\temp** directory.

Create Local IBM FileNet Groups

On the node servers, create the following local groups::

Table 3: Local IBM 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

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.

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 2000

- Act as part of the operating System
- Log on as a service
- Increase quotas
- Replace a process token

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 116 to repeat these procedures on the node 2 server.

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 – Redeploying Workplace after Records Manager is Installed

CAUTION When you install RM, the files **rmapi.jar** and **rmapiresources.jar** are copied to the **WEB-INF/lib** folder of the *deployed* Workplace application. If you ever redeploy Workplace, you must perform the following steps after it has been redeployed.

1. Stop the web application server.
2. Copy the files **rmapi.jar** and **rmapiresources.jar** from ***<RM_install_path>*/FileNet/RecordsManager/WEB-INF/lib** to the **WEB-INF/lib** folder of the *deployed* Workplace application.
3. Start the web application server.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user’s responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM’s future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

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

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

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

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, 5724-S00, , 5724-S19

Printed in USA

GC31-5491-01