IBM FileNet P8 Platform

Version 4.0

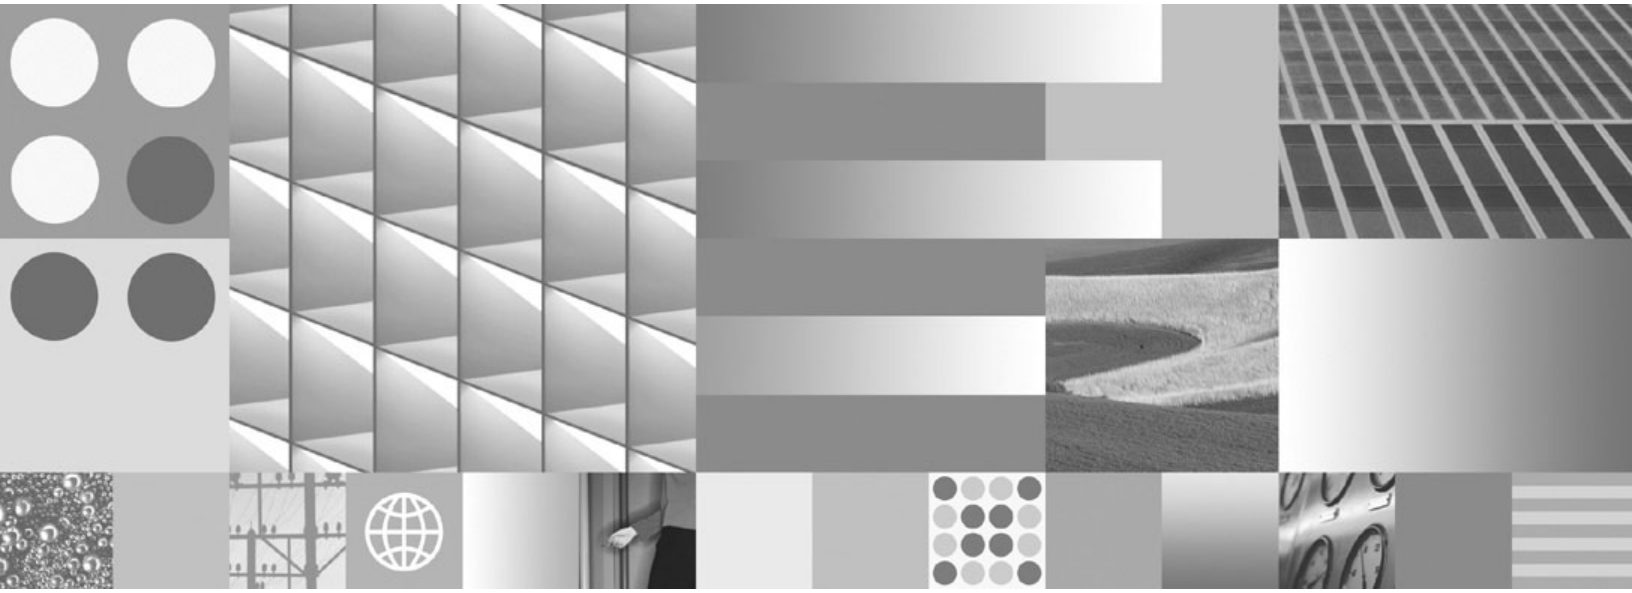

Disaster Recovery Technical Notice

GC31-5486-00

IBM FileNet P8 Platform

Version 4.0

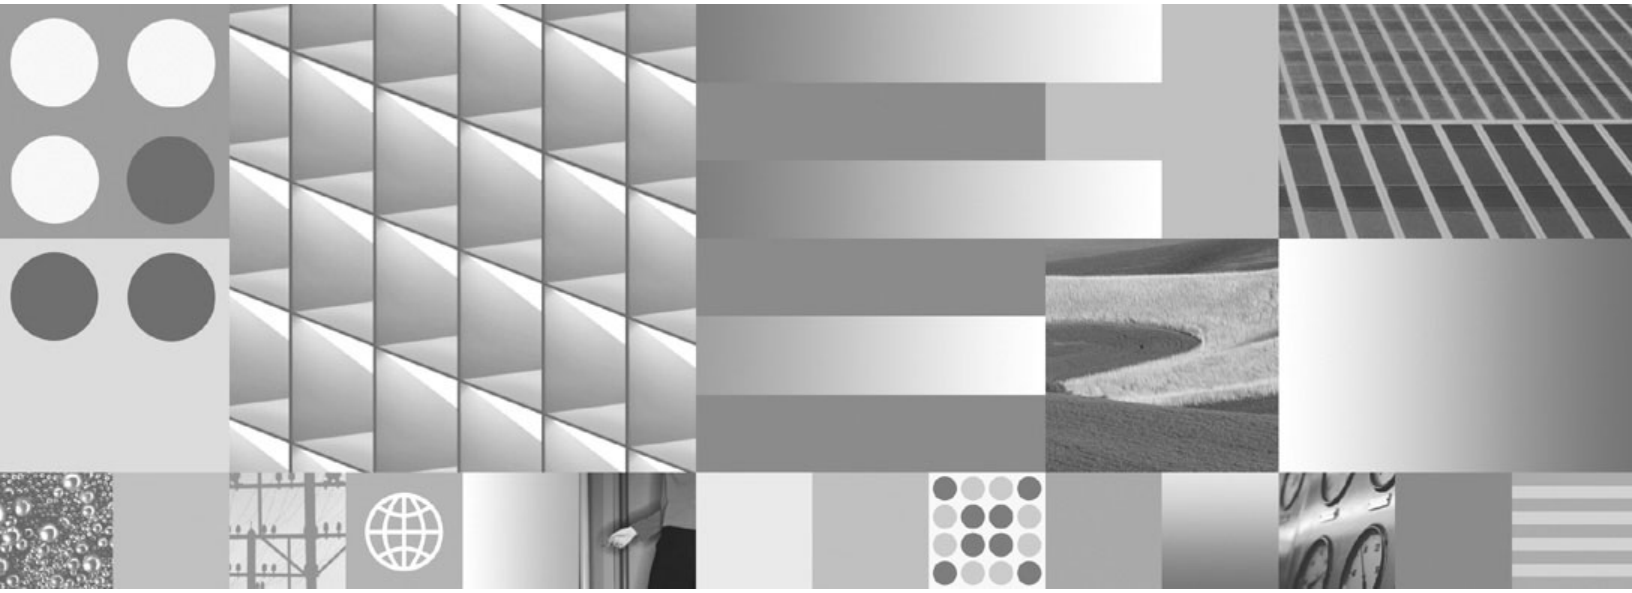

Disaster Recovery Technical Notice

GC31-5486-00

Note
Before using this information and the product it supports, read the information in "Notices" on page 43.

This edition applies to version 4.0.0 of IBM FileNet Business Process Manager (product number 5724 R76), version 4.0.0 of IBM FileNet Content Manager (product number 5724 R81), and to all subsequent releases and modifications until otherwise indicated in new editions.

Revision Log

Date	Revision
08/08	Added standard IBM cover pages and Trademarks section.
12/07	Made various changes throughout document to meet IBM requirements.
04/07	Initial document.

Contents

Table of Figures

FileNet P8 Platform Disaster Recovery

This document describes the technical issues related to configuring the FileNet® P8 Platform for disaster recovery.

CAUTION The information in this document is intended as overview material. For more detailed instructions, see your third-party and FileNet documentation. The document is *not* intended as a step-by-step guide detailing the use of the many third-party software and tools required to create and configure a disaster recovery solution.

Audience

The intended audience for this technical notice is the experienced system, database, network, and other technical administrator who will play a part in designing and constructing a disaster recovery solution.

The procedures in this document require a working administrative knowledge of the FileNet P8 system.

Accessing IBM FileNet Documentation

To access documentation for IBM FileNet products:

1. Navigate to the Information Management support page (www.ibm.com/software/data/support).
2. Select the appropriate IBM FileNet product from the "Select a category" list.
3. From the Product Support page, click **Product Documentation** under Learn.
4. From the Product Documentation page
   a. If necessary, click the Doc Link for the appropriate component product to display the document list.
   b. Click the icon in the appropriate release column to access the document you need.

Additional Documentation

Consult the appropriate documentation, including user manuals, administration guides, and solution guides, before implementing a disaster recovery solution.

FileNet Documentation

- *IBM FileNet P8 Platform 4.0.x Installation and Upgrade Guide* (PDF)
- *IBM FileNet P8 Platform 4.0.x High Availability Technical Notice* (PDF)
- *IBM FileNet P8 Platform 4.0.x Hardware and Software Requirements* (PDF)

To download IBM FileNet documentation from the IBM support page, see Accessing IBM FileNet Documentation.

Symantec Veritas Documentation (http://www.symantec.com/enterprise/support/index.jsp)

- NetBackup Enterprise Server Version 6.0
- Veritas Storage Foundation Enterprise HA Version 5.0
- Veritas Volume Manager Version 5.0
- Veritas File System Version 5.0
- Veritas Volume Replicator Version 5.0

- Veritas Storage Foundation Enterprise HA Version 5.0
- Veritas Cluster Server Version 5.0

Oracle Documentation (http://www.oracle.com)

Oracle 10g RDBMS

Microsoft® Documentation (http://www.microsoft.com)

- Windows® 2003 DNS Server
- Windows 2003 Active Directory

Terms and Concepts

What is Disaster Recovery?

Business continuity is the practice of leveraging procedure and infrastructure to enable business operations to continue in the event of any kind of failure, up to and including a full disaster. Disaster recovery is the subset of business continuation dealing with disaster, while high availability is a subset of business continuity dealing with local component failures. Disasters can be natural occurrences such as floods or earthquakes, or products of man, such as armed conflicts, riots, or construction accidents.

While unplanned downtime is always costly in terms of lost productivity and unfulfilled business transactions, extended interruptions severely impact the bottom line and can cripple a business' profitability.

NOTE From a technical standpoint, a disaster is any extraordinary occurrence that disrupts day to day operational activity by disabling or destroying the production system, and is beyond the control of the business.

Disaster Recovery and High Availability

FileNet P8 Platform ECM/BPM solutions are engineered to impart the utmost reliability. However, factors external to FileNet P8 can threaten the accessibility of critical business data. To ensure availability, IBM supports high availability solutions for its FileNet P8 products.

Disaster recovery builds upon the principles of high availability, and improves the availability of clusters by duplicating data to another cluster in an alternate data center. Although disaster recovery is a solution to a specific problem, how to eliminate a data center or campus as a single point of failure is not a stand-alone solution for ensuring availability.

High availability solutions provide business continuity in the face of localized failures, such as a single server failure or hard disk crash. Disaster recovery solutions, on the other hand, provide for business continuity in the face of natural or man-made disasters that cause the loss of an entire production system. While the goal of both high availability and disaster recovery solutions is the same—keeping your 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 might 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.

For more information on High Availability implementations of FileNet P8, see the *IBM FileNet P8 High Availability Technical Notice*.

Implementing Disaster Recovery Environments

Disaster recovery environments should ideally be implemented in a staged approach.

1. Identifying Requirements

   The first stage is to identify availability requirements and weigh them against the cost of implementation. This process should specifically identify the business needs for a disaster recovery environment. If the need is in sharp imbalance with the costs involved, a disaster recovery project could be inappropriate for the business.

   NOTE For a detailed discussion of this topic, see "Appendix Two: Identifying Requirements" on page 33.

2. Project Planning

   The second stage is to create a plan to fulfill the requirements identified in the first stage. The disaster recovery project must have a clearly defined scope and must meet business objectives. Clearly defined goals and objectives enable project managers and architects to gauge progress and measure success.

   NOTE For a detailed discussion of this topic, see “Appendix Three: Project Planning” on page 38.

3. Implementation

   The steps for setting up FileNet P8 Platform for disaster recovery are described in “Case Study: Veritas Volume Manager 5.0, Veritas Cluster Server 5.0” on page 18.

   NOTE As requirements vary from business to business, and depending on the third party products used, the solution from this case study should not be seen as step-by-step instructions, but rather as a reference for architecting a FileNet P8 deployment for disaster recovery.

4. Verification and Documentation

   Once the environment is implemented, the fourth and final stage is to measure the success of the project against business objectives. The final verification of components and acceptance testing determines if the system is able to perform under real-world conditions. Each cluster, global service group, and global application must be administratively verified and released to the user community for acceptance testing. Availability and performance issues raised during testing should be investigated and resolved before the environment is released to the general user community for production use.

   NOTE For a detailed discussion of this topic, see “Appendix Four: Verification and Documentation” on page 40.

Common Disaster Scenarios

The goal of a successful disaster recovery initiative is to eliminate a corporation’s exposure to risk, regardless of circumstance. While natural disasters can be a contributing factor, human error presents a far greater risk to data center availability. The table below contains examples of common disasters.

Natural Disasters	Man-made Disasters
• Floods	• Construction accidents
• Earthquakes	• War
• Wildfires	• Fires
• Hurricanes	• Engineering failures

Components of a Disaster Recovery Solution

A disaster recovery environment ensures business continuity by duplicating services in an alternate geographic location using data replication and global clusters. If a primary site becomes unavailable to application users, an alternate site takes over and resumes operations. This procedure is known as a global or site failover. Generally, a site failover would not be initiated unless a disaster interrupts application availability for a prolonged or indeterminate period of time.

A global application framework leverages fault tolerant resources across multiple clusters in multiple locations using resources such as:

- Individual servers with redundant components which are routinely backed up in case of system failures.
- Fault tolerant storage.
- Two or more servers grouped together in a highly available cluster that mitigates the risk of a complete server or operating system failure.
- Sites composed of logical groups of clusters at a common location.
- Primary sites replicating dynamic application data to disaster recovery sites.
- Global Applications consisting of logically grouped cluster resources which can be migrated between primary and disaster recovery sites.

The following is a high-level list of the various components that are leveraged to create a disaster recovery environment.

Primary Production Site

Disaster recovery needs a highly available solution operating at the primary production site to ensure that simple server failures do not trigger an unnecessary site failover. For more information on setting up a highly available production environment, see the *IBM FileNet P8 High Availability Technical Notice*.

Backup and Restore: Laying the Foundation

Regularly back up the production environment, ensuring that all of the data across all applications are synchronized. Store cloned copies of backup media offsite. Backups allow data to be restored when files are lost due to human error, file system corruption, or disk failure.

To guarantee consistency, users must not modify the files in the file system while the backup is running. Maintaining the consistency of application components is complicated by the distributed nature of the FileNet P8 environment.

The case study includes an example of how to set up low-impact backups using Veritas NetBackup. Backups of production data should be tested at least once before a go-live date (when the server is released to the user community for production use) and once or twice a year thereafter. This practice ensures that day to day administration of servers and applications does not jeopardize the operational integrity of the backup infrastructure.

For more information, see the *IBM FileNet P8 online documentation*, Backup and Restore topic.

WARNING Do not store documentation critical to the recovery of FileNet P8 disaster recovery components exclusively in FileNet P8. Should a catastrophe strike and data must be recovered from tape, details must be readily available for functional groups to begin restoring the components to the environment.

Backup

The foundation of a disaster recovery environment rests on the automated backup of application data. Backing up data in distributed applications, where disparate components must be in sync, presents a significant technical challenge.

For example, files in a Content Engine file store must match the transactional content of the database. In order to obtain a reliable backup, all components, Content Engine, and the database must be backed up at the same time while transactions are suspended. If these components run on separate servers, the backup procedure must coordinate the shutdown, backup, and restart of these components.

Snapshot technology, described below, can be used to minimize the impact of backups and accommodate a narrow backup window.

Snapshots

A snapshot is a logical mechanism for creating a copy of data at an instantaneous point-in-time. A snapshot can then be mounted and backed up while the production data is once again available for applications. Backups can then stretch beyond the backup window without impacting production hours.

In the case of FileNet P8, all applications (including support software like Oracle) must be shut down, the snapshot created, and then the applications restarted allowing the users to access the production data. You then back up the snapshot instead of the actual production environment.

Many snapshot solutions exist, and although the snapshot method can vary, the underlying principles for implementing a backup procedure remain the same.

Volume Management: Fault Tolerant Storage

Data critical enough for backup should also be stored on fault tolerant storage. The primary function of fault tolerant storage is to ensure that disk failures do not preclude users from accessing data and applications. Fault tolerance is especially important in the enterprise where large numbers of disks are used to meet performance and capacity demands of the user community.

As with most high availability concepts, storage must be configured in a manner which allows multiple failures to occur and repairs to be conducted without disrupting running applications. The case study in this technical notice uses Veritas Volume Manager to maximize the availability of the storage system.

Conduct your storage failure tests when you initially deploy the infrastructure, and retest when significant changes in volume layout occur.

Hardware Redundancy: Hardening Server and Data Center Infrastructure

Identify and eliminate single points of failure in the data center. Hardened infrastructure should include fault tolerant heating, ventilating, and air conditioning (HVAC), uninterruptible power supplies (UPS), power distribution units (PDU), electrical circuits, and network access. Hardware redundancy at a server level prevents component failures from interrupting the availability of its applications. Fault tolerant servers use hot swappable, redundant host bus adapters (HBA), network cards, and power supplies.. System memory slots should be populated with registered error-correcting code (ECC) modules only.

Before a server is truly production-ready, test all component failover scenarios, note and address any deficiencies. Conduct these tests at the time the server is built or when significant infrastructure changes mandate recertification.

Clusters: Ensuring Application Availability

Regardless of the degree of redundancy asserted in a single server, chassis failures remain a risk to the availability of a server's applications. Clusters mitigate this risk by pairing two or more servers together. When a server failure occurs, applications fail over to another server node which assumes the

responsibility of the failed node. Users are largely unaware of the transition between server nodes. This concept is known as application transparency.

Truly transparent application clusters involve little observable interruption to the user during failover. High availability for FileNet P8 Platform is covered in the *IBM FileNet P8 High Availability Technical Notice*. Prior to implementing data replication, clusters should be tested thoroughly for proper failover and recovery.

Disaster Recovery Site

Data Replication: Duplicating Data to Alternate Sites

The core component of disaster recovery is data replication. The risk to proprietary enterprise data is mitigated by replicating the data to an alternate location in real time. As users commit data to FileNet P8 it is also written to a storage system at an alternate site. In the event of a disaster the duplicated data allows application users to resume work using the duplicate copy, ideally with little downtime. Implemented properly, data duplication in this manner is suitable for use with FileNet P8 applications.

The case study included in this technical notice uses Veritas Volume Replicator to ensure that FileNet P8 data is available on a global scale. Before a disaster recovery environment is taken into production you should verify the replicated data by halting replication at the primary site and testing the data with the application at the disaster recovery site.

Global Applications: Automating Business Continuity

Globally managed clusters sharing duplicate data in multiple sites can migrate business critical operations to an alternate location, in a matter of minutes. Centralized management of clusters greatly simplifies the orchestration of a site-wide failover. This capability is especially powerful when used with distributed applications such as the FileNet P8 platform. Users of globally managed applications are largely unaware that the FileNet P8 components and infrastructure have moved to a backup site at an alternate location.

Rehearse and conduct disaster recovery drills and site failover tests before going into production. You should also retest these annually.

Both Sites

Infrastructure Uniformity

You must have software uniformity throughout a disaster recovery environment. Any difference between one node and the next corresponding node diminishes the likelihood that a failover will execute smoothly when disaster strikes. Even minor differences are not only difficult to troubleshoot but also complicate patching and render regular maintenance unnecessarily difficult.

Fix pack levels, service packs, and configuration files must match exactly for all applications on each node. The same principle also applies to applications and their configuration files.

Uniformity is also part of best practices at the hardware level, but cost reasons could prohibit identical solutions.

Example:

Your production site might contain ten web servers to support a large user community for day-to-day activities. The disaster recovery site where servers are simply on standby mode might contain fewer physical servers. The servers in the disaster recovery site remain in standby until a disaster event warrants its activation.

When the disaster recovery site is brought online it will typically stay online for a short duration or until the primary production site comes back online. Hence, it might be difficult to justify the same hardware investment and maintenance costs for the disaster recovery site as for the primary production site. If the disaster recovery site has lower capacity than the primary site, you must account for this after a failover by for example limiting the use of the system to the most critical functions only.

CAUTION If you make modifications to software at the production site, you must match the modifications at the disaster recovery site as well.

Time Zone Uniformity

Transactions and batch jobs are often time dependent and thus each system should be time synchronized. If the disaster recover site is in a time zone other than the one the user community normally operates in, be sure users understand the time offset.

Example:

If a batch job is scheduled to occur at 2:00 AM at the primary site, to minimize the impact on users, and the disaster recovery site is in a different time zone, you must account for this when scheduling when the batch job will run at the disaster recovery site.

Scaling: Identifying an Appropriate DR Strategy

Global application management provides a powerful and flexible framework for failing an application from one site to another. However, flexibility can create challenges with physically separate storage and the coordination of application failovers.

Application Failover Transparency

In FileNet P8, Content Engine, Process Engine, Application Engine, and eForms are examples of horizontally scaled applications. These applications use a load balancing mechanism or a fault tolerant cluster address to maintain the availability of the services they provide. Users and applications access services by specifying the hostname of a load balancer or proxy, which forwards requests to one of a group of servers running the application.

FileNet P8 users access Workplace through Application Engine, which requires users to be transparently redirected to an alternate site in the event of a disaster. This function is handled using dynamic domain name system (DNS) updates. In the case of Application Engine, the interface point is the URL used to

connect to Workplace. The hostname portion of the URL uses the DNS alias of the global application instead of the cluster hostname.

Content Engine and Application Engine are farmed by using an application server's internal clustering mechanism. This precludes the ability to manage resources not deployed in the application server itself. From a global application standpoint the global failover of Content Engine and Application Engine is facilitated using a dynamic DNS update. When an event triggers a global failover, the application server is not manipulated. Instead, the alias or canonical name (CNAME) of the load balancer is switched to direct users to the alternate site.

The concept of DNS aliases also applies to the Process Engine region ID configured in FileNet Enterprise Manager (EM) or the Content Engine web services URL. However, in regards to the Application Engine URI in the WcmApiConfig.properties file, the CEMP application server cluster address will be used. This address is a comma-separated value using the application server supported protocol such as WebLogic T3 and WebSphere® CORBALOC cluster server address, listing all servers in the farm.

Data Availability Characteristics

Clusters share data between cluster nodes by allocating logical volumes in a common disk enclosure between cluster participants. However, in a disaster recovery environment each cluster is geographically separated, which stretches the connectivity requirements of disaster recovery beyond the capabilities of storage area network (SAN) and small computer system interface (SCSI) attached storage.

You must replicate data from the primary site across a wide area network (WAN) or virtual private network (VPN) to a disaster recovery site. The replication strategy largely depends on the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) of the business in the context of the FileNet P8 application.

Data relating to FileNet P8 that requires replication include, but are not limited to, databases, file stores, AE configuration files, and Verity collections.

Case Study: Veritas Volume Manager 5.0, Veritas Cluster Server 5.0

This study outlines the steps required for deploying FileNet P8 for disaster recovery using the Veritas Storage Solutions 5.0 product line. It also includes instructions for configuring DNS and load balancing software. Users seeking a low impact backup strategy can also use Veritas NetBackup Enterprise Server.

For a high-level view of the disaster recovery environment implemented in this case study see Figure 4: Case Study Disaster Recovery Environment on page 32.

For a more detailed view of the Global Application Configuration of same environment, see Figure 1: Global Application Configuration on page 19.

Products Used in this Study

- Backup and Recovery

  NetBackup Enterprise Server Version 6.0
- Logical Volume Management
  - Veritas Storage Foundation Enterprise HA Version 5.0
  - Veritas Volume Manager Version 5.0
  - Veritas File System Version 5.0
  - Veritas Volume Replicator Version 5.0
- Clustering
  - Veritas Storage Foundation Enterprise HA Version 5.0
  - Veritas Cluster Server Version 5.0
- FileNet P8 Platform
  - Content Engine
  - Process Engine
  - Application Engine
  - Content Search Engine
  - eForms
- Oracle 10g R2 RDBMS
- WebSphere® Network Deployment 6.0.2.17

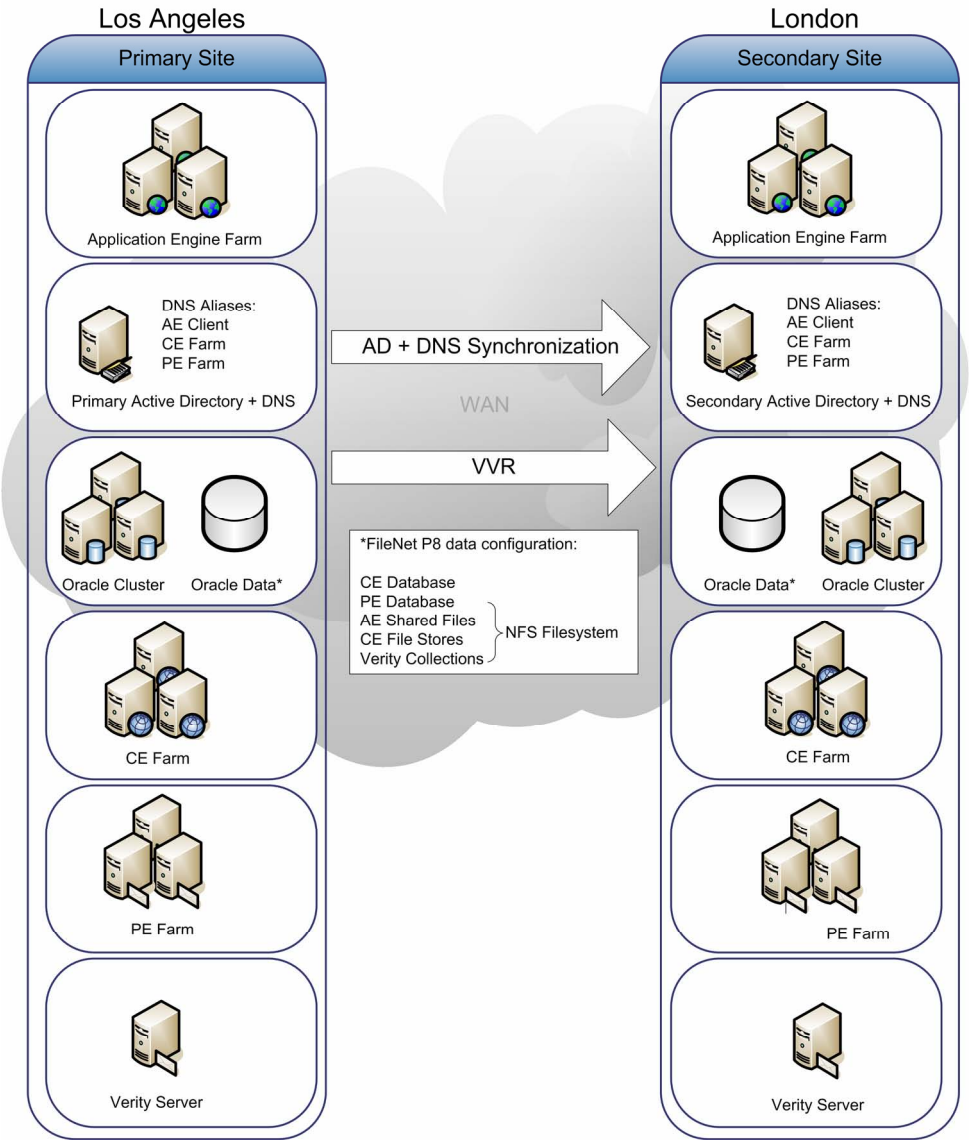

Figure 1: Global Application Configuration

Overview

For this case study, each site (primary and disaster recovery) consists of:

- A two node WebSphere cluster for CE.
- A two node PE farm.
- A two node WebSphere cluster for AE.
- A two node Oracle cluster.

Each server points to a Windows 2003 Active directory domain, for user authentication and DNS resolution. All of the relevant virtual hostnames are CNAME's in this DNS. The primary domain controller (DC) and DNS Server are located at the primary site, and the disaster recovery DC and DNS Server at the disaster recovery site. All of the shared data, including the two Oracle databases (one for CE and one for PE), the AE shared folder, CE file stores, and Verity collections, resides on one Veritas logical volume that is replicated from the primary site to the disaster recovery site. This single file system is accessible to

all servers with NFS. The two Oracle clusters are members of a global cluster that handles the VVR failover and the DNS updates.

To fail over from the primary to disaster recovery site, use the following sequence:

1. Shut down all FileNet P8 software at the primary site (CE, PE, and AE).
2. Use the Oracle global cluster to migrate all of the shared data from the primary site to the disaster recovery site, change the direction of VVR replication (so the disaster recovery is now the master), and perform the DNS updates (see description of this process below).
3. Verify that the shared file system (hosted by the global Oracle cluster) is mounted on all relevant servers at the disaster recovery site.
4. Verify that the DNS cache on each client system has been updated.

   Perform an “ipconfig /flushdns” on Windows systems, or the appropriate command for your particular UNIX® platform.
5. Start the FileNet P8 components (CE, PE, and AE) at the disaster recovery site.

Implementation Task List

The following is a general outline of the major installation tasks that must be performed.

NOTES

- These are the general steps followed in this case study, and my not accurately reflect the steps required for a particular disaster recovery solution.
- To successfully complete these tasks you must have access to all relevant third-party and IBM FileNet documentation.

To set up a disaster recovery environment

1. Install Veritas NetBackup Enterprise Server on the system designated to be your backup server.
2. Install Oracle.

   On the four nodes designated for the Oracle clusters (two nodes at the primary site and two at the disaster recovery site) and data replication, do the following:

   a. Install Veritas NetBackup Client

   b. Install Veritas volume manager.

      Create a Veritas volume group for the Oracle databases, the AE shared folder, the Verity Collections, and the file stores.

      Verify that block sizes and volume lengths do not differ between the sites, as this can introduce unexpected behavior after site failover.

   c. Use NFS to share the AE directory, the Verity collections, and the file store directory

   d. Install Veritas Cluster Server.

   e. Install Oracle RDBMS.

   f. Install Veritas Enterprise Agent for Oracle.

   g. Configure the two Oracle Clusters:

      i. One cluster at the primary site.

      ii. One cluster at the disaster recovery site.

      Test failover at each site.

h. Install Veritas Volume Replicator.

i. Install the Enterprise Agent for Veritas Volume Replicator.

j. Configure VVR to replicate the entire Veritas volume group from the primary site to the remote disaster recovery site.

k. Create the Content Engine database instance on the replicated volume group.

l. Create the Process Engine database instance on the replicated volume group

i. Configure the global cluster, uniting the two Oracle clusters.

ii. Configure the replicated volume group (See the VVR section for details on the global cluster).

3. Install Content Engine.

a. Install WebSphere Application Server Network Deployment on the four CE nodes (two at the primary site and two at the disaster recovery site).

b. Install CE on the four CE nodes (two at the primary site and two at the disaster recovery site).

c. Set up load balancing/proxy for CE.

There will be two load balancers, one at the primary site, balancing the load between those two CE nodes, and one at the disaster recovery site, balancing that load. Configure the CNAME alias in DNS to point to this load balancer.

NOTE Content Engine load balancing using software or hardware is only supported for CE web services connectivity (HTTP). Load balancing or proxy redirection of the EJB transport is not supported.

d. (Optional) Configure File Stores.

If you are going to use file stores, create a file system on the replicated volume group (managed by the Oracle cluster) and export it with NFS. Mount this directory on each CE node. Create your file stores here.

4. Install Application Engine.

a. Install WebSphere Application Server, Network Deployment, on the four AE nodes (two at the primary site and two at the disaster recovery site).

b. Install AE on the four AE nodes (two at the primary site and two at the disaster recovery site).

Create a file system on the replicated volume group (managed by the Oracle cluster) and export it with NFS. Mount this directory on each AE node. When the installation wizard asks where to locate the configuration files, point to this directory.

c. Set up load balancing/proxy for AE.

There will be two load balancers, one at the primary site, balancing the load between those two AE nodes, and one at the disaster recovery site, balancing that load. Configure the CNAME alias in DNS to point to this load balancer.

5. Install Content Search Engine on each Content Search Engine node.

6. Install Process Engine.

a. Install PE on all four PE nodes (two at the primary site and two at the disaster recovery site.) Configure each pair to be a PE farm.

b. Configure all of the PE servers to belong to one farm configuration.

CAUTION All PE nodes must be members of the same farm. This allows the farm nodes on the disaster recovery site to pick up work from the primary site in the event of a failover, and vice-versa.

WARNING Only nodes on the currently active site (primary or disaster recovery) should be running at any given moment.

7. Test and verify site failover.
8. Configure and test NetBackup policies.

Configuration Specific to the FileNet P8 Platform

Veritas Cluster Server Global Cluster Option (VCS)

At the top of the disaster recovery environment sits the Veritas Cluster Server Global Cluster Option, which centralizes the administration of all VCS clusters in the disaster recovery environment and manages global applications. A global application is a logical set of service groups spanning multiple VCS clusters. Global Clusters leverage these cluster service groups to control and monitor distributed applications such as Oracle.

After all highly available clusters have been implemented a global cluster is configured to coordinate the failover of cluster service groups to a disaster recovery site. This creates a cluster service group. VCS does not directly interact with any FileNet applications. Instead, the cluster service group manipulates VCS and VVR to react to events across the site. At the site level a highly available global cluster service group controls the operation of the global cluster. The disaster recovery site also employs a global cluster service group. The two service groups share information about the state of the site and health of global applications using heartbeats. When fault occurs, this service group coordinates the failover of global applications.

Hardware

Hardware requirements for the VCS software are minimal and the software itself should not impact the performance of other applications. VCS does require additional infrastructure to create a global application.

The application service group makes a change to DNS to steer users to the correct site. This feature is important due to the logistic complications of moving a virtual IP across a WAN. You must maintain DNS servers to manage resolution of hostnames to IP addresses. For specific information related to DNS, see the “Windows 2003 DNS servers“ section.

Software

Before attempting to configure a global cluster, verify that the preliminary testing of the base components, such as replication and volume groups, is complete. These steps are absolutely critical to the proper operation and configuration of a global cluster. If the data doesn’t replicate properly or one of the clusters doesn’t failover and return cleanly, a functional disaster recovery site failover will fail..

Installation and Configuration

Figure 2: VCS configuration illustrates the VCS configuration used in this case study. This is the configuration of one of the sites only. The other site must reproduce this VCS configuration exactly (except for the properties of the DNS aliases and the other resources unique to each site) or a global failover will fail.

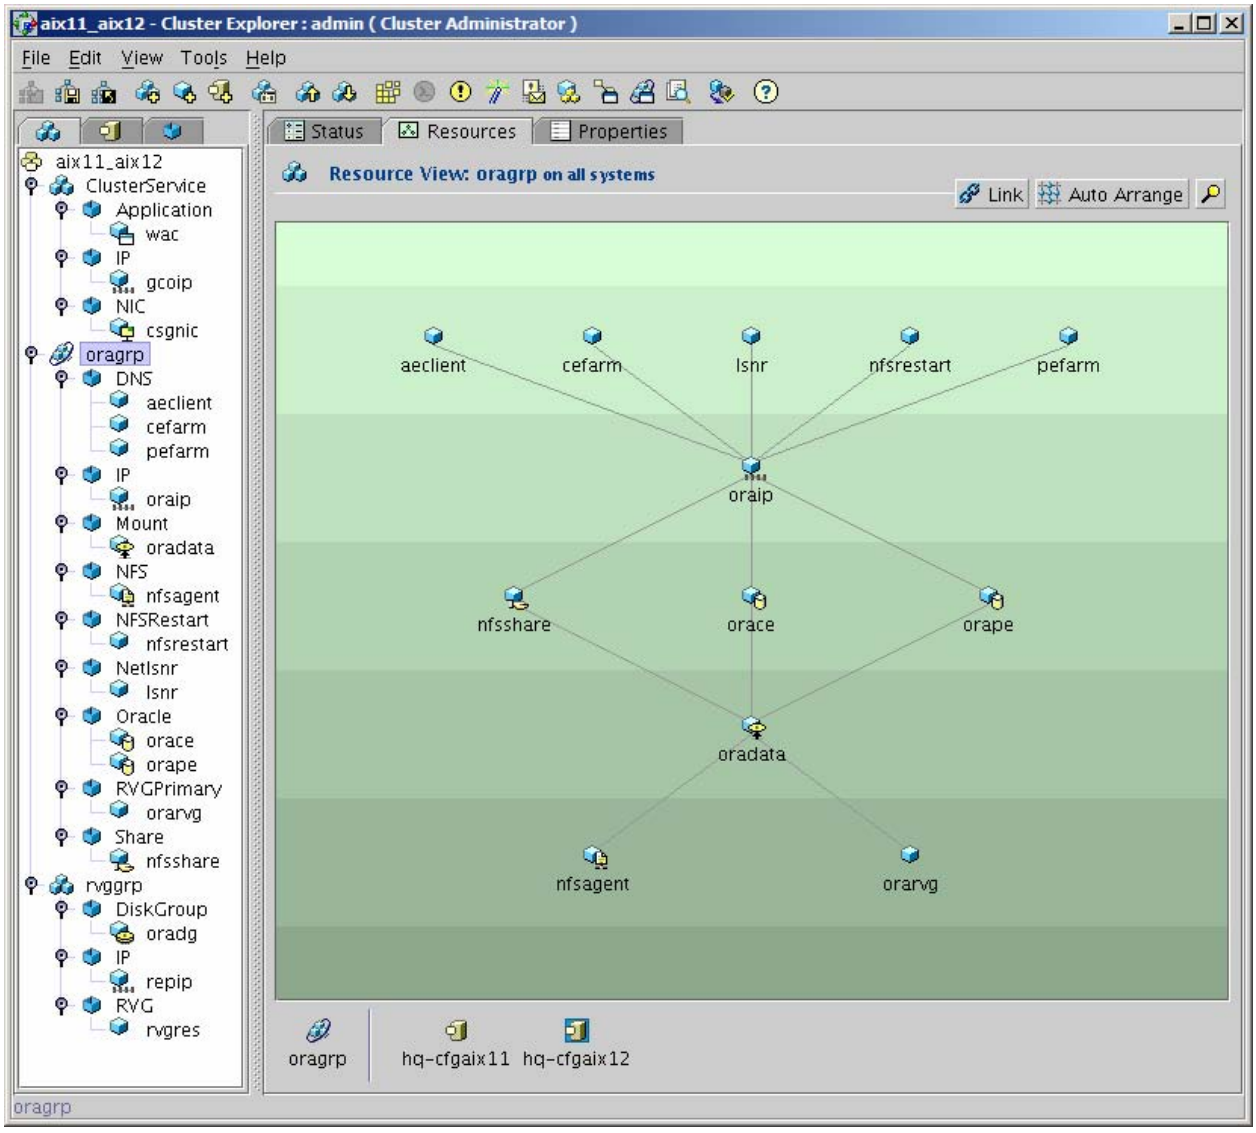

Figure 2: VCS configuration

This configuration contains three service groups:

1. The ClusterService group

   The ClusterService group is configured when VCS is configured to be global, and allows the two clusters at the two separate sites to communicate with each other. The three resources in this group include:

   - The Network Interface Card (NIC) they will use to communicate.
   - An IP address for cluster communication only.
   - A wide-area connector (WAC) application, which VCS uses for cluster communication.

2. The rvggroup

   The rvggroup is the group that controls VVR replication. If a cluster fails over to the other local node (for example, NOT a global failover), this group switches ownership of the replicated volume group to the new node and maintains replication to the remote site. If the cluster fails over to the other global cluster, this group switches primary ownership of the replicated volume group to the remote site and reverses the direction of replication. The three resources are:

- The Veritas Disk group (used in replication).
- The IP address used for replication (note there should also be a dedicated NIC for this resource, as replication traffic can be heavy).
- The replicated volume resource (rvgres). Refer to VCS and VVR documentation for details on how to configure this group.

3. The oragrp

   The oragrp is the group that controls the Oracle databases and the DNS aliases. It is the group that is expanded in the diagram above, showing the resource dependencies.

   - The nfsagent resource controls operating system NFS daemons. The orarvg resource is the primary replicated volume group resource.
   - The oradata resource mounts the file system that resides on the replicated volume group. This file system contains all of the shared data, including the Oracle databases, the AE shared files, the CE filestores, and the Verity collections.
   - The nfsshare resource shares this file system with NFS.
   - The two Oracle resources, orace and orepe, are the CE and PE oracle databases.
   - The oraip resource is the IP address used by oracle clients to connect to the Oracle database. (Note that again there should be a dedicated NIC for this resource).
   - The lsnr resource is the Oracle listener.
   - The nfsrestart resource allows clients of the NFS directories to maintain the NFS mounts in the event of a local cluster failover.
   - The three DNS resources, aeclient, cefarm, and pefarm are the three DNS CNAME aliases that get updated in the event of a global cluster failover. If all FileNet P8 applications are configured using these aliases (for example, PE connection points use the pefarm alias, and so on), then no reconfiguration of FileNet P8 will be required upon global failover. End users will use the aeclient alias to connect to workplace, as follows: http://aeclient/Workplace. This alias will always point to whichever site is running.

The rvgroup and oragrp are linked with a dependency of “online local hard.” Refer to VCS and VVR documentation for details on this, and all other aspects of VCS and VVR configuration.

Veritas Volume Replicator (VVR)

Veritas Volume Replicator (VVR) is a core technology for disaster recovery. VVR replicates volumes by intercepting block-level write() system calls to volumes in the Replicated Volume Group (RVG) and duplicating them to peer clusters at disaster recovery sites. The major concern in VVR replication is the difference in I/O throughput between local volumes and WAN-replicated volumes. To address this problem, VVR has two replication modes: synchronous I/O and asynchronous I/O.

- Synchronous I/O suspends write operations until all blocks written at the primary site have been replicated to disaster recovery sites. Various issues such as disk and WAN latency, however, might impact the performance of this mode of replication. The main benefit with this mode is that the remote site is always guaranteed to be synchronized with the primary site.
- Asynchronous I/O allows the write() operation to return as soon as the blocks have been queued for replication. The impact of asynchronous writes on I/O is minimal, but WAN latency often causes the disaster recovery site to be several I/O operations behind the primary. With asynchronous I/O, you risk losing transactions that have not yet been replicated in the event of a site failure.

Hardware

VVR does not require additional hardware.

WAN connectivity to the remote site should be redundant. WAN connectivity with a single point of failure greatly diminishes the reliability of the disaster recovery environment and exposes the environment to the risk of becoming “split-brained”. A split-brain condition occurs when the cluster server master at each site believes the other site has faulted, and that its node is now the primary site. To avoid split-brain, use redundant links for replication and cluster server communication.

Installation and Configuration

Installing VVR for use with FileNet P8 components is covered in each component’s respective section of the Case Study. In general, a VVR secondary replication log (SRL) should be configured with the same parameters as the volumes it replicates. For example, if File Store data is stored on a logical volume that spans six spindles, the SRL should do the same. This is critical to maintaining optimal performance of FileNet P8 components.

Windows 2003 Active Directory

This case study uses Microsoft Windows active directory for user authentication.

Hardware

Two servers will be required, one at each site.

See the Microsoft website for hardware consideration details, for running Microsoft Windows 2003 active directory servers.

Software

The Microsoft Windows active directory Server is part of the Windows 2003 Server operating system.

Installation and Configuration

For detailed information on installation and configuration of Windows active directory, see Microsoft documentation.

Windows 2003 DNS Server

This case study focuses on the configuration of a Microsoft Windows 2003 DNS server. The DNS server directs users to the correct site after a site failover occurs.

Hardware

Two servers will be required, one at each site.

See the Microsoft website for hardware consideration details, for running Microsoft Windows 2003 DNS servers.

Software

The Microsoft Windows 2003 DNS Server is part of the Windows 2003 Server operating system.

Installation and Configuration

This document covers the high-level steps necessary to configure DNS for disaster recovery of the FileNet P8 Platform. For detailed information on the installation and configuration of Windows DNS servers, see Microsoft documentation.

For this case study CNAME aliases are created on the DNS server for each globally connected component. The Content Engine, Process Engine, and Application Engine farms will all have CNAME aliases to switch between physical IP addresses or load balancer virtual addresses.

NOTE Each CNAME alias must be able to accept updates from any IP addresses that might initiate an update. In this case this should be the cluster server master at either site.

Oracle 10g R2

Disaster recovery configurations for Oracle are beyond the scope of this document. See the Symantec and Oracle web sites for instructions on how to configure Oracle with Veritas.

Content Engine

Set up a Content Engine application server farm at each site as described in the *IBM FileNet P8 Platform High Availability Technical Notice*.

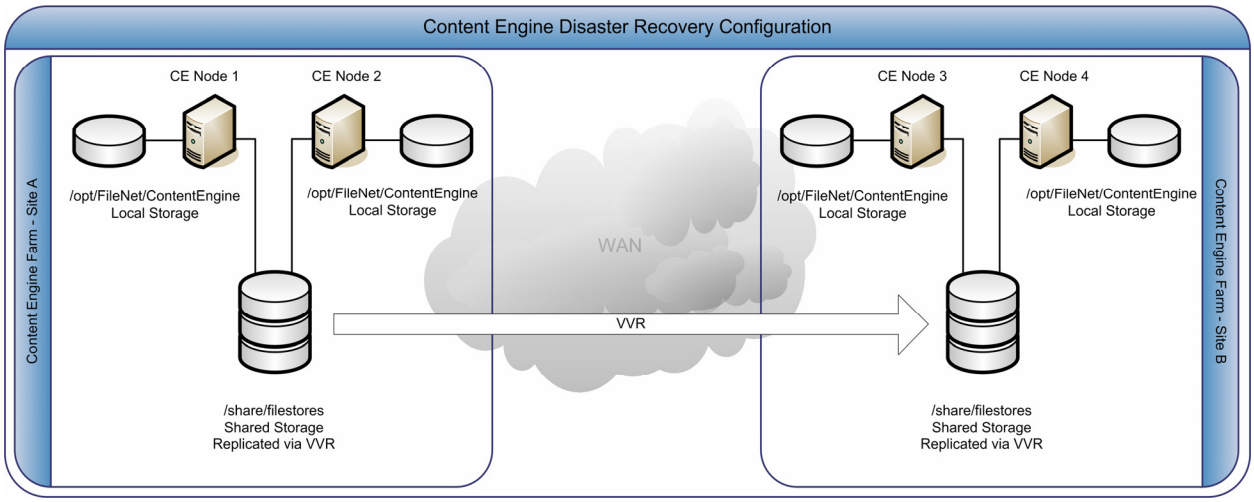

Figure 3: Content Engine Disaster Recovery Configuration

Hardware

Content Engine requires no additional software other than that which is required for two highly available CE farms.

All replication activity should occur over a dedicated network interface to minimize the impact replication has on the application and vice-versa.

For details on the hardware requirements for Content Engine, see the *IBM FileNet P8 Platform Installation Guide* and the *IBM FileNet P8 Platform High Availability Technical Notice*.

Software

Content Engine does not require any additional software for disaster recovery other than third party products to ensure data replication and global clustering.

Installation and Configuration

Installing Content Engine for disaster recovery environments is similar to installing in a highly available environment. When installing Content Engine create all of the nodes in the same FileNet P8 domain. For this case study we use four nodes. The first two nodes are at the primary site part of one application server cluster and the other two nodes are at the disaster recovery site part of a second application server cluster. Configure both application server clusters to point to its local database cluster server address.

The file stores resource used by object stores within CE, use an NFS resource. This NFS resource is also part of volume replication by VVR. The NFS resource must be mounted on the CE nodes when a given site is online. The NFS mount uses a local cluster server address.

The hostname that points to the load balancer must be a CNAME alias in DNS. When FileNet P8 runs on the primary site, this alias will point to the load balancer, or proxy, on that site. When a failover to the remote site occurs, this CNAME alias is updated to point to the load balancer on the remote site. AE will use this hostname to find CE for Component Manager. AE primary connection to CE will be using application server specific cluster address such as WebSphere CORBALOC or WebLogic T3.

Process Engine

PE is installed in a farm configuration, on both the primary and disaster recovery site, as documented in the *IBM FileNet P8 Platform High Availability Technical Notice*.

Hardware

PE in a disaster recovery environment does not require any additional hardware beyond that of high availability. Refer to the *IBM FileNet P8 Platform High Availability Technical Notice* for details concerning farming PE at each site.

Software

PE requires no additional software other than that which is required for two highly available PE farms.

Installation and Configuration

Install PE so that all servers are members of the same farm environment. This allows the PE farm nodes at one site to perform work on objects created from farm nodes at a remote site in the event of a failover.

The only additional configuration steps beyond what is covered in the *IBM FileNet P8 Platform High Availability Technical Notice* is that the hostname that points to the load balancer must be a CNAME alias in DNS. When FileNet P8 is running on the primary site, this alias will point to the load balancer on that site. When a failover to the remote site occurs, this CNAME alias is updated to point to the load balancer on the remote site. Use this alias on CE when configuring the PE region ID. This way, CE will always find the PE system that is currently running.

Update the hosts file for each PE server with the NCH name of each of the PE servers, to allow proper communication between PE servers on different subnets.

NOTE This is a precautionary step in case the administrator would like to bring online PE servers on the remote site for testing/verifying, while the primary server is running.

Example of hosts file entries:

10.15.16.101	peserver1	peserver1-filenet-nch-server
10.15.8.182	peserver2	peserver2-filenet-nch-server

Application Engine

Set up an AE web farm at each site as described in the *IBM FileNet P8 Platform High Availability Technical Notice*. Set up each AE farm with redundant load balancers, to ensure that a single failed load balancer does not cause the entire global application to be migrated to the disaster recovery site. If the load balancers have more than two ports, cross-connect them to a switch and to the AE servers. Most modern load balancers use proprietary HA mechanisms, which should be enabled. Since AE does not integrate directly with Veritas Cluster Server or Global Cluster Manager, no further configuration is required.

Hardware

AE in a disaster recovery environment does not require any additional hardware beyond that of high availability. Refer to the *IBM FileNet P8 Platform High Availability Technical Notice* for details concerning farming Application Engine at each site.

Software

AE does not require any additional software or configuration for disaster recovery.

Installation and Configuration

AE requires little additional installation as it does not participate in a cluster. The exception to this rule is the shared location of several AE configuration files. IBM recommends that this share be part of the replicated volume group used by the Oracle database.

In cases where AE runs on a UNIX platform, use a highly available replicated NFS resource on a UNIX cluster that participates as part of a cluster server global application. The NFS resource will be used to share a common set of AE configuration files such as the bootstrap.properties file. This NFS resource is part of volume replication by VVR. The NFS resource must be mounted on the AE nodes when a given site is online. The NFS mount uses a local cluster server address.

The hostname that points to the load balancer or proxy must be a CNAME alias in DNS. When FileNet P8 is running on the primary site, this alias will point to the load balancer on that site. When a failover to the remote site occurs, this CNAME alias is updated to point to the load balancer on the remote site. Clients will use this alias to find Workplace. For example, instead of http://Hostname/Workplace, clients will use http://alias/Workplace. Since this alias will point to the load balancer at whatever site is running FileNet P8, clients will always be able to find it.

Content Search Engine

Content Search Engine is a dynamic application where the data of the application can change. The data In this case are the Verity collections.

Hardware

IBM recommends that all replication activity for the collections occur over a dedicated network interface to minimize the impact replication has on the application and vice-versa.

The Content Search Engine in a disaster recovery environment consists of two Master servers, one at each site. The collections are replicated from one site to the other using Veritas VVR.

Hardware requirements for Content Search Engine are detailed in the *IBM FileNet P8 Platform Installation Guide* and the *IBM FileNet P8 Platform High Availability Technical Notice*.

Software

Content Search Engine does not require any additional software for disaster recovery other than third party products to ensure data replication and global clustering.

Installation and Configuration

Installation and configuration of Content Search Engine is documented in the *IBM FileNet P8 Platform Installation and Upgrade Guide*.

CAUTION Both production and disaster recovery Content Search Engines in a Disaster Recovery environment must be configured to be Master servers, and only one will be active at any given time. Both servers will use the shared mounted location of the collections. In addition, there can be more than one Content Search Engine server in the primary site. You must configure an identical suite of servers in the disaster recovery site.

After a site failover, perform the following tasks using EM at the domain level:

1. Update the Verity master server hostname and port on the Verity Domain Configuration tab to point to the Content Search Engine on the disaster recovery site.
2. Update the list of broker servers on the Verity Server tab to point to the servers available on the secondary Content Search Engine.
3. Update the index areas for each object store to reflect the search and index servers available on the disaster recovery site's Content Search Engine. This update also triggers the import of the collections to a secondary Content Search Engine if it doesn't already exist. The index on the collections will persist to the secondary Content Search Engine so a re-index isn't necessary.

When the primary site is back online, remove the collections from the remote site's Content Search Engine, using Verity's dashboard. When you remove the collection, only remove the reference of the collection in Verity and not the actual directory and file system. The purpose for doing this is to ensure that the Content Search Engine is in a clean state for when another site failover occurs.

Veritas NetBackup Enterprise Server

Users who will be implementing Global Clusters might wish to use the NetBackup Policy Tools to orchestrate the shutdown and startup of application components.

Hardware

No additional backup hardware, beyond what is required for a production site, is required for a disaster recovery configuration.

Software

If you use Veritas Volume Manager for logical volume management, no additional software is required for this solution beyond that of traditional backups. However, depending on how Volume Manager was bundled, extra Veritas licenses might be required.

Installation and Configuration

Implement the Volume Manager snapshots after the Global Cluster Option (GCO) and Veritas Volume Replicator (VVR) have been configured, but before the go-live date. This will ensure that all prerequisites for the environment are met and that backup requirements are fully realized.

To further shrink backup times, consider Veritas Fast Resync when implementing a snapshot solution. For more information about Veritas Fast Resynch, see the *Veritas Volume Manager documentation*.

Backup and Recovery

For procedures on administering backups and restores with NetBackup, see the Veritas NetBackup Administration guide.

NOTE For this case study “cold” backup and restore was tested, where all FileNet P8 software and all support software (such as Oracle) was shut down prior to making the backup.

Appendix One: Case Study Environment

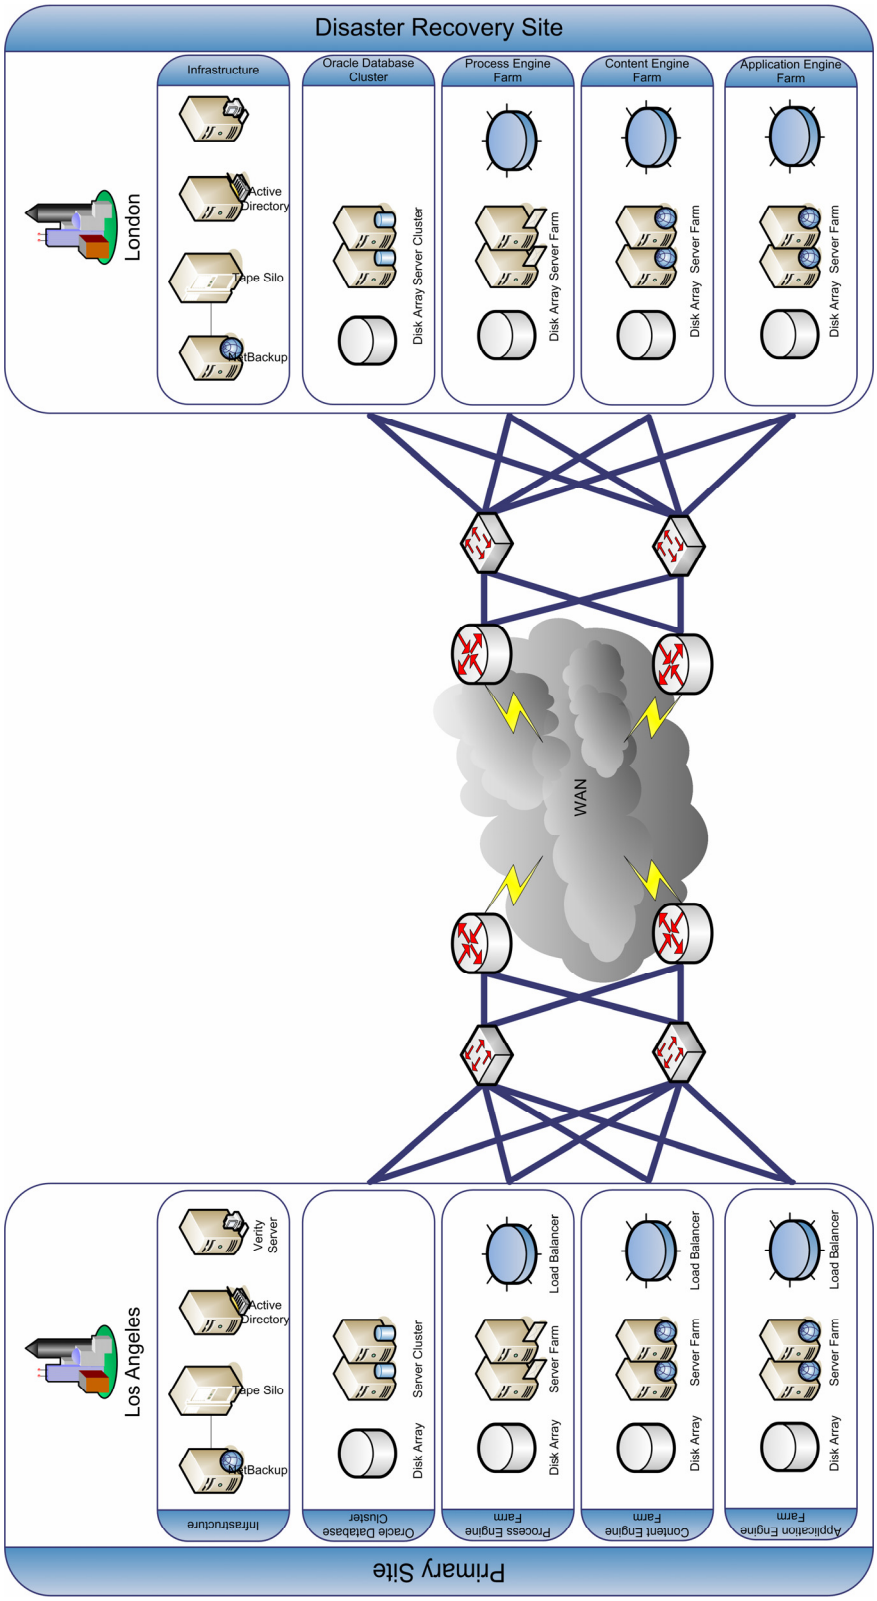

Figure 4: Case Study Disaster Recovery Environment

Appendix Two: Identifying Requirements

The first step in planning a disaster recovery environment is to create a prioritized specification of the business objectives that a disaster recovery initiative must fulfill. This specification ensures that project deliverables are in line with business objectives and that the disaster recovery architecture addresses the problem. Also clearly distinguish between essential and non-essential objectives.

- An essential objective might be to enable the business to resume normal operations within a half an hour of a site failure.
- A non-essential objective might be to ease administration of operational procedures in an effort to reduce cost.

A cost-benefit analysis is essential to planning a disaster recovery environment, and will quantify the necessary capital required to fulfill business recovery objectives and weigh them against the potential financial impact of a disaster. As with most IT expenditures, disaster recovery costs are justified on the basis of Return on Investment (ROI) of the expenditure. ROI is the difference between the total potential loss of revenue and the Total Cost of Ownership (TCO) of the project.

A disaster recovery project must be free of conjecture and speculation. Without clear justifications, a disaster recovery initiative can produce an under-funded project with unrealistic deliverables. Changes in design to meet evolving business objectives or budgetary constraints during project implementation might introduce shortcomings and design flaws, effectively defeating the benefits intended by the project. Therefore, evaluate the TCO before implementing the disaster recovery project. Failure to complete the proper calculations can produce a disaster recovery environment that falls short of business continuance objectives and severely undercuts the projected ROI.

The following is a general guide for calculating the ROI of a disaster recovery environment.

To calculate ROI

1. Estimate the impact of a complete outage of business operations.

   Calculate the potential revenue loss over a projected period of time required to resolve the outage. A rough figure can be calculated using the total net revenues over the last eight quarters divided by the total operational hours in the same period. This figure is the projected revenue per hour. Then multiply this figure by the total number of projected hours of downtime.

   The projected time period should be a realistic prediction of the duration of time between the occurrence of a disaster and the resumption of operations. Outage experience should weigh heavy in this calculation.

   Consider the following factors when estimating potential downtime.

   - Geological data should be sought to gauge the risk of sinkholes, earthquakes, mudslides, and volcanoes. Scrutinizing geological factors help identify the level of potential damage related to geologic activity. Geologic history should also be consulted when estimating a recovery timeframe based on these factors. For example, corporations in the Silicon Valley area should evaluate downtime based on metrics realized during the last major earthquake.
   - The same holds true for meteorological activity. For example, corporations in Florida should consider the frequency and impact of tropical storms when projecting downtime.
   - Another factor to consider is the frequency of construction in the area. If there is a disproportionate amount of construction activity in the area the exposure to risk is greater than for data centers in stable areas.
   - Indirect factors. Disasters can also indirectly create potential periods of downtime, such as rolling power outages and access to key business partners. For example, a business's key supplier might be unable to use nearby roads or highways to deliver product. The target facility would be considered unavailable because of the indirect factor contributing to the inaccessibility of the facility. Adjust your downtime projections to account for these indirect factors.

Calculating downtime projections

a. Factors Influencing Downtime Costs:
    i. Costs of employees who are unable to perform essential job functions.
        ◊ Compensation
        ◊ Benefits
        ◊ Cost of office space
        ◊ Equipment leases
    ii. Fiduciary losses caused by business interruption.
        ◊ Lost sales
        ◊ Sliding market share
        ◊ Loss of customer confidence
        ◊ Stock valuation
        ◊ Legal liabilities public and private
b. Factors Influencing Recovery Costs:
    i. Cost of infrastructure replacement.
        ◊ Raised floors
        ◊ Network Infrastructure
        ◊ Server hardware
        ◊ Tape silos
        ◊ Software Licensing
        ◊ Software Media
    ii. Labor costs.
        ◊ Consulting Services
        ◊ Data recovery logistics (i.e. restoring from tape)

2. Calculate the Recovery Time Objective (RTO).

   The RTO is the delta of time between when the user community loses access to the application to when the application is available again. The design of the underlying infrastructure will be the greatest controllable factor to fulfilling an RTO. As such it significantly influences TCO.

   All replication technologies assume the inherent risk of data loss; however, most replication products allow administrators to assert some degree of control over the replication behavior. The data in a dynamic application is constantly changing, and every one of these changes must be replicated to the alternate site to ensure both sites are in sync. In a perfect world the data would always be identical at both sites; however, in the real world data must travel over a WAN which introduces a rate limitation on the data throughput to the disaster recovery site. When the commit rate at the primary site eclipses the capabilities of the WAN to replicate to the disaster recovery site, the primary must either wait for data to sync up or queue the operation and return control to the application.

To address this limitation Veritas Volume Replicator (VVR) supports synchronous and asynchronous replication and storage snapshots.

- Synchronous Replication

  When the data commit rate is lesser than the throughput of the WAN connection, use synchronous writes to ensure that both sites contain the exact same data. However, distances and the cost of high bandwidth WAN connectivity sometimes render synchronous replication impractical.

- Asynchronous Replication

  The alternative is asynchronous replication. Under Veritas Volume Manager, replicated volume groups are known as a replicated data set (RDS). When the operating system receives a write() system call on a volume in an RDS, the kernel directs the vxio driver to write n bytes to a specified file descriptor. When vxio completes, write() returns the number of bytes that were written. A write to an RDS in synchronous mode requires each and every write() call to be replicated to the remote site before it returns the length of the write.

  These write operations travel over a WAN which is far slower than local disk and severely impedes performance of write() calls on synchronously replicated volumes. To circumvent this limitation, VVR can handle write() operations asynchronously. When vxio receives a write() on an asynchronous RDS, write() returns the number of bytes written as soon as it hits the system buffer cache and the secondary replication link (SRL) is updated. At this time write() returns the number bytes written, returning control to the application. This write is then copied to the local volume and replicated to the disaster recovery site.

  This function works well, but distributed applications require consistency across multiple servers. This problem is addressed using VVR snapshots. A VVR snapshot, when taken on each server while IO is halted, guarantees consistency across all FileNet P8 components. Various strategies exist for leveraging snapshots with replication. Consult the Veritas Volume Replicator documentation for strategies concerning VVR snapshots.

Due to the constraints on replicated data ,you must allocate sufficient WAN resources to satisfy the recovery point objective. The slower the replication performance the higher the potential is for lost transactions. For detailed information on VVR and the implications of varying replication configurations consult the Veritas Volume Replicator documentation.

3. Calculate the Total Cost of Ownership (TCO).

   Calculating TCO enables the business to measure the costs of mitigating the risk of an outage. The TCO of a disaster recovery environment is the sum of the expenditures required for implementation. These expenses include:

   - Tangible costs such as hardware and software.
   - Intangible costs such as the impact on manpower to address regularly assigned workloads.

   While calculating TCO, begin planning the physical aspects of the disaster recovery environment. This process helps to ensure the environment meets the RPO and RTO while keeping TCO below the ROI. In other words this process aids in designing an environment that addresses business requirements while keeping implementation expenditures below the projected cost of a disaster. When the TCO outpaces the ROI, the TCO can be reduced by experimenting with various hardware and software configurations without sacrificing business objectives. When TCO cannot be reduced to meet the ROI, you should re-evaluate the business objectives. .

   There are many creative ways to cut costs without sacrificing availability or business objectives. For example, consider the hardware at each site. Four three-node clusters could be consolidated to two N-M clusters, thus reducing the cost of hardware and administration dramatically. You could also control costs by scaling in a manner appropriate for the application. Clusters requiring high floating point performance should be scaled vertically and application components using 32-bit binaries often perform better on 32-bit architecture which is often less expensive than 64-bit architecture.

When calculating the factors influencing cost, also consider an employee's ability to address formal workloads. If an administrator is pulled from regular duties to contribute to a disaster recovery project, a back fill of responsibilities might be warranted. Given the previous, evaluate the administrative costs against the administrator to infrastructure ratio the business would prefer to maintain. For example, if the industry average for servers to administrators is fourteen servers per administrator, allocate resources to maintain this ratio. This number, of course, fluctuates with the type of operating system, the type of server, and how the servers are scaled. For example, in a UNIX environment, maintaining a web farm with many identical servers is less costly than the same number of servers running with unlike configurations.

a. Possible Implementation Costs
    i. Software
        - ◊ Licensing
        - ◊ Media
        - ◊ Service Contracts
    ii. Hardware
        - ◊ UPS
        - ◊ PDU
        - ◊ Storage
        - ◊ Servers
        - ◊ Peripherals
        - ◊ Switches
        - ◊ Routers
        - ◊ Firewalls
        - ◊ Power
        - ◊ Cabling
        - ◊ KVM and console switches
        - ◊ Racks
        - ◊ Service Contracts
    iii. Labor
        - ◊ Consultants
        - ◊ Contractors
        - ◊ Administrators
        - ◊ Managers
        - ◊ Architects
        - ◊ Electricians
b. Possible Recurring Costs
    i. Software
        - ◊ Service Contracts
        - ◊ Upgrades

ii. Hardware
    - ◊ Service Contracts
    - ◊ Upgrades
    - ◊ Incident Response

iii. Labor
    - ◊ Recurring Consultation
    - ◊ Contractors

4. Calculate the ROI (differenc e of cost of downtime and TCO).

   This number is the potential operational surplus realized by implementing a disaster recovery environment.

Appendix Three: Project Planning

The design and implementation of a disaster recovery environment requires careful planning and coordination. Many groups and individuals, each with different levels of expertise and responsibility, must work together to successfully deliver a disaster recovery environment that satisfies business objectives.

High-level Overview

A high-level overview that breaks down each component greatly simplifies planning and allows project managers to delegate the design of more technical aspects to those best qualified to handle them. For instance, Figure 4: Case Study Disaster Recovery Environment on page 32 is a high-level view of a disaster recovery environment. From this diagram we can deduce the number of servers required, the network infrastructure involved, and which components can be consolidated. From here tasks can be delegated to responsible groups.

For example, the design requires four WAN connections, which immediately identifies the network engineering group as the most appropriate resource for planning that portion of the disaster recovery infrastructure.

PE and CE each require an Oracle database instance as part of the installation. Changes in this configuration could influence the hardware and software required for the installation. Decide the name of the instances and the manner in which they are clustered ahead of time.

An overview enables project managers to identify opportunities for groups to collaborate on cross functional tasks and devise solutions which best meet the requirements set by the business.

Project Worksheet

Use a worksheet with configuration details and contact information for responsible administrators to iron out the technological minutiae of planning a disaster recovery environment. Make sure this worksheet is specific to business needs and tailored to the facilitation of the operational groups executing the plan.

Statement of Work

Use a written statement of work for each server and network component to plan each system implementation, to help scope each task and identify the resources needed to perform an installation.

Site Survey and Data center Preparation

A disaster recovery environment typically begins with a well-established primary data center in a location strategically significant to the business. This location requirement is usually not required for alternate data centers. This allows disaster recovery architects great flexibility when developing an alternate site.

When the primary data center is small or there is no established primary site, evaluate which site would be most suited as the primary. Although each data center should assert redundancy wherever possible, if budget constraints compel a reduction of infrastructure, the primary site should be given priority.

A well-hardened data center not only includes hardened servers network infrastructure, it also includes redundant UPS, diesel electric generators, and environmental systems. If possible, redundant WAN connectivity should be installed, preferably with separate providers, or even better, separate providers over separate local loops. Wherever possible, eliminate single points of failure.

While disaster recovery sites normally require preparation before the installation of hardware or software, take measures to protect from perils specific to a geographic location.

Example:

Data centers located close to fault lines where earthquakes are prevalent should consider reinforcing server and telecom racks. This not only includes reinforcing racks to the walls with ladders, but also ensuring that equipment is properly bolted into a rack appropriate for the application and that the units are physically installed as per the manufacturer's instructions. Fastening loose equipment is not only an easy way to reduce occupational hazard to employees, it also is the single most inexpensive way to prevent unplanned outages due to small tremors or accidentally pushing a server out of its rack.

Appendix Four: Verification and Documentation

The final step in developing a production-ready disaster recovery environment is a rigorous test of failover components and of global application functionality. The purpose behind these exercises is to discover ahead of time what issues arise when unexpected failures occur and to document procedures for handling emergencies. Testing should be conducted with a phased approach.

1. Administrator acceptance.

   Do the administrators grasp the process, understand the effort involved, understand the scope of the tasks for which they responsible?

2. User acceptance.

   Do the users find the application runs as it should? Are users able to do their jobs before, during, and after a failover?

3. Business acceptance.

   Does the disaster recovery environment meet business objectives? Does the executive body understand the impact to the business in a failure scenario? Does the environment deliver on executive expectations?

All of these questions must be answered definitively before you rely on a disaster recovery environment for production use.

Administrative Acceptance

The administrators responsible for each portion of the environment should be given ample time to test their components. This will ensure that the responsible personnel will have had at least minimum exposure to problem resolution in a disaster recovery scenario. For example DBAs must be allowed to test Oracle before and after a restore from tape. This will allow them to experience the process, understand the effort involved, and give scope to the tasks for which they responsible. Moreover, if unexpected behaviors surface during testing, it will allow DBAs time to investigate and resolve the issues without the pressures of the business units who depend on them. This process not only helps identify unexpected inconsistencies ahead of time, but also gives administrators the experience to resolve issues pertaining to disaster recovery and gauge potential maintenance windows.

As each functional group conducts its tests, other groups should carefully document anything that could not be retrieved with a standard backup including volume layout, the location of backup media, procedures for recalling backup media from offsite tape storage vendors, the location of the install media, license keys, WAN connectivity vendor and circuit information, and so on.

NOTE Administrative acceptance testing should include verifying licensing and support agreements. If demonstration keys were used during implementation they should be replaced with permanent keys as soon as possible. Moreover support agreements should be verified and information relevant to contacting technical support should be available to all responsible parties as well as management.

User Acceptance

Once each functional group signs off on their area of responsibility, test failovers while a subset of the user community exercises FileNet P8. Have users log into AE and perform tasks as they would during a regular workday including checking in documents, launching workflows, and changing document properties. It might also be useful for FileNet P8 administrators who can exercise the more advanced features of FileNet P8 to participate as well. Fail over one component at a time; sometimes with user knowledge, sometimes without and have them give feedback. If any deficiencies are uncovered, address them and retest the component.

Once the user community is satisfied with component level failovers, retest with component failovers at the site level and then test failovers of the entire site. Again, collect user feedback and address issues as they arise.

Executive Acceptance

Once the administrative and user communities are satisfied with failover performance, the final acceptance test is to measure the success of the project against the expectations of the business. Executive acceptance is important to instill confidence that the disaster recovery environment accomplishes the task of protecting share holder interests and is a fiducially sound and responsible investment. A disaster recovery environment as grand and complex as one that supports FileNet P8 requires the business to exude the utmost confidence in the ability of its people and the products that support global availability. This exercise will ensure that when a catastrophe erupts, the business can be wholly confident that its technical staff and their disaster recovery investment will weather the storm and survive in the face of unfortunate circumstances.

There are two areas of evaluation.

1. The first area should evaluate project performance. Did the project deliver the intend result?
2. The second area should evaluate the total cost of ownership against the return on investment. This process will identify the final value the project effort represents.

Project Performance

Compare the project plan against the project results. Identify bottlenecks and set backs, discuss them, and objectively evaluate their impact to the project and its deliverables. Areas which identify difficulties relating to hardware or software implementation, for instance, should be evaluated to determine root cause. In many cases, this process can reveal staffing or educational shortcomings, and identify opportunities to strengthen future project management and execution. Compare the recovery time objectives (RTO) and recovery point objectives (RPO) to experiences during the administrative and user acceptance testing.

Total Cost of Ownership

Finally, compare the actual value of the project to the projected value. This involves a recalculation of the TCO using actual receipts, invoices, and purchase orders. Compare the TCO against the projected TCO. If the implementation went smoothly and cost estimation was reasonably inline with cost experience, the actual TCO should fall below or close to the projected TCO. If the difference between actual and projected TCO vary greatly then project performance should reevaluated to identify where the problem is. Try to identify under-estimations or miscalculations to account for the difference, or if the difference is caused by unexpected difficulties or even scope creep.

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.

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

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

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

U.S. Patents Disclosure

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.

Program Number: 5724 R76, 5724 R81

Printed in USA

GC31-5486-00