AUTHORIZED DOCUMENTATION

Administration Guide for NetWare® 6.5 SP8

Novell®
Business Continuity Clustering

1.1 SP2

September 23, 2010

www.novell.com

Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For Novell Trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide

This guide describes how to install, configure, and manage Novell® Business Continuity Clustering 1.1 Support Pack 2 (SP2) for NetWare® 6.5 SP8 in combination with Novell Cluster Services™ 1.8.5 for NetWare clusters.

Audience

This guide is intended for intended for anyone involved in installing, configuring, and managing Novell Cluster Services™ for NetWare in combination with Novell Business Continuity Clustering for NetWare.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html (http://www.novell.com/documentation/feedback.html) and enter your comments there.

Documentation Updates

The latest version of this *Novell Business Continuity Clustering 1.1 SP2 Administration Guide for NetWare 6.5 SP8* is available on the Business Continuity Clustering Documentation Web site (http://www.novell.com/documentation/bcc/).

Additional Documentation

For information about Novell Cluster Services for NetWare, see the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

For information about Novell Identity Manager 3.5.1, see the Identity Management Documentation Web site (http://www.novell.com/documentation/idm35/).

For information about NetWare 6.5 SP8, see the NetWare 6.5 Documentation Web site (http://www.novell.com/documentation/nw65/).

For information about using Novell Business Continuity Clustering 1.2 for OES 2 Linux, see the BCC Documentation Web site (http://www.novell.com/documentation/bcc/).

For information about OES 2 SP1 for Linux, see the OES 2 Documentation Web site (http://www.novell.com/documentation/oes2/).

Overview

1

As corporations become more international, fueled in part by the reach of the Internet, the requirement for service availability has increased. Novell® Business Continuity Clustering (BCC) offers corporations the ability to maintain mission-critical (24x7x365) data and application services to their users while still being able to perform maintenance and upgrades on their systems.

In the past few years, natural disasters (ice storms, earthquakes, hurricanes, tornadoes, and fires) have caused unplanned outages of entire data centers. In addition, U.S. federal agencies have realized the disastrous effects that terrorist attacks could have on the U.S. economy when corporations lose their data and the ability to perform critical business practices. This has resulted in initial recommendations for corporations to build mirrored or replicated data centers that are geographically separated by 300 kilometers (km) or more. (The minimum acceptable distance is 200 km.)

Many companies have built and deployed geographically mirrored data centers. The problem is that setting up and maintaining the multiple centers is a manual process that takes a great deal of planning and synchronizing. Even configuration changes must be carefully planned and replicated. One mistake and the redundant site is no longer able to effectively take over in the event of a disaster.

This section identifies the implications for disaster recovery, provides an overview of some of the network implementations today that attempt to address disaster recovery, and describes how Business Continuity Clustering can improve your disaster recovery solution by providing specialized software that automates cluster configuration, maintenance, and synchronization across two to four geographically separate sites.

1.1 Disaster Recovery Implications

The implications of disaster recovery are directly tied to your data. Is your data mission critical? In many instances, critical systems and data drive the business. If these services stop, the business stops. When calculating the cost of downtime, some things to consider are

- File transfers and file storage
- E-mail, calendaring, and collaboration
- Web hosting
- Critical databases
- Productivity
- Reputation

Continuous availability of critical business systems is no longer a luxury, it is a competitive business requirement.The Gartner Group estimates that 40% of enterprises that experience a disaster will go out of business in five years, and only 15% of enterprises have a full-fledged business continuity plan that goes beyond core technology and infrastructure.

The cost to the business for each one hour of service outage includes the following:

- Income loss measured as the income-generating ability of the service, data, or impacted group
- Productivity loss measured as the hourly cost of impacted employees
- Recovery cost measured as the hourly cost of IT personnel to get services back online
- Future lost revenue because of customer and partner perception

1.2 Disaster Recovery Implementations

Stretch clusters and cluster of clusters are two approaches for making shared resources available across geographically distributed sites so that a second site can be called into action after one site fails. To use these approaches, you must first understand how the applications you use and the storage subsystems in your network deployment can determine whether a stretch cluster or cluster of clusters solution is possible for your environment.

- Section 1.2.1, “LAN-Based versus Internet-Based Applications,” on page 12
- Section 1.2.2, “Host-Based versus Storage-Based Data Mirroring,” on page 12
- Section 1.2.3, “Stretch Clusters vs. Cluster of Clusters,” on page 13

1.2.1 LAN-Based versus Internet-Based Applications

Traditional LAN applications require a LAN infrastructure that must be replicated at each site, and might require relocation of employees to allow the business to continue. Internet-based applications allow employees to work from any place that offers an Internet connection, including homes and hotels. Moving applications and services to the Internet frees corporations from the restrictions of traditional LAN-based applications.

By using Novell exteNd Director portal services, iChain®, and ZENworks®, all services, applications, and data can be rendered through the Internet, allowing for loss of service at one site but still providing full access to the services and data by virtue of the ubiquity of the Internet. Data and services continue to be available from the other mirrored sites.

1.2.2 Host-Based versus Storage-Based Data Mirroring

For clustering implementations that are deployed in data centers in different geographic locations, the data must be replicated between the storage subsystems at each data center. Data-block replication can be done by host-based mirroring for synchronous replication over short distances up to 10 km. Typically, replication of data blocks between storage systems in the data centers is performed by SAN hardware that allows synchronous mirrors over a greater distance.

For stretch clusters, host-based mirroring is required to provide synchronous mirroring of the SBD (split-brain detector) partition between sites. This means that stretch-cluster solutions are limited to distances of 10 km.

Table 1-1 compares the benefits and limitations of host-based and storage-based mirroring.

Table 1-1 *Comparison of Host-Based and Storage-Based Data Mirroring*

Capability	Host-Based Mirroring	Storage-Based Mirroring
Geographic distance between sites	Up to 10 km	Can be up to and over 300 km. The actual distance is limited only by the SAN hardware and media interconnects for your deployment.
Mirroring the SBD partition	An SBD can be mirrored between two sites.	Yes, if mirroring is supported by the SAN hardware and media interconnects for your deployment.
Synchronous data-block replication of data between sites	Yes	Yes, requires a Fibre Channel SAN or iSCSI SAN.
Failover support	No additional configuration of the hardware is required.	Requires additional configuration of the SAN hardware.
Failure of the site interconnect	LUNs can become primary at both locations (split brain problem).	Clusters continue to function independently. Minimizes the chance of LUNs at both locations becoming primary (split brain problem).
SMI-S compliance	If the storage subsystems are not SMI-S compliant, the storage subsystems must be controllable by scripts running on the nodes of the cluster.	If the storage subsystems are not SMI-S compliant, the storage subsystems must be controllable by scripts running on the nodes of the cluster.

1.2.3 Stretch Clusters vs. Cluster of Clusters

A stretch cluster and a cluster of clusters are two clustering implementations that you can use with Novell Cluster Services™ to achieve your desired level of disaster recovery. This section describes each deployment type, then compares the capabilities of each.

Novell Business Continuity Clustering automates some of the configuration and processes used in a cluster of clusters. For information, see Section 1.3, “Novell Business Continuity Clusters,” on page 19.

- “Stretch Clusters” on page 14
- “Cluster of Clusters” on page 14
- “Comparison of Stretch Clusters and Cluster of Clusters” on page 16
- “Evaluating Disaster Recovery Implementations for Clusters” on page 18

Stretch Clusters

A stretch cluster consists of a single cluster where the nodes are located in two geographically separate data centers. All nodes in the cluster must be in the same Novell eDirectory™ tree, which requires the eDirectory replica ring to span data centers. The IP addresses for nodes and cluster resources in the cluster must share a common IP subnet.

At least one storage system must reside in each data center. The data is replicated between locations by using host-based mirroring or storage-based mirroring. For information about using mirroring solutions for data replication, see Section 1.2.2, “Host-Based versus Storage-Based Data Mirroring,” on page 12. Link latency can occur between nodes at different sites, so the heartbeat tolerance between nodes of the cluster must be increased to allow for the delay.

The split-brain detector (SBD) is mirrored between the sites. Failure of the site interconnect can result in LUNs becoming primary at both locations (split brain problem) if host-based mirroring is used.

In the stretch-cluster architecture shown in Figure 1-1, the data is mirrored between two data centers that are geographically separated. The server nodes in both data centers are part of one cluster, so that if a disaster occurs in one data center, the nodes in the other data center automatically take over.

Figure 1-1 *Stretch Cluster*

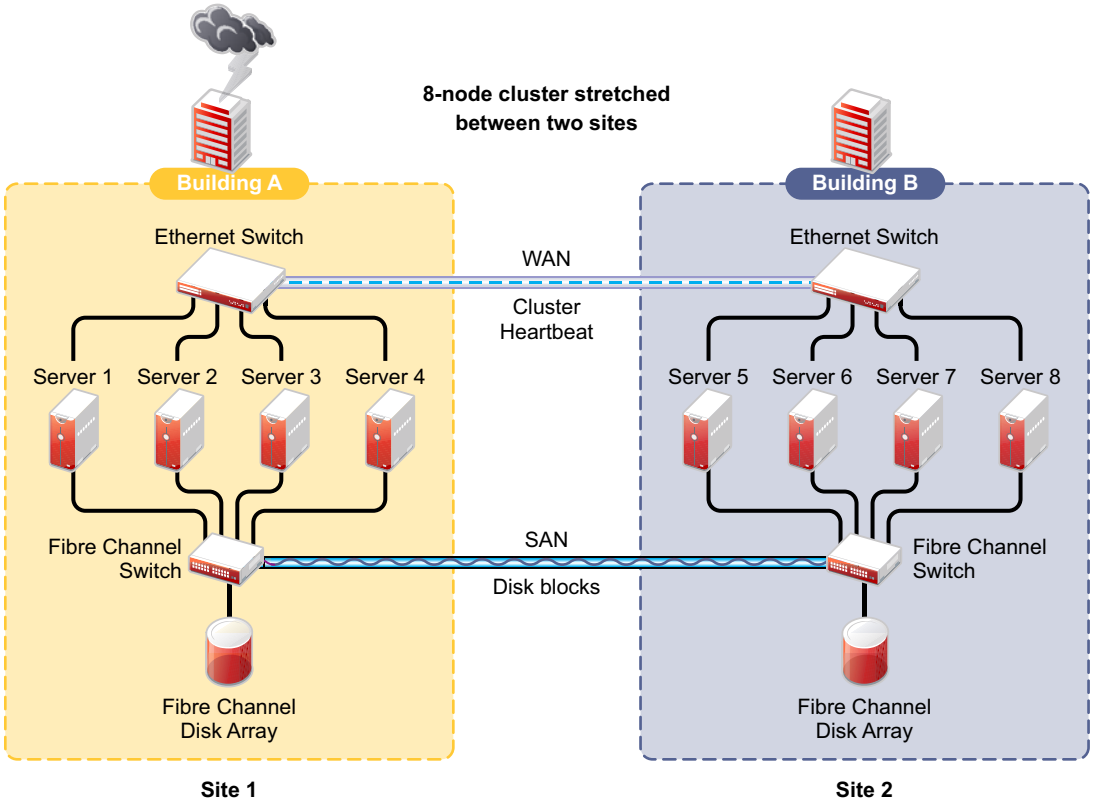

Cluster of Clusters

A cluster of clusters consists of multiple clusters in which each cluster is located in a geographically separate data center. Each cluster can be in different Organizational Unit (OU) containers in the same eDirectory tree, or in different eDirectory trees. Each cluster can be in a different IP subnet.

A cluster of clusters provides the ability to fail over selected cluster resources or all cluster resources from one cluster to another cluster. For example, the cluster resources in one cluster can fail over to separate clusters by using a multiple-site fan-out failover approach. A given service can be provided by multiple clusters. Resource configurations are replicated to each peer cluster and synchronized manually. Failover between clusters requires manual management of the storage systems and the cluster.

Nodes in each cluster access only the storage systems co-located in the same data center. Typically, data is replicated by using storage-based mirroring. Each cluster has its own SBD partition. The SBD partition is not mirrored across the sites, which minimizes the chance for a split-brain problem occurring when using host-based mirroring. For information about using mirroring solutions for data replication, see Section 1.2.2, “Host-Based versus Storage-Based Data Mirroring,” on page 12.

In the cluster-of-clusters architecture shown in Figure 1-2, the data is synchronized by the SAN hardware between two data centers that are geographically separated. If a disaster occurs in one data center, the cluster in the other data center takes over.

Figure 1-2 *Cluster of Clusters*

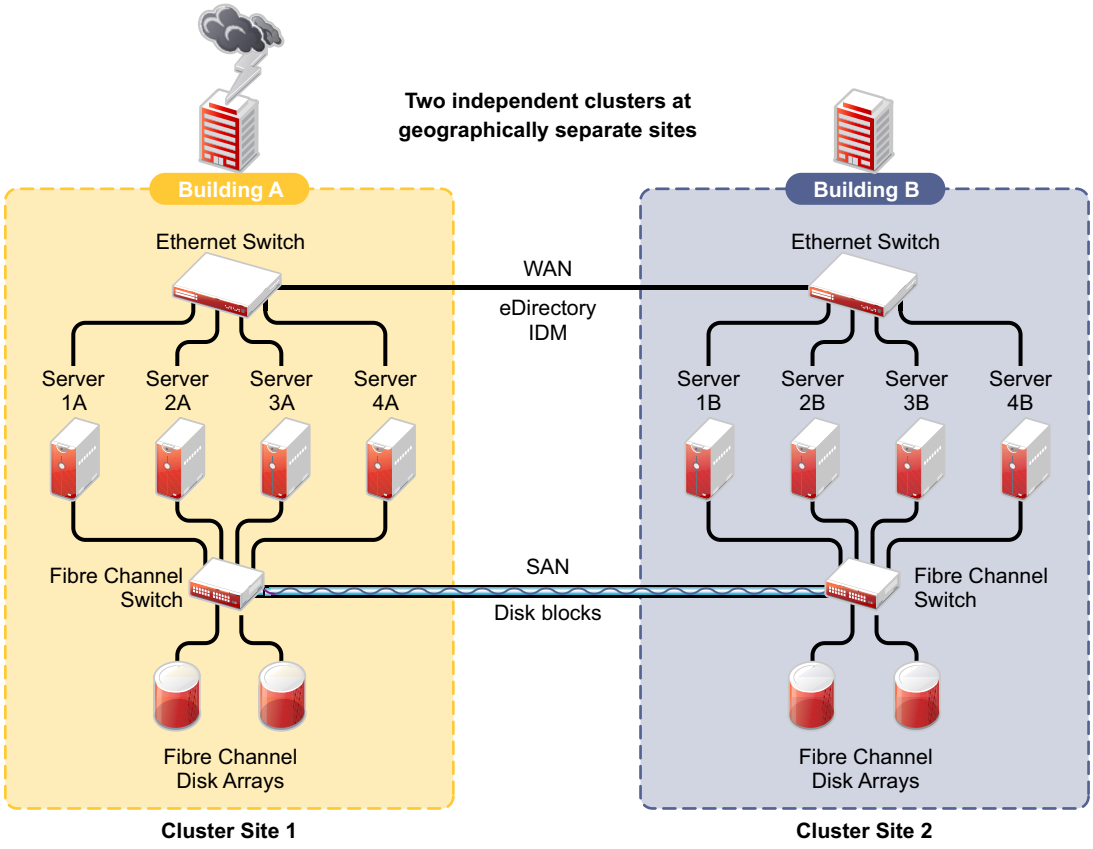

Comparison of Stretch Clusters and Cluster of Clusters

Table 1-2 compares the capabilities of a stretch cluster and a cluster of clusters.

Table 1-2 *Comparison of Stretch Cluster and Cluster of Clusters*

Capability	Stretch Cluster	Cluster of Clusters
Number of clusters	One	Two or more
Number of geographically separated data centers	Two	Two or more
eDirectory trees	Single tree only; requires the replica ring to span data centers.	One or multiple trees
eDirectory Organizational Units (OUs)	Single OU container for all nodes. As a best practice, place the cluster container in an OU separate from the rest of the tree.	Each cluster can be in a different OU. Each cluster is in a single OU container. As a best practice, place each cluster container in an OU separate from the rest of the tree.
IP subnet	IP addresses for nodes and cluster resources must be in a single IP subnet. Because the subnet spans multiple locations, you must ensure that your switches handle gratuitous ARP (Address Resolution Protocol).	IP addresses in a given cluster are in a single IP subnet. Each cluster can use the same or different IP subnet. If you use the same subnet for all clusters in the cluster of clusters, you must ensure that your switches handle gratuitous ARP.
SBD partition	A single SBD is mirrored between two sites by using host-based mirroring, which limits the distance between data centers to 10 km.	Each cluster has its own SBD. Each cluster can have an on-site mirror of its SBD for high availability. If the cluster of clusters uses host-based mirroring, the SBD is not mirrored between sites, which minimizes the chance of LUNs at both locations becoming primary.
Failure of the site interconnect if using host-based mirroring	LUNs might become primary at both locations (split brain problem).	Clusters continue to function independently.
Storage subsystem	Each cluster accesses only the storage subsystem on its own site.	Each cluster accesses only the storage subsystem on its own site.

Capability	Stretch Cluster	Cluster of Clusters
Data-block replication between sites For information about data replication solutions, see Section 1.2.2, “Host-Based versus Storage-Based Data Mirroring,” on page 12.	Yes; typically uses storage-based mirroring, but host-based mirroring is possible for distances up to 10 km.	Yes; typically uses storage-based mirroring, but host-based mirroring is possible for distances up to 10 km.
Clustered services	A single service instance runs in the cluster.	Each cluster can run an instance of the service.
Cluster resource failover	Automatic failover to preferred nodes at the other site.	Manual failover to preferred nodes on one or multiple clusters (multiple-site fan-out failover). Failover requires additional configuration.
Cluster resource configurations	Configured for a single cluster	Configured for the primary cluster that hosts the resource, then the configuration is manually replicated to the peer clusters.
Cluster resource configuration synchronization	Controlled by the master node	Manual process that can be tedious and error-prone.
Failover of cluster resources between clusters	Not applicable	Manual management of the storage systems and the cluster.
Link latency between sites	Can cause false failovers. The cluster heartbeat tolerance between master and slave must be increased to as high as 30 seconds. Monitor cluster heartbeat statistics, then tune down as needed.	Each cluster functions independently in its own geographical site.

Evaluating Disaster Recovery Implementations for Clusters

Table 1-3 illustrates why a cluster of cluster solution is less problematic to deploy than a stretch cluster solution. Manual configuration is not a problem when using Novell Business Continuity Clustering for your cluster of clusters.

Table 1-3 *Advantages and Disadvantages of Stretch Clusters versus Cluster of Clusters*

	Stretch Cluster	Cluster of Clusters
Advantages	◆ It automatically fails over when configured with host-based mirroring. ◆ It is easier to manage than separate clusters. ◆ Cluster resources can fail over to nodes in any site.	◆ eDirectory partitions don't need to span the cluster. ◆ Each cluster can be in different OUs in the same eDirectory tree, or in different eDirectory trees. ◆ IP addresses for each cluster can be on different IP subnets. ◆ Cluster resources can fail over to separate clusters (multiple-site fan-out failover support). ◆ Each cluster has its own SBD. Each cluster can have an on-site mirror of its SBD for high availability. If the cluster of clusters uses host-based mirroring, the SBD is not mirrored between sites, which minimizes the chance of LUNs at both locations becoming primary.
Disadvantages	◆ The eDirectory partition must span the sites. ◆ Failure of site interconnect can result in LUNs becoming primary at both locations (split brain problem) if host-based mirroring is used. ◆ An SBD partition must be mirrored between sites. ◆ It accommodates only two sites. ◆ All IP addresses must reside in the same subnet.	◆ Resource configurations must be manually synchronized. ◆ Storage-based mirroring requires additional configuration steps.

	Stretch Cluster	Cluster of Clusters
Other Considerations	◆ Host-based mirroring is required to mirror the SBD partition between sites. ◆ Link variations can cause false failovers. ◆ You could consider partitioning the eDirectory tree to place the cluster container in a partition separate from the rest of the tree. ◆ The cluster heartbeat tolerance between master and slave must be increased to accommodate link latency between sites. You can set this as high as 30 seconds, monitor cluster heartbeat statistics, and then tune down as needed. ◆ Because all IP addresses in the cluster must be on the same subnet, you must ensure that your switches handle ARP. Contact your switch vendor or consult your switch documentation for more information.	◆ Depending on the platform used, storage arrays must be controllable by scripts that run on NetWare® if the SANs are not SMI-S compliant.

1.3 Novell Business Continuity Clusters

A Novell Business Continuity Clustering cluster is an automated cluster of Novell Cluster Services clusters. It is similar to what is described in “Cluster of Clusters” on page 14, except that the cluster configuration, maintenance, and synchronization have been automated by adding specialized software.

BCC supports up to four peer clusters. The sites are geographically separated mirrored data centers, with a high availability cluster located at each site. Configuration is automatically synchronized between the sites. Data is replicated between sites. All cluster nodes and their cluster resources are monitored at each site. If one site goes down, business continues through the mirrored sites.

The business continuity cluster configuration information is stored in eDirectory. eDirectory schema extensions provide the additional attributes required to maintain the configuration and status information of BCC enabled cluster resources. This includes information about the peer clusters, the cluster resources and their states, and storage control commands.

BCC is an integrated set of tools to automate the setup and maintenance of a business continuity infrastructure. Unlike competitive solutions that attempt to build stretch clusters, BCC uses a cluster of clusters. Each site has its own independent clusters, and the clusters in each of the geographically separate sites are each treated as “nodes” in a larger cluster, allowing a whole site to do fan-out failover to other multiple sites. Although this can currently be done manually with a cluster of clusters, BCC automates the system by using eDirectory and policy-based management of the resources and storage systems.

Novell Business Continuity Clustering software provides the following advantages over typical cluster-of-clusters solutions:

- Supports up to four clusters with up to 32 nodes each.
- Integrates with shard storage hardware devices to automate the failover process through standards-based mechanisms such as SMI-S.
- Uses Novell Identity Manager technology to automatically synchronize and transfer cluster-related eDirectory objects from one cluster to another, and between trees.
- Provides the capability to fail over as few as one cluster resource, or as many as all cluster resources.
- Includes intelligent failover that allows you to perform site failover testing as a standard practice.
- Provides scripting capability that allows enhanced storage management control and customization of migration and fail over between clusters.
- Provides simplified business continuity cluster configuration and management by using the browser-based Novell iManager management tool. iManager is used for the configuration and monitoring of the overall system and for the individual resources.

1.4 BCC Deployment Scenarios

There are several Business Continuity Clustering deployment scenarios that can be used to achieve the desired level of disaster recovery. Three possible scenarios include:

- Section 1.4.1, “Two-Site Business Continuity Cluster Solution,” on page 20
- Section 1.4.2, “Multiple-Site Business Continuity Cluster Solution,” on page 21
- Section 1.4.3, “Low-Cost Business Continuity Cluster Solution,” on page 22

1.4.1 Two-Site Business Continuity Cluster Solution

The two-site business continuity cluster deploys two independent clusters at geographically separate sites. Each cluster can support up to 32 nodes. The clusters can be designed in one of two ways:

- **Active Site/Active Site:** Two active sites where each cluster supports different applications and services. Either site can take over for the other site at any time.
- **Active Site/Passive Site:** A primary site in which all services are normally active, and a secondary site which is effectively idle. The data is mirrored to the secondary site, and the applications and services are ready to load if needed.

The active/active deployment option is typically used in a company that has more than one large site of operations. The active/passive deployment option is typically used when the purpose of the secondary site is primarily testing by the IT department. Replication of data blocks is typically done by SAN hardware, but it can be done by host-based mirroring for synchronous replication over short distances up to 10 km.

Figure 1-3 shows a two-site business continuity cluster that uses storage-based data replication between the sites. BCC uses eDirectory and Identity Manager to synchronize cluster information between the two clusters.

Figure 1-3 *Two-Site Business Continuity Cluster*

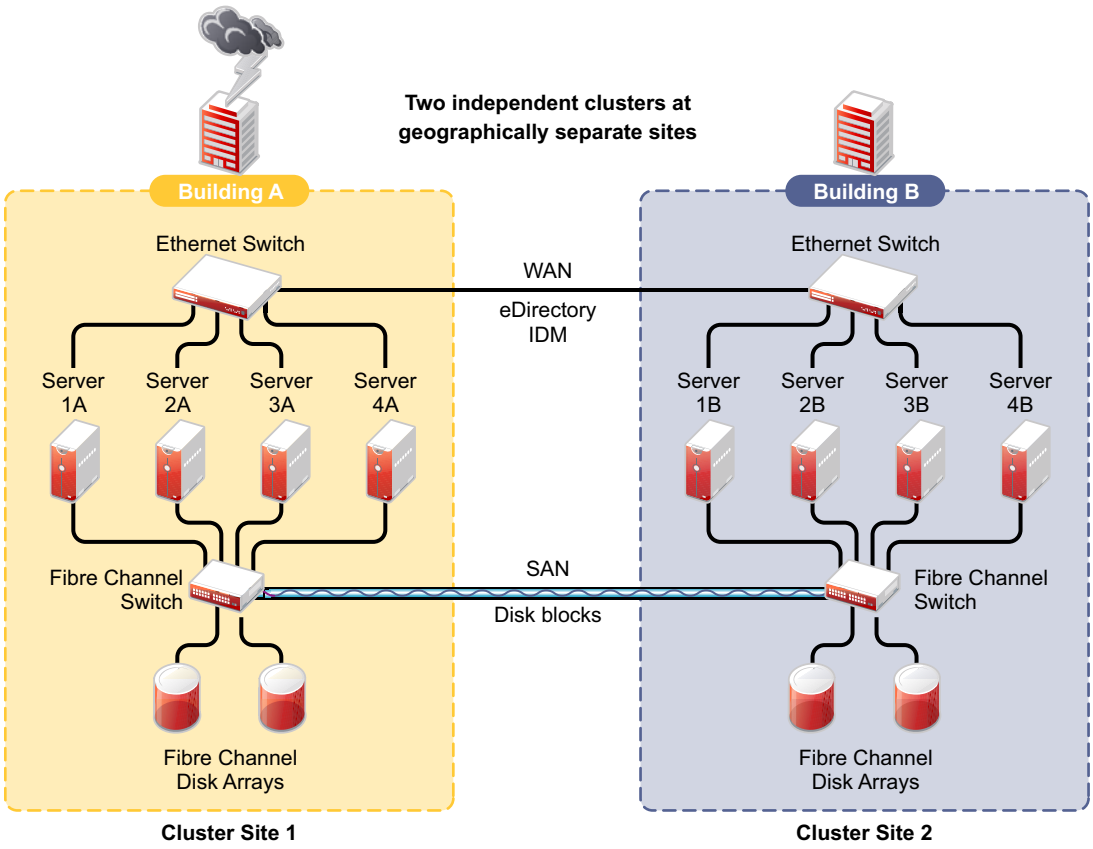

1.4.2 Multiple-Site Business Continuity Cluster Solution

The multiple-site business continuity cluster is a large solution capable of supporting up to four sites. Each cluster can support up to 32 nodes. Services and applications can do fan-out failover between sites. Replication of data blocks is typically done by SAN hardware, but it can be done by host-based mirroring for synchronous replication over short distances up to 10 km.

Figure 1-4 depicts a four-site business continuity cluster that uses storage-based data replication between the sites. BCC uses eDirectory and Identity Manager to synchronize cluster information between the two clusters.

Figure 1-4 *Multiple-Site Business Continuity Cluster*

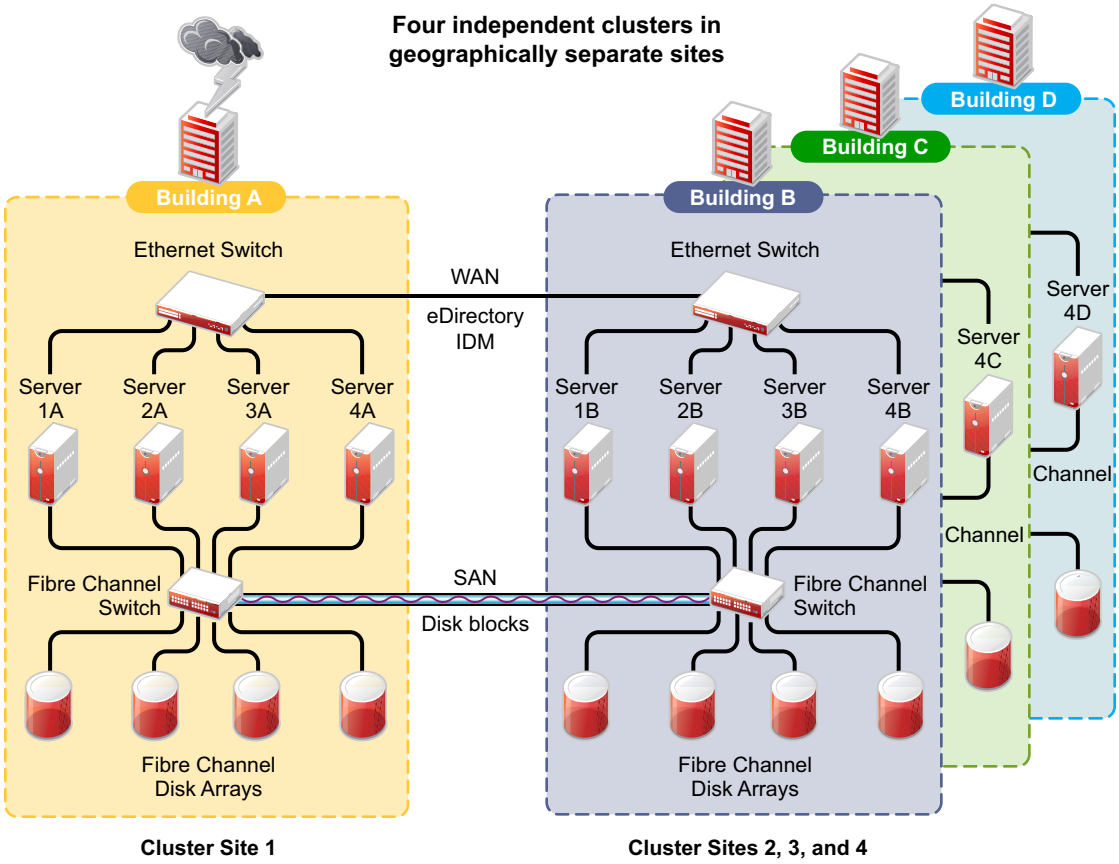

Using additional products, all services, applications, and data can be rendered through the Internet, allowing for loss of service at one site but still providing full access to the services and data by virtue of the ubiquity of the Internet. Data and services continue to be available from the other mirrored sites. Moving applications and services to the Internet frees corporations from the restrictions of traditional LAN-based applications. Traditional LAN applications require a LAN infrastructure that must be replicated at each site, and might require relocation of employees to allow the business to continue. Internet-based applications allow employees to work from any place that offers an Internet connection, including homes and hotels.

1.4.3 Low-Cost Business Continuity Cluster Solution

The low-cost business continuity cluster solution is similar to the previous two solutions, but replaces Fibre Channel storage arrays with iSCSI storage arrays. Data block mirroring can be accomplished either with iSCSI-based block replication, or host-based mirroring. In either case, snapshot technology can allow for asynchronous replication over long distances. However, the low-cost solution does not necessarily have the performance associated with higher-end Fibre Channel storage arrays.

1.5 Key Concepts

The key concepts in this section can help you understand how Business Continuity Clustering manages your business continuity cluster.

1.5.1 Business Continuity Clusters

A cluster of two to four Novell Cluster Services clusters that are managed together by Business Continuity Clustering software. All nodes in every peer cluster are running the same operating system.

1.5.2 Cluster Resources

A cluster resource is a cluster-enabled shared disk that is configured for Novell Cluster Services. It is also BCC-enabled so that it can be migrated and failed over between nodes in different peer clusters.

1.5.3 Landing Zone

The landing zone is an eDirectory context in which the objects for the Virtual Server, the Cluster Pool, and the Cluster Volume are placed when they are created for the peer clusters. You specify the landing zone context when you configure the Identity Manager drivers for the business continuity cluster.

1.5.4 BCC Drivers for Identity Manager

Business Continuity Clustering requires a special Identity Manager driver that uses an Identity Vault to synchronize the cluster resource configuration information between the peer clusters. If the peer clusters are in different eDirectory trees, an additional BCC driver helps synchronize user information between the trees. For information, see Section 6.1.1, “Configuring the Identity Manager Drivers and Templates,” on page 66.

What's New for BCC 1.1 for NetWare

2

This section describes the enhancements made to Novell® Business Continuity Clustering 1.1 for NetWare®.

2.1 What's New for BCC 1.1 SP2

Business Continuity Clustering 1.1 SP2 for NetWare SP8 provides the following enhancements and changes:

- Support for NetWare 6.5 SP8
- Support for Novell Cluster Services™ 1.8.5 for NetWare
- Support for Identity Manager 3.5.1
- Support for 64-bit architectures
- Support for Novell eDirectory™ 8.8
- Support for Novell iManager 2.7.2

2.2 What's New for BCC 1.1 SP1

Business Continuity Clustering 1.1 SP1 for NetWare 6.5 SP6 provides the following enhancements and changes:

- Support for NetWare 6.5 SP6
- Support for Identity Manager 3.*x*

2.3 What's New for BCC 1.1

Business Continuity Clustering 1.1 for NetWare 6.5 SP5 provides the following enhancements and changes as compared to BCC 1.0 for NetWare:

- Support for NetWare 6.5 SP5
- Support for Identity Manager 2.*x*
- Changed inter-cluster communication from the NCP™ (NetWare Control Protocl) port 524 to the CIM ports 5988 and 5989
- Storage Management Initiative (SMI-S) CIM support
  - Standards-based management of the SAN for automatic LUN failover
  - Support for most SANs (such as Xiotech*, EMC*, HP*, IBM*, and so on)

- Automatic failover
  - No need for administrator intervention
  - Based on a configurable minimum number of nodes or a percentage of nodes
  - Extensible monitoring framework
  - Disabled by default

Planning a Business Continuity Cluster

3

Use the guidelines in this section to design your Novell® Business Continuity Clustering solution. The success of your business continuity cluster depends on the stability and robustness of the individual peer clusters. BCC cannot overcome weaknesses in a poorly designed cluster environment.

3.1 Determining Design Criteria

The design goal for your business continuity cluster is to ensure that your critical data and services can continue in the event of a disaster. Design the infrastructure based on your business needs.

Determine your design criteria by asking and answering the following questions:

- ❑ What are the key services that drive your business?
- ❑ Where are your major business sites, and how many are there?
- ❑ What services are essential for business continuance?
- ❑ What is the cost of down time for the essential services?
- ❑ Based on their mission-critical nature and cost of down time, what services are the highest priority for business continuance?
- ❑ Where are the highest-priority services currently located?
- ❑ Where should the highest-priority services be located for business continuance?
- ❑ What data must be replicated to support the highest-priority services?
- ❑ How much data is involved, and how important is it?

3.2 Best Practices

The following practices help you avoid potential problems with your BCC:

- IP address changes should always be made on the Protocols page of the iManager cluster plug-in, not in load and unload scripts.

  This is the only way to change the IP address on the virtual NCP™ server object in eDirectory™.

- Ensure that eDirectory and your clusters are stable before implementing BCC.
- Engage Novell Consulting.
- Engage a consulting group from your SAN vendor.
- The cluster node that hosts the Identity Manager driver should have a full read/write eDirectory™ replica with the following containers in the replica:
  - Driver set container
  - Cluster container
  - (Parent) container where the servers reside
  - Landing zone container
  - User object container
- Ensure that you have full read/write replicas of the entire tree at each data center.

3.3 LAN Connectivity Guidelines

The primary objective of LAN connectivity in a cluster is to provide uninterrupted heartbeat communications. Use the guidelines in this section to design the LAN connectivity for each of the peer clusters in the business continuity cluster:

- Section 3.3.1, “VLAN,” on page 28
- Section 3.3.2, “NIC Teaming,” on page 28
- Section 3.3.3, “IP Addresses,” on page 29
- Section 3.3.4, “Name Resolution,” on page 29
- Section 3.3.5, “IP Addresses for BCC-Enabled Cluster Resources,” on page 29

3.3.1 VLAN

Use a dedicated VLAN (virtual local area network) for each cluster.

The cluster protocol is non-routable, so you cannot direct communications to specific IP addresses. Using a VLAN for the cluster nodes provides a protected environment for the heartbeat process and ensures that heartbeat packets are exchanged only between the nodes of a given cluster.

When using a VLAN, no foreign host can interfere with the heartbeat. For example, it avoids broadcast storms that slow traffic and result in false split-brain abends.

3.3.2 NIC Teaming

Use NIC teaming for adapters for LAN fault tolerance. NIC teaming combines Ethernet interfaces on a host computer for redundancy or increased throughput. It helps increase the availability of an individual cluster node, which helps avoid or reduce the occurrences of failover caused by slow LAN traffic.

When configuring Spanning Tree Protocol (STP), ensure that Portfast is enabled, or consider Rapid Spanning Tree. The default settings for STP inhibit the heartbeat for over 30 seconds whenever there is a change in link status. Test your STP configuration with Novell Cluster Services™ running to make sure that a node is not cast out of the cluster when a broken link is restored.

Consider connecting cluster nodes to access switches for fault tolerance.

3.3.3 IP Addresses

Use a dedicated IP address range for each cluster. You need a unique static IP address for each of the following components of each peer cluster:

- Cluster (master IP address)
- Cluster nodes
- Cluster resources that are not BCC-enabled (file system resources and service resources such as DHCP, DNS, SLP, FTP, and so on)
- Cluster resources that are BCC-enabled (file system resources and service resources such as DHCP, DNS, SLP, FTP, and so on)

Plan your IP address assignment so that it is consistently applied across all peer clusters. Provide an IP address range with sufficient addresses for each cluster.

3.3.4 Name Resolution

In BCC 1.1 and later, the master IP addresses are stored in the NCS:BCC Peers attribute. Ensure that SLP is properly configured for name resolution.

3.3.5 IP Addresses for BCC-Enabled Cluster Resources

Use dedicated IP address ranges for BCC-enabled cluster resources. With careful planning, the IP address and the name of the virtual server for the cluster resource never need to change.

The IP address of an inbound cluster resource is transformed to use an IP address in the same subnet of the peer cluster where it is being cluster migrated. You define the transformation rules to accomplish this by using the Identity Manager driver's search and replace functionality. The transformation rules are easier to define and remember when you use strict IP address assignment, such as using the third octet to identify the subnet of the peer cluster.

3.4 SAN Connectivity Guidelines

The primary objective of SAN (storage area network) connectivity in a cluster is to provide solid and stable connectivity between cluster nodes and the storage system. Before installing Novell Cluster Services and Novell Business Continuity Clustering, make sure the SAN configuration is established and verified.

Use the guidelines in this section to design the SAN connectivity for each of the peer clusters in the business continuity cluster:

- Use host-based multipath I/O management.
- Use redundant SAN connections to provide fault-tolerant connectivity between the cluster nodes and the shared storage devices.
- Connect each node via two fabrics to the storage environment.

- Use a minimum of two mirror connections between storage environments over different fabrics and wide area networks.
- Make sure the distance between storage subsystems is within the limitations of the fabric used given the amount of data, how the data is mirrored, and how long applications can wait for acknowledgement. Also make sure to consider support for asynchronous versus synchronous connections.

3.5 Storage Design Guidelines

Use the guidelines in this section to design the shared storage solution for each of the peer clusters in the business continuity cluster.

- Use a LUN device as the failover unit for each BCC-enabled cluster resource. Multiple pools per LUN are possible, but are not recommended. A LUN cannot be concurrently accessed by servers belonging to different clusters. This means that all resources on a given LUN can be active in a given cluster at any given time. For maximum flexibility, we recommend that you create only one cluster resource per LUN.
- If you use multiple LUNs for a given shared NSS pool, all LUNs must fail over together. We recommend that you use only one LUN per pool, and only one pool per LUN.
- Data must be mirrored between data centers by using host-based mirroring or storage-based mirroring. Storage-based mirroring is recommended.
- When using host-based mirroring, make sure that the mirrored partitions are accessible for the nodes of only one of the BCC peer clusters at any given time. If you use multiple LUNs for a given pool, each segment must be mirrored individually. In large environments, it might be difficult to determine the mirror state of all mirrored partitions at one time. You must also make sure that all segments of the resource fail over together.

3.6 eDirectory Design Guidelines

Your Novell eDirectory solution for each of the peer clusters in the business continuity cluster must consider the following configuration elements. Make sure your approach is consistent across all peer clusters.

3.6.1 Object Location

Cluster nodes and Cluster objects can exist anywhere in the eDirectory tree. The virtual server object, cluster pool object, and cluster volume object are automatically created in the eDirectory context of the server where the cluster resource is created and cluster-enabled. You should create cluster resources on the master node of the cluster.

3.6.2 Cluster Context

Place each cluster in a separate Organizational Unit (OU). All server objects and cluster objects for a given cluster should be in the same OU.

Figure 3-1 *Cluster Resources in Separate OUs*

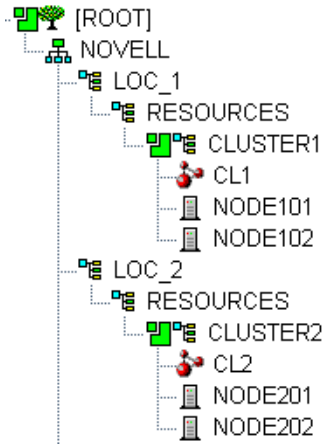

3.6.3 Partitioning and Replication

Partition the cluster OU and replicate it to dedicated eDirectory servers holding a replica of the parent partition and to all cluster nodes. This helps prevent resources from being stuck in an NDS® Sync state when a cluster resource's configuration is modified.

3.6.4 Objects Created by the BCC Drivers for Identity Manager

When a resource is BCC-enabled, its configuration is automatically synchronized with every peer cluster in the business continuity cluster by using customized Identity Manager drivers. The following eDirectory objects are created in each peer cluster:

- Cluster Resource object
- Virtual Server object
- Cluster Pool object
- Cluster Volume object

The Cluster Resource object is placed in the Cluster object of the peer clusters where the resource did not exist initially. The Virtual Server, Cluster Pool, and Cluster Volume objects are stored in the landing zone. Search-and-replace transform rules define cluster-specific modifications such as the IP address.

3.6.5 Landing Zone

Any OU can be defined as the BCC landing zone. Use a separate OU for the landing zone than you use for a cluster OU. The cluster OU for one peer cluster can be the landing zone OU for a different peer cluster.

3.6.6 Naming Conventions for BCC-Enabled Resources

Develop a cluster-independent naming convention for BCC-enabled cluster resources. It can become confusing if the cluster resource name refers to one cluster and is failed over to a peer cluster.

You can use a naming convention for resources in your BCC as you create those resources. Changing existing names of cluster resources is less straightforward and can be error prone.

For example, when cluster-enabling NSS pools the default naming conventions used by NSS are:

Cluster Resource: `poolname_SERVER`
Cluster-Enabled Pool: *`clustername_poolname`*`_POOL`
Cluster-Enabled Volume: *`clustername_volumename`*
Virtual Server: *`clustername_poolname`*`_SERVER`

Instead, use names that are independent of the clusters and that are unique across all peer clusters. For example, replace the *clustername* with something static such as BCC.

Cluster Resource: `poolname_SERVER`
Cluster-Enabled Pool: `BCC_`*`poolname`*`_POOL`
Cluster-Enabled Volume: `BCC_`*`volumename`*
Virtual Server: `BCC_`*`poolname`*`_SERVER`

Resources have an identity in each peer cluster, and the names are the same in each peer cluster. For example, Figure 3-2 shows the cluster resource identity in each of two peer clusters.

Figure 3-2 *Cluster Resource Identity in Two Clusters*

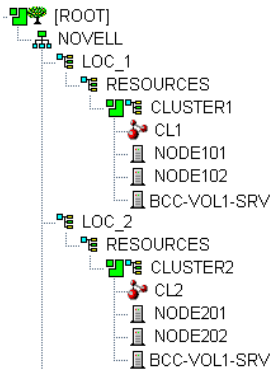

3.7 Cluster Design Guidelines

Your Novell Cluster Services solution for each of the peer clusters in the business continuity cluster must consider the following configuration guidelines. Make sure your approach is consistent across all peer clusters.

- IP address assignments should be consistently applied within each peer cluster and for all cluster resources.
- Ensure that IP addresses are unique across all BCC peer clusters.

- Volume IDs must be unique across all peer clusters. Each cluster node automatically assigns volume ID 0 to volume SYS and volume ID 1 to volume _ADMIN. Cluster-enabled volumes use high volume IDs, starting from 254 in decending order. Novell Client uses the volume ID to access a volume.

  When existing clusters are configured and enabled within the same business continuity cluster, the volume IDs for the existing shared volumes might also share the same volume IDs. To resolve this conflict, manually edit the load script for each volume that has been enabled for business continuity and change the volume IDs to unique values for each volume in the business continuity cluster.
- BCC configuration should consider the configuration requirements for each of the services supported across all peer clusters.
- Create failover matrixes for each cluster resource so that you know what service is supported and which nodes are the preferred nodes for failover within the same cluster and among the peer clusters.

Installing Business Continuity Clustering

4

This section describes how to install, set up, and configure Novell® Business Continuity Clustering 1.1 SP2 for NetWare® 6.5 SP8 for your specific needs:

4.1 Requirements for BCC 1.1 SP2 for NetWare

The requirements in this section must be met prior to installing Novell Business Continuity Clustering software.

4.1.1 Business Continuity Clustering Licensing

Novell Business Continuity Clustering software requires a license agreement for each business continuity cluster. For purchasing information, see Novell Business Continuity Clustering (http://www.novell.com/products/businesscontinuity/howtobuy.html).

4.1.2 Business Continuity Cluster Component Locations

Figure 4-1 illustrates where the various components needed for a business continuity cluster are installed.

Figure 4-1 *Business Continuity Cluster Component Locations*

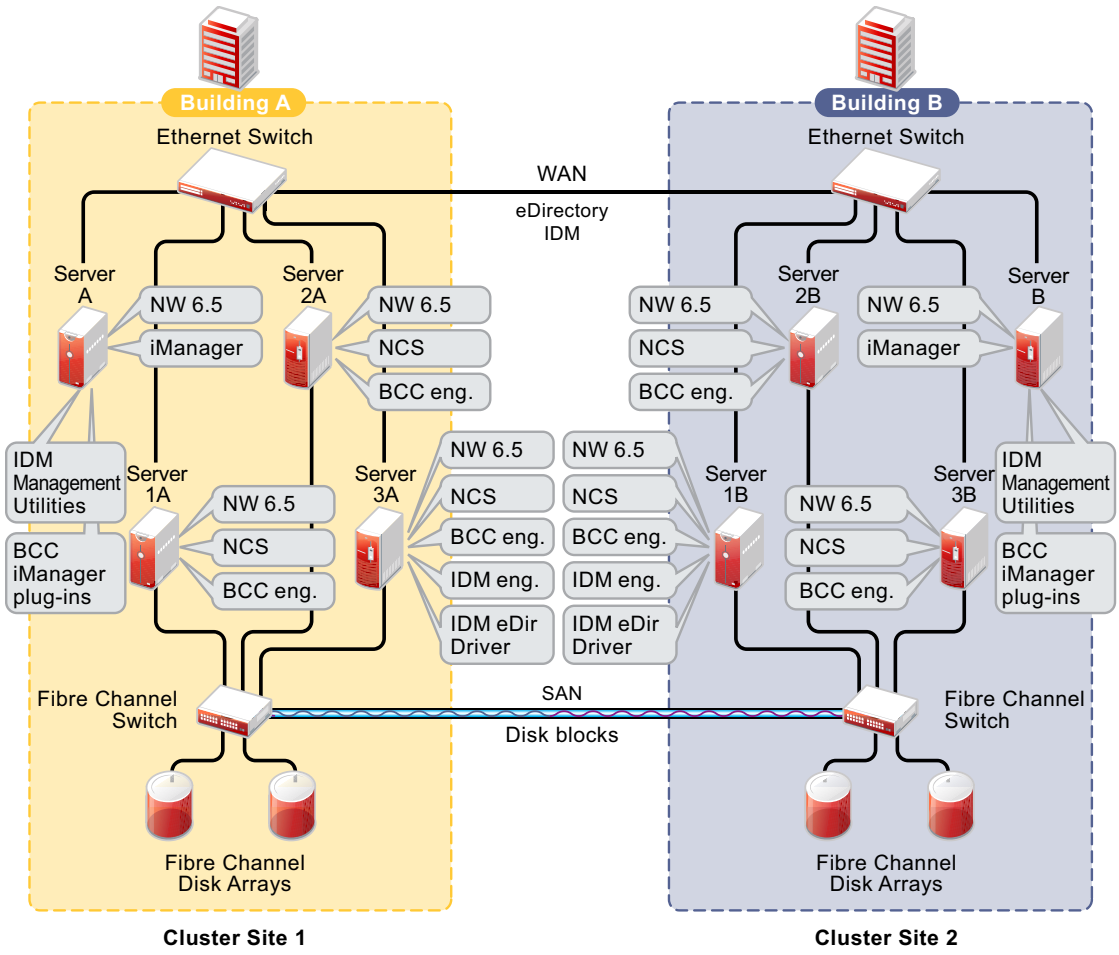

Figure 4-1 uses the following abbreviations:

BCC: Novell Business Continuity Clustering 1.1 SP2 for NetWare
NCS: Novell Cluster Services 1.8.5 for NetWare
IDM:Identity Manager 3.5.1 Bundle Edition
eDir: Novell eDirectory 8.8
NW 6.5: NetWare 6.5 SP8

4.1.3 NetWare 6.5 SP8

NetWare® 6.5 Support Pack 8 must be installed and running on every node in each peer cluster that will be part of the business continuity cluster.

See the *NW65 SP8: Installation Guide* for information on installing and configuring NetWare 6.5 SP8.

4.1.4 Novell Cluster Services 1.8.5 for NetWare

You need two to four clusters with Novell Cluster Services™ 1.8.5 (the version that ships with NetWare 6.5 SP8 installed and running on each node in the cluster.

See the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide* for information on installing, configuring, and managing Novell Cluster Services.

Consider the following when preparing your clusters for the business continuity cluster:

- “Cluster Names” on page 37
- “Storage” on page 37
- “eDirectory” on page 37
- “Peer Cluster Credentials” on page 38

Cluster Names

Each cluster must have a unique name, even if the clusters reside in different Novell eDirectory™ trees. Clusters must not have the same name as any of the eDirectory trees in the business continuity cluster.

Storage

The storage requirements for Novell Business Continuity Clustering software are the same as for Novell Cluster Services. For more information, see the following in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*:

- “Hardware Requirements”
- “Shared Disk System Requirements”
- “Rules for Using Disks in a Shared Storage Space”

Some storage vendors require you to purchase or license their CLI (Command Line Interface) separately. The CLI for the storage system might not initially be included with your hardware.

Also, some storage hardware may not be SMI-S compliant and cannot be managed by using SMI-S commands.

eDirectory

The recommended configuration is to have each cluster in the same eDirectory tree but in different OUs (Organizational Units). This guide focuses on the single-tree setup.

BCC 1.1 SP2 for NetWare also supports a business continuity cluster with clusters in two eDirectory trees. See Appendix B, “Implementing a Multiple-Tree BCC,” on page 125 for more information.

Peer Cluster Credentials

To add or change peer cluster credentials, you must access iManager on a server that is in the same eDirectory tree as the cluster where you are adding or changing peer credentials.

4.1.5 Novell eDirectory 8.8

Novell eDirectory 8.8 is supported with Business Continuity Clustering 1.1 Support Pack 1. See the eDirectory 8.8 documentation (http://www.novell.com/documentation/edir88/) for more information.

Rights Need for Installing BCC

The first time that you install the Business Continuity Clustering engine software in an eDirectory tree, the eDirectory schema is automatically extended with BCC objects.

IMPORTANT: The user who installs BCC must have the eDirectory credentials necessary to extend the schema.

If the eDirectory administrator username or password contains special characters (such as $, #, and so on), make sure to escape each special character by preceding it with a backslash (\) when you enter credentials.

Rights Needed for Individual Cluster Management

The BCC Administrator user is not automatically assigned the rights necessary to manage all aspects of each peer cluster. When managing individual clusters, you must log in as the Cluster Administrator user. You can manually assign the Cluster Administrator rights to the BCC Administrator user for each of the peer clusters if you want the BCC Administrator user to have all rights.

Rights Needed for BCC Management

Before you install BCC, create the BCC Administrator user and group identities in eDirectory to use when you manage the BCC. For information, see Section 4.3, “Configuring a BCC Administrator User,” on page 45.

Rights Needed for Identity Manager

The node where Identity Manager is installed must have an eDirectory full replica with at least read/write access to all eDirectory objects that will be synchronized between clusters. You can also have the eDirectory master running on the node instead of the replica.

The replica does not need to contain all eDirectory objects in the tree. The eDirectory full replica must have at least read/write access to the following containers in order for the cluster resource synchronization and user object synchronization to work properly:

- The Identity Manager driver set container.
- The container where the Cluster object resides.
- The container where the Server objects reside.

  If Server objects reside in multiple containers, this must be a container high enough in the tree to be above all containers that contain Server objects.

The best practice is to have all Server objects in one container.

- The container where the cluster Pool objects and Volume objects are placed when they are synchronized to this cluster. This container is referred to as the *landing zone*. The NCP™ Server objects for the virtual server of a BCC-enabled resource are also placed in the landing zone.
- The container where the User objects reside that need to be synchronized. Typically, the User objects container is in the same partition as the cluster objects.

IMPORTANT: Full eDirectory replicas are required. Filtered eDirectory replicas are not supported with this version of Business Continuity Clustering software.

4.1.6 SLP

You must have SLP (Server Location Protocol) set up and configured properly on each server node in every cluster. Typically, SLP is installed as part of the eDirectory installation and setup when you install the server operating system for the server. For information, see "Implementing the Service Location Protocol" (http://www.novell.com/documentation/edir88/edir88/data/ba5lb4b.html) in the *Novell eDirectory 8.8 Administration Guide*.

4.1.7 Identity Manager 3.5.1 Bundle Edition

The Identity Manager 3.5.1 Bundle Edition is required for synchronizing the configuration of the peer clusters in your business continuity cluster. It is not involved in other BCC management operations such as migrating cluster resources within or across peer clussters.

Before you install Business Continuity Clustering on the cluster nodes, make sure that Identity Manager and the Identity Manager driver for eDirectory are installed on one node in each peer cluster that you want to be part of the business continuity cluster.

The same Identity Manager installation program that is used to install the Identity Manager engine is also used to install the Identity Manager eDirectory driver and management utilities. See "Business Continuity Cluster Component Locations" on page 36 for information on where to install Identity Manager components.

Downloading the Bundle Edition

The bundle edition is a limited release of Novell Identity Manager 3.5.1 for NetWare 6.5 SP8 that allows you to use the Identity Manager software, the eDirectory driver, and the Identity Manager management tools for Novell iManager 2.7.2. BCC driver templates are applied to the eDirectory driver to create BCC-specific drivers that automatically synchronize BCC configuration information between the Identity Manager nodes in peer clusters. To download the Bundle Edition, go to to the Identity Manager 3.5.1 Bundle Edition download site (http://download.novell.com/Download?buildid=hEOxV3rys2M~).

Credential for Drivers

The Bundle Edition requires a credential that allows you to use drivers beyond an evaluation period. The credential can be found in the BCC license. In the Identity Manager interface in iManager, enter the credential for each driver that you create for BCC. You must also enter the credential for the matching driver that is installed in a peer cluster. You can enter the credential, or put the credential in a file that you point to.

Identity Manager Engine and eDirectory Driver

BCC requires Identity Manager 3.5.1 or later to run on one node in each of the clusters that belong to the business continuity cluster. (Identity Manager was formerly called DirXML®.) Identity Manager should not be set up as clustered resource.

For information about installing and configuring Identity Manager 3.5.1, see the Identity Manager 3.5.1 documentation Web site (http://www.novell.com/documentation/idm35/).

The node where the Identity Manager engine and the eDirectory driver are installed must have an eDirectory full replica with at least read/write access to all eDirectory objects that will be synchronized between clusters. This does not apply to all eDirectory objects in the tree. For information about the eDirectory full replica requirements, see Section 4.1.5, “Novell eDirectory 8.8,” on page 38.

Identity Manager Driver for eDirectory

On the same node where you install the Identity Manager engine, install the following:

- **Single Tree:** One instance of the Identity Manager driver for eDirectory.
- **Two Trees:** Two instances of the Identity Manager driver for eDirectory. The eDirectory driver must be installed twice on each node—once for Tree A and once for Tree B.

For information about installing the Identity Manager driver for eDirectory, see *Identity Manager 3.5.1 Driver for eDirectory: Implementation Guide* (http://www.novell.com/documentation/idm35drivers/edirectory/data/bktitle.html).

Identity Manager Management Utilities

The Identity Manager management utilities must be installed on the same server as Novell iManager. The Identity Manager utilities and iManager can be installed on a cluster node, but installing them on a non-cluster node is the recommended configuration. For information about iManager requirements for BCC, see Section 4.1.8, “Novell iManager 2.7.2,” on page 41.

IMPORTANT: Identity Manager plug-ins for iManager require that eDirectory is running and working properly in the tree. If the plug-in does not appear in iManager, make sure that the eDirectory daemon (`ndsd`) is running on the server that contains the eDirectory master replica.

To restart `ndsd` on the master replica server, enter the following command at its terminal console prompt as the `root` user:

```
rcndsd restart
```

4.1.8 Novell iManager 2.7.2

Novell iManager 2.7.2 (the version released with NetWare 6.5 SP8, or later version) must be installed and running on a server in the eDirectory tree where you are installing Business Continuity Clustering software. You need to install the BCC plug-in, the Clusters plug-in, and the Storage Management plug-in in order to manage the BCC in iManager. As part of the install process, you must also install plug-ins for the Identity Manager role that are management templates for configuring a business continuity cluster.

For information about installing and using iManager, see the Novell iManager 2.7 documentation Web site (http://www.novell.com/documentation/imanager27/).

The Identity Manager management utilities must be installed on the same server as iManager. You can install iManager and the Identity Manager utilities on a cluster node, but installing them on a non-cluster node is the recommended configuration. For information about Identity Manager requirements for BCC, see Section 4.1.7, “Identity Manager 3.5.1 Bundle Edition,” on page 39.

See “Business Continuity Cluster Component Locations” on page 36 for specific information on where to install Identity Manager components.

4.1.9 Storage-Related Plug-Ins for iManager 2.7.2

The Clusters plug-in (ncsmgmt.npm) has been updated from the release in NetWare 6.5 SP8 to provide support for this release of Business Continuity Clustering. You must install the Clusters plug-in and the Storage Management plug-in (`storagemgmt.npm`). Other storage-related plug-ins are Novell Storage Services™ (NSS) (`nssmgmt.npm`), Novell AFP (`afpmgmt.npm`), Novell CIFS (`cifsmgmt.npm`), Novell Distributed File Services (`dfsmgmt.npm`), and Novell Archive and Version Services (`avmgmt.npm`). NSS is required in order to use shared NSS pools as cluster resources. The other services are optional.

IMPORTANT: The Storage Management plug-in module (`storagemgmt.npm`) contains common code required by all of the other storage-related plug-ins. Make sure that you include `storagemgmt.npm` when installing any of the others. If you use more than one of these plug-ins, you should install, update, or remove them all at the same time to make sure the common code works for all plug-ins.

The storage-related plug-ins are available as a zipped download on the Novell Downloads (http://download.novell.com/) Web site.

1. On the iManager server, if the OES 2 version of the storage-related plug-ins are installed, or if you upgraded this server from OES 2 Linux or NetWare 6.5 SP7, log in to iManager, then uninstall all of the storage-related plug-ins that are currently installed, including `storagemgmt.npm`.

   This step is necessary for upgrades only if you did not uninstall and reinstall the storage-related plug-ins as part of the upgrade process.
2. Copy the new .npm files into the iManager plug-ins location, manually overwriting the older version of the plug-in in the packages folder with the newer version of the plug-in.
3. In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.
4. Restart Tomcat by entering the following commands at a system console prompt:

```
tc4stop
tomcat4
```

5 Restart Apache by entering the following command at a system console prompt:

```
ap2webrs
```

4.1.10 Windows Workstation

The Business Continuity Clustering installation program is run from a Windows* workstation. Prior to running the installation program:

- The Windows workstation must have the latest Novell Client™ software installed.
- You must be authenticated to the eDirectory tree where the cluster resides.

4.1.11 OpenWBEM

OpenWBEM must be configured to start in `autoexec.ncf`, and be running when you manage the cluster with Novell iManager. For information on setup and configuration, see the *NW 6.5 SP8: OpenWBEM Services Administration Guide*.

Port 5989 is the default setting for secure HTTP (HTTPS) communications. If you are using a firewall, the port must be opened for CIMOM communications.

Beginning in NetWare 6.5 SP7, the Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a username and password) between iManager and the `_admin` volume on the server that you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

If CIMOM is not currently running when you click *OK* or *Finish* for the task that sends the sensitive information, you get an error message explaining that the connection is not secure and that CIMOM must be running before you can perform the task.

IMPORTANT: If you receive file protocol errors, it might be because WBEM is not running.

To check the status of WBEM:

1 As the Admin user or equivalent, enter the following at the server console:

```
modules owcimomd
```

To start WBEM:

1 As the Admin user or equivalent, enter the following at the server console:

```
openwbem
```

4.1.12 Shared Disk Systems

For Business Continuity Clustering, a shared disk system is required for each peer cluster in the business continuity cluster. See “Shared Disk System Requirements” in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

In addition to the shared disks in an original cluster, you need additional shared disk storage in the other peer clusters to mirror the data between sites as described in Section 4.1.13, “Mirroring Shared Disk Systems Between Peer Clusters,” on page 43.

4.1.13 Mirroring Shared Disk Systems Between Peer Clusters

The Business Continuity Clustering software does not perform data mirroring. You must separately configure either storage-based mirroring or host-based file system mirroring for the shared disks that you want to fail over between peer clusters. Storage-based synchronized mirroring is the preferred solution.

IMPORTANT: Use whatever method is available to implement storage-based mirroring or host-based file system mirroring between the peer clusters for each of the shared disks that you plan to fail over between peer clusters.

For information about how to configure host-based file system mirroring for Novell Storage Services pool resources, see Appendix D, “Configuring Host-Based File System Mirroring for NSS Pools,” on page 135.

For information about storage-based mirroring, consult the vendor for your storage system or see the vendor documentation.

4.1.14 LUN Masking for Shared Devices

LUN masking is the ability to exclusively assign each LUN to one or more host connections. With it, you can assign appropriately sized pieces of storage from a common storage pool to various servers. See your storage system vendor documentation for more information on configuring LUN masking.

When you create a Novell Cluster Services system that uses a shared storage system, it is important to remember that all of the servers that you grant access to the shared device, whether in the cluster or not, have access to all of the volumes on the shared storage space unless you specifically prevent such access. Novell Cluster Services arbitrates access to shared volumes for all cluster nodes, but cannot protect shared volumes from being corrupted by non-cluster servers.

Software included with your storage system can be used to mask LUNs or to provide zoning configuration of the SAN fabric to prevent shared volumes from being corrupted by non-cluster servers.

IMPORTANT: We recommend that you implement LUN masking in your business continuity cluster for data protection. LUN masking is provided by your storage system vendor.

4.1.15 Link Speeds

For real-time mirroring, link speeds should be 1 GB or better, the Fibre Channel cable length between sites should be less than 200 kilometers, and the links should be dedicated.

Many factors should be considered for distances greater than 200 kilometers, some of which include:

- The amount of data being transferred

- The bandwidth of the link
- Whether or not snapshot technology is being used

4.1.16 Ports

If you are using a firewall, the ports must be opened for OpenWBEM and the Identity Manager drivers.

Table 4-1 *Default Ports for the BCC Setup*

Product	Default Port
OpenWBEM	5989
eDirectory driver	8196
Cluster Resources Synchronization driver	2002 (plus the ports for additional instances)
User Object Synchronization driver	2001 (plus the ports for additional instances)

4.1.17 Web Browser

When using iManager, make sure your Web browser settings meet the requirements in this section.

- "Web Browser Language Setting" on page 44
- "Web Browser Character Encoding Setting" on page 44

Web Browser Language Setting

The iManager plug-in might not operate properly if the highest priority Language setting for your Web browser is set to a language other than one of iManager's supported languages. To avoid problems, in your Web browser, click *Tools* > *Options* > *Languages*, then set the first language preference in the list to a supported language.

Refer to the Novell iManager documentation (http://www.novell.com/documentation/imanager27/) for information about supported languages.

Web Browser Character Encoding Setting

Supported language codes are Unicode (UTF-8) compliant. To avoid display problems, make sure the Character Encoding setting for the browser is set to Unicode (UTF-8) or ISO 8859-1 (Western, Western European, West European).

In a Mozilla browser, click *View* > *Character Encoding*, then select the supported character encoding setting.

In an Internet Explorer browser, click *View* > *Encoding*, then select the supported character encoding setting.

4.1.18 BASH

BASH must be installed on all nodes that participate in the business continuity cluster. The BASH shell does not need to be running, only installed.

4.1.19 LIBC

You must have the latest LIBC patch installed. This is currently libcsp6X. See *LIBC Update NetWare 6.5 SP6 9.00.05 (Technical Information Document # 5003460)* (http://support.novell.com/docs/Readmes/InfoDocument/patchbuilder/readme_5003460.html).

4.1.20 autoexec.ncf File

The `sys:\system\autoexec.ncf` file must be modified so that the call to `sys:/bin/unixenv.ncf` is before the calls to `openwbem.ncf` and `ldbcc.ncf`.

4.2 Downloading the Business Continuity Clustering Software

Before you install Novell Business Continuity Clustering, download and copy the software to a directory on your Windows workstation. To download Novell Business Continuity Clustering 1.1 SP2 for NetWare 6.5 SP8, go to the Novell Business Continuity Clustering download site (http://download.novell.com/Download?buildid=bdkmSxRgKVk~).

4.3 Configuring a BCC Administrator User

The BCC Administrator user is a trustee of each of the peer Cluster objects in the business continuity cluster. During the install, you specify an existing user to be the BCC Administrator user. This user should have at least Read and Write rights to the All Attribute Rights property on the Cluster object of the remote cluster. The user should also have rights to the `sys:/tmp` directory.

- Section 4.3.1, “Creating the BCC Administrator User,” on page 45
- Section 4.3.2, “Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects,” on page 46
- Section 4.3.3, “Assigning Trustee Rights for the BCC Administrator User to the _ADMIN Volume,” on page 46
- Section 4.3.4, “Assigning Trustee Rights for the BCC Administrator User to the sys:\tmp Directory,” on page 47

4.3.1 Creating the BCC Administrator User

The BCC Administrator user will be a trustee of each of the peer cluster objects in the business continuity cluster. Identify an existing user, or create a new user, who you want to use as the BCC Administrator user.

4.3.2 Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects

Assign trustee rights to the BCC Administrator user for each cluster that you plan to add to the business continuity cluster.

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the NetWare server or Windows server where you have installed iManager and the Identity Manager preconfigured templates for iManager.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the *Roles and Tasks* column, click *Rights*, then click the *Modify Trustees* link.
4. Specify the Cluster object name, or browse and select it, then click OK.
5. If the BCC Administrator user is not listed as a trustee, click the *Add Trustee* button, browse and select the User object, then click *OK*.
6. Click *Assigned Rights* for the BCC Administrator user, and then ensure the *Read* and *Write* check boxes are selected for the *All Attributes Rights* property.
7. Click *Done* to save your changes.
8. Repeat Step 3 through Step 7 for the other clusters in your business continuity cluster.

4.3.3 Assigning Trustee Rights for the BCC Administrator User to the _ADMIN Volume

You must also ensure that the BCC Administrator user has file system rights to the `_ADMIN:\Novell\Cluster` directory of each of the nodes in your BCC. This is necessary because the `_ADMIN` volume is virtual, and is created each time the server starts. For this reason, you cannot assign eDirectory trustee rights to the `_ADMIN` volume.

To assign BCC Administrator user file system rights to the `_ADMIN:\Novell\Cluster` directory:

1. Open the `sys:\etc\trustrees.xml` file
2. Add a trustee entry for the BCC Administrator user that assigns Read, Write, Modify, and File Scan (RWMF) rights to the `_ADMIN:\Novell\Cluster` directory.
3. Repeat this process on all NetWare nodes that are part or your BCC.

   The trustee entry could be similar to the following entry:

```
<addTrustee>
  <name>BCCAdmin.users.lab.acme_tree</name>
  <fileName>_ADMIN:\Novell\Cluster</fileName>
  <rights>
    <read/>
    <write/>
    <fileScan/>
    <modify/>
   </rights>
</addTrustee>
```

   Note the following items with this example:

   - The `<name>` element is the BCC Administrator user. The tree name is required.

- The <filename> element must be _ADMIN:\Novell\Cluster
- The rights must be RWMF.
- You must add the trustee entry to all the NetWare nodes in your BCC.

The following is an example of a complete trustees.xml file. Note the multiple trustee entries. For this reason you should edit this file and add the BCC entry rather than copy the file from server to server.

```
<specialTrustees>
  <addTrustee>
    <name>BCCAdmin.users.lab.acme_tree</name>
    <fileName>_ADMIN:\Novell\Cluster</fileName>
    <rights>
      <read/>
      <write/>
      <fileScan/>
      <modify/>
    </rights>
  </addTrustee>
  <addTrustee>
    <context/>
    <name>[public]</name>
    <fileName>_admin:manage_nss\files.cmd</fileName>
    <rights>
      <read/>
      <write/>
      <fileScan/>
    </rights>
    <background/>
  </addTrustee>
</specialTrustees>
```

After the trustees.xml file has been modified on all NetWare nodes, the NetWare nodes must be rebooted. This can be done in a rolling fashion. You should start with the node that has the highest IP address first and work down in IP address order. This speeds the rate at which the Novell Cluster Services master node acquires the change.

4.3.4 Assigning Trustee Rights for the BCC Administrator User to the sys:\tmp Directory

You must also ensure that the BCC Administrator user is a trustee with Read, Write, Create, Erase, Modify, and File Scan access rights to the sys:\tmp directory on every node in your NetWare clusters.

IMPORTANT: If you are concerned about denial of service attacks with the BCC Administrator user, you can set a quota of 5 MB for that user. This can prevent the BCC Administrator user from filling the sys: volume by copying an excessive number of files to the sys:\tmp directory.

To assign BCC Administrator user file system rights to the sys:\tmp directory:

1. Open the sys:\etc\trustrees.xml file
2. Add a trustee entry for the BCC Administrator user that assigns Read, Write, Create, Erase, Modify, and File Scan (RWCEMF) rights to the sys:\tmp directory.
3. Repeat this process on all NetWare nodes that are part or your BCC.

The trustee entry could be similar to the following entry:

```
<addTrustee>
  <name>BCCAdmin.users.lab.acme_tree</name>
  <fileName>sys:\tmp</fileName>
  <rights>
    <read/>
    <write/>
    <create/>
    <erase/>
    <fileScan/>
    <modify/>
  </rights>
</addTrustee>
```

Note the following items with this example:

- The `<name>` element is the BCC Administrator user. The tree name is required.
- The `<filename>` element must be `sys:\tmp`
- The rights must be `RWCEMF`.
- You must add the trustee entry to all the NetWare nodes in your BCC.

IMPORTANT: Make sure that you edit each `trustees.xml` file on each cluster node to add the BCC entry rather than copy the file from server to server.

After the `trustees.xml` file has been modified on all NetWare nodes, the NetWare nodes must be rebooted. This can be done in a rolling fashion. You should start with the node that has the highest IP address first and work down in IP address order. This speeds the rate at which the Novell Cluster Services master node acquires the change.

4.4 Installing and Configuring the Novell Business Continuity Clustering Software

It is necessary to run the Novell Business Continuity Clustering installation program when you want to:

- Install the Business Continuity Clustering engine software on cluster nodes for the clusters that will be part of a business continuity cluster.

  The Business Continuity Clustering installation installs to only one cluster at a time. You must run the installation program again for each NetWare cluster that you want to be part of a business continuity cluster.

- Install the BCC-specific Identity Manager templates for iManager snap-ins on either a NetWare 6.5 SP5 or SP6 server or a Windows server.

  The templates add functionality to iManager so you can manage your business continuity cluster. You must have previously installed iManager on the server where you plan to install the templates.

IMPORTANT: Before you begin, make sure your setup meets the requirements specified in Section 4.1, “Requirements for BCC 1.1 SP2 for NetWare,” on page 35. The BCC Administrator user and group must already be configured as specified in Section 4.3, “Configuring a BCC Administrator User,” on page 45.

- Section 4.4.1, “Installing the BCC Engine,” on page 49
- Section 4.4.2, “Installing the Identity Manager Templates,” on page 50

4.4.1 Installing the BCC Engine

You must install the Business Continuity Clustering engine software on each cluster node for the clusters that will be part of a business continuity cluster. You install the software on the nodes of one cluster at a time.

To install and configure Business Continuity Clustering, complete the following steps:

1. From the directory on your Windows workstation where you copied the Business Continuity Clustering software, run `install.exe`.

   For download information, see Section 4.2, “Downloading the Business Continuity Clustering Software,” on page 45.
2. Continue through the installation wizard until you get to the page that prompts you to select the components to install.
3. Select one of the *Identity Manager Templates for iManager* installation options, select the *Novell Business Continuity Clustering* component, then click *Next*.

   The templates add functionality to iManager so you can manage your business continuity cluster. You must have previously installed iManager on the server where you plan to install the templates.

   Identity Manager Templates for NetWare iManager Servers: Installs the templates on a NetWare iManager server. You will be asked to specify the NetWare server where the templates will be installed later in the installation.

   Identity Manager Templates for Windows iManager Servers: Installs the templates on the local Windows iManager server. You will be asked to specify the path to Tomcat (a default path is provided) on the Windows server later in the installation.

   Novell Business Continuity Clustering: Installs the core Business Continuity Clustering engine files. This core software must be installed on all nodes in each Novell Cluster Services cluster that will be part of a business continuity cluster.
4. Do one of the following:

   - **NetWare iManager Server:** If you chose to install the Identity Manager iManager templates on a NetWare server, specify the name of the eDirectory tree and the fully distinguished name for the server where you want to install the templates. Then click *Next*.

     If you don’t know the fully distinguished name for the server, you can browse and select it.
   - **Windows iManager Server:** If you chose to install the Identity Manager iManager templates on a Windows server, specify the path to Tomcat (a default path is provided) on the server. Then click *Next*.
5. Continue through the Upgrade Reminder page, then specify the name of the eDirectory tree and the fully distinguished name for the cluster where you want to install the core software files.

If you don't know the fully distinguished name for the cluster, you can browse and select it.

6 Select the servers in the cluster where you want to install the core software files for the Business Continuity Clustering product.

All servers currently in the cluster you specified are listed and are selected by default.

You can choose to automatically start Business Continuity Clustering software on each selected node after the installation is complete. If Business Continuity Clustering software is not started automatically after the installation, you can start it manually later by rebooting the cluster server or by entering `LDBCC` at the server console.

7 Enter the name and password of an eDirectory user (or browse and select one) with sufficient rights to manage your BCC. This name should be entered in eDirectory dot format. For example, `admin.servers.novell`.

This user should have at least Read and Write rights to the All Attribute Rights property on the Cluster object of the remote cluster. For information, see Section 4.3, "Configuring a BCC Administrator User," on page 45.

8 Continue through the final installation page, then restart the cluster nodes where Identity Manager is running and where you have upgraded `libc.nlm`.

Restarting the cluster nodes can be performed in a rolling fashion in which one server is restarted while the other servers in the cluster continue running. Then another server is restarted, and then another, until all servers in the cluster have been restarted.

This lets you keep your cluster up and running and lets your users continue to access the network while cluster nodes are being restarted.

9 Repeat the above steps for each Novell Cluster Services cluster that will be part of the business continuity cluster.

4.4.2 Installing the Identity Manager Templates

After the install, you can use the Business Continuity Clustering install program to install the Identity Manager templates on additional iManager servers in the same tree as the business continuity cluster.

1 From the directory on your Windows workstation where you copied the Business Continuity Clustering software, run `install.exe`.

For download information, see Section 4.2, "Downloading the Business Continuity Clustering Software," on page 45.

2 Continue through the installation wizard until you get to the page that prompts you to select the components to install.

3 Select one of the *Identity Manager Templates for iManager* installation options, deselect the *Novell Business Continuity Clustering* component, then click *Next*.

The templates add functionality to iManager so you can manage your business continuity cluster. You must have previously installed iManager on the server where you plan to install the templates.

Identity Manager Templates for NetWare iManager Servers: Installs the templates on a NetWare iManager server. You will be asked to specify the NetWare server where the templates will be installed later in the installation.

Identity Manager Templates for Windows iManager Servers: Installs the templates on the local Windows iManager server. You will be asked to specify the path to Tomcat (a default path is provided) on the Windows server later in the installation.

4 Do one of the following:

- **NetWare iManager Server:** If you chose to install the Identity Manager iManager templates on a NetWare server, specify the name of the eDirectory tree and the fully distinguished name for the server where you want to install the templates. Then click *Next*.

  If you don't know the fully distinguished name for the server, you can browse and select it.

- **Windows iManager Server:** If you chose to install the Identity Manager iManager templates on a Windows server, specify the path to Tomcat (a default path is provided) on the server. Then click *Next*.

5 Continue through the final installation page.

4.5 What's Next

After you have installed BCC on every node in each cluster that you want to be in the business continuity cluster, continue with configuring the BCC. For information, see Chapter 6, "Configuring Business Continuity Clustering Software," on page 65.

Upgrading Business Continuity Clustering for NetWare

5

Novell® Business Continuity Clustering (BCC) 1.1 SP2 for NetWare® 6.5 SP8 supports upgrades from Novell Cluster Services™ clusters that are running BCC 1.1 SP1 for NetWare 6.5 SP6 or from clusters running BCC 1.0 (which is available on NetWare only).

BCC 1.2 for OES 2 SP1 Linux supports conversion from BCC 1.1 SP2 on NetWare. In order to convert BCC clusters from NetWare to Linux clusters, you must first upgrade existing BCC 1.0 or BCC 1.1 SP1 for NetWare clusters to BCC 1.1 SP2 for NetWare. For information about converting to BCC 1.2 for Linux, see "Converting BCC Clusters from NetWare to Linux" in the *BCC 1.2: Administration Guide for OES 2 SP1 Linux*.

This section covers two upgrade scenarios:

5.1 Guidelines for Upgrading

Use the guidelines in this section to upgrade clusters one peer cluster at a time.

5.1.1 Requirements

BCC 1.1 SP2 for NetWare requires that every node in each peer cluster be upgraded to NetWare 6.5 SP8 with the latest patches for NetWare and Novell Cluster Services. For information, see the following resources:

- **NetWare:** "Upgrading to NetWare 6.5 SP6" in the *NW65 SP8: Installation Guide* for information about how to upgrade the NetWare 6.5 SP6 servers to NetWare 6.5 SP8
- **Cluster Services:** "Installation and Setup" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide* for information about what is required for Novell Cluster Services for NetWare
- **Business Continuity Clustering:** Section 4.1, "Requirements for BCC 1.1 SP2 for NetWare," on page 35 for information about what is required for Business Continuity Clustering 1.1 SP2 for NetWare
- **Identity Manager:** *Identity Manager 3.51 Installation Guide* (http://www.novell.com/documentation/idm35/install/data/front.html). For Identity Manager configuration requirements, see Section 4.1.7, "Identity Manager 3.5.1 Bundle Edition," on page 39.

5.1.2 Performing a Rolling Cluster Upgrade

Performing a rolling upgrade of NetWare and applying the latest patches lets you keep your cluster up and running and lets your users continue to access the network while the upgrade is being performed.

For example, during a rolling cluster upgrade, one server is upgraded from NetWare SP6 to SP8 and the latest patches are applied while the other servers in the cluster continue running a previous support pack of NetWare 6.5. Then another server is upgraded and patched, and then another, until all servers in the cluster have been upgraded to NetWare 6.5 SP8 with the latest patches.

After upgrading NetWare and applying the latest patches, reboot the server to automatically load Cluster Services software.

During the upgrade process, cluster pools, volumes, and resources fail over from the server being upgraded to other servers in the cluster. After a cluster server is upgraded and brought back online, the pools, volumes, and resources that failed over to other servers in the cluster during the upgrade process fail back to the upgraded server.

5.2 Disabling BCC 1.0, Upgrading Servers to NetWare 6.5 SP8, then Enabling BCC 1.1 SP2

The safest and most straightforward way for you to upgrade from BCC 1.0 to BCC 1.1 SP2 is to disable BCC on the peer clusters, upgrade to NetWare 6.5 SP8, re-install BCC 1.1 SP2, and re-configure BCC for the clusters.

This approach leaves you without the BCC capability during the upgrade; however, the clustering high availability is still active.

1. Create a worksheet for the Identity Manager driver configuration for each peer cluster where you note the ports, landing zones, BCC drivers and driver sets, and certificates that are currently in use.

   You can use this information later when you re-create the drivers.
2. Stop the BCC drivers from running on the Identity Manager nodes of every peer cluster.
3. On the Identity Manager node in each peer cluster, delete the Identity Manager drivers and driver sets for the cluster.
4. Disable BCC for each of the cluster resources on all nodes in every peer cluster.
5. Disable BCC for each peer cluster.
6. Clean up the landing zones in each peer cluster.
7. Uninstall BCC 1.0 from each node in every peer cluster.
8. Perform a rolling cluster upgrade for all nodes in a peer cluster:
   - **8a** Issue the `cluster leave` command on one node in the cluster.
   - **8b** Update NetWare 6.5 from version SP5 or SP6 to version SP8, and apply the latest patches for NetWare and Novell Cluster Services.

     For information, see the *NW65 SP8: Installation Guide*.
   - **8c** Issue the `cluster join` command on the node.
   - **8d** Install BCC 1.1 SP2 on each node in the cluster. Do not start BCC at this time.

8e If the node is running iManager, update the following iManager plug-ins by uninstalling the existing NPM files, and then re-installing the correct iManager NPMs.

Novell Archive and Version Services (`arkmgmt.npm`)

Novell Cluster Services (`ncsmgmt.npm`)

Novell Storage Services™ (`nssmgmt.npm`)

Storage Management (`storagemgmt.npm`)

8f Repeat Step 3a through Step 3e for each node in the cluster.

9 On the Identity Manager node in the cluster, upgrade Identity Manager from version 2.*x* to 3.51.

For information, see "Upgrading" (http://www.novell.com/documentation/idm35/install/data/ampxjxi.html) in the *Identity Manager 3.5.1 Installation Guide*.

9a Before you begin, make sure that the Identity Manager node meets the 3.5.1 upgrade requirements.

9b Upgrade Identity Manager from version 2.*x* to 3.51.

IMPORTANT: Do not start the drivers yet.

The upgrade updates the Identity Manager software and its plug-ins for iManager 2.7.2 on the same node.

9c Stop Tomcat 5 by issuing `tc5stop` command.

9d Reset Apache 2 by using the `ap2webrs` command.

9e Start Tomcat 5 by using the `tomcat` command.

9f Wait at least 5 minutes for the iManager changes to take effect, then restart iManager.

9g In iManager, verify that *Identity Manager* is available.

IMPORTANT: In iManager, the Identity Manager plug-ins are displayed as version 3.6. The 3.6 version functions properly with Identity Manager 3.5.1.

---

10 On the Identity Manager node in each peer cluster, create new Identity Manager driver sets and the BCC drivers with the new BCC 1.1 SP2 templates.

IMPORTANT: Do not create a new partition when creating the driver set.

You can use the information you documented in Step 1 as a guide for creating the drivers and driver sets. If desired, you can use the same names and port numbers that you used for BCC 1.0.

11 Start the new Identity Manager drivers for each peer cluster.

12 Re-enable BCC for the peer clusters.

13 Re-enable BCC and configure BCC for the cluster resources.

5.3 Upgrading Clusters from BCC 1.0 to BCC 1.1 SP2 for NetWare

Upgrading from BCC 1.0 to BCC 1.1 SP2 while leaving the BCC configuration in place is a two-phased process where you must perform an intermediate upgrade to BCC 1.1 SP1 on NetWare 6.5 SP6 as follows:

Initial Configuration	Intermediate Configuration	Final Configuration
NetWare 6.5 SP5	NetWare 6.5 SP6	NetWare 6.5 SP8
BCC 1.0	BCC 1.1 SP1 for NetWare	BCC 1.1 SP2 for NetWare
Identity Manager 2.*x*	Identity Manager 2.*x*	Identity Manager 3.5.1

You must perform the upgrade sequence on every node in each peer cluster that is part of your existing business continuity cluster.

IMPORTANT: Before you begin, review Section 4.1, “Requirements for BCC 1.1 SP2 for NetWare,” on page 35.

During a BCC upgrade, Business Continuity Clustering 1.0 clusters are unable to communicate with Business Continuity Clustering 1.1 clusters. This condition exists temporarily until the upgrade has been completed. If an actual disaster were to occur during the upgrade, a Business Continuity Clustering 1.0 cluster can be failed over to a Business Continuity Clustering 1.1 cluster.

- Section 5.3.1, “Upgrading the BCC Cluster from 1.0 to 1.1 SP1,” on page 56
- Section 5.3.2, “Upgrading the BCC Cluster from 1.1 SP1 to 1.1 SP2,” on page 58

5.3.1 Upgrading the BCC Cluster from 1.0 to 1.1 SP1

Perform a rolling cluster upgrade for each peer cluster from BCC 1.0 to BCC 1.1 SP1 to achieve the intermediate configuration as follows:

Initial Configuration	Intermediate Configuration
NetWare 6.5 SP5	NetWare 6.5 SP6
BCC 1.0	BCC 1.1 SP1 for NetWare
Identity Manager 2.*x*	Identity Manager 2.*x*
iManager 2.5 or 2.6	iManager 2.6

1. Create a worksheet for the Identity Manager driver configuration for each peer cluster where you note the ports, landing zones, BCC drivers and driver sets, and certificates that are currently in use.

   You can use this information later in Step 5b when you re-create the drivers.
2. Stop the Identity Manager (formerly DirXML) drivers from running on the Identity Manager node of every peer cluster in the business continuity cluster.
3. Perform a rolling cluster upgrade for all nodes in a peer cluster:
   - **3a** Issue the `cluster leave` command on one node in the cluster.
   - **3b** Update NetWare 6.5 from version SP5 to SP6, and apply the latest patches for NetWare and Novell Cluster Services.

     You can perform a script upgrade by using NWCONFIG.

     For information, see the *NW65 SP8: Installation Guide*.

See the Business Continuity Clustering 1.1 Readme file (http://www.novell.com/documentation/bcc/pdfdoc/readme/readme111.pdf) for instructions on mandatory patches.

3c Issue the `cluster join` command on the node.

3d Upgrade Business Continuity Clustering from version 1.0 to 1.1 SP1 by installing BCC 1.1 SP1.

IMPORTANT: Do not start BCC at this time.

The Business Continuity Clustering 1.1 installation program automatically detects if Business Continuity Clustering 1.0 is installed and performs the necessary updates to convert 1.0 to 1.1. This includes searching eDirectory for SAN scripts and updating those scripts to be SMI-S compliant.

For installation instructions for BCC 1.1 SP1, see "Running the Business Continuity Cluster Installation Program" (http://www.novell.com/documentation/bcc/bcc_administration_nw/data/ht05s5vv.html#h7qplroj) in the *Novell Business Continuity Clustering 1.1 for NetWare Administration Guide* (http://www.novell.com/documentation/bcc/bcc_administration_nw/data/bktitle.html).

3e If the node is running iManager, update the following iManager plug-ins by uninstalling the existing NPM files, and then re-installing the correct iManager NPMs.

Novell Archive and Version Services (`arkmgmt.npm`)

Novell Cluster Services (`ncsmgmt.npm`)

Novell Storage Services™ (`nssmgmt.npm`)

Storage Management (`storagemgmt.npm`)

3f Repeat Step 3a through Step 3e for each node in the cluster.

4 Repeat the rolling cluster upgrade in Step 3 for each peer cluster.

5 On the Identity Manager node in each peer cluster, delete and re-create the Identity Manager driver sets for the cluster:

5a Delete the Identify manager 2.*x* driver sets.

5b Create new Identity Manager driver sets and the BCC drivers with the new BCC 1.1 SP1 templates.

IMPORTANT: Do not create a new partition when creating the driver set.

You can use the information you documented in Step 1 as a guide for creating the drivers and driver sets. If desired, you can use the same names and port numbers that you used for BCC 1.0.

5c Repeat Step 5a through Step 5b for each node in the peer cluster

6 Issue the `ldbcc.ncf` command on each server node to have them join the BCC.

7 Start the Identity Manager drivers.

8 BCC disable and enable each BCC resource to make sure the BCC attributes are updated.

If a BCC enabled resource is missing BCC attributes, try deleting the eDirectory™ Cluster Resource objects for the pool resource, then re-create the cluster resource to get it back to a usable state in the BCC.

9 Reset the cluster peer credentials between clusters.

The BCC Administrator user credentials that were set for BCC 1.0 do not work with BCC 1.1. A fully distinguished eDirectory name (FDN) was required for BCC 1.0, but BCC 1.1 requires only the BCC administrator name.

For instructions on resetting BCC Administrator user credentials, see Section 7.2, “Changing Cluster Peer Credentials,” on page 82.

IMPORTANT: Make sure the administrator username meets the requirements specified in Section 4.3, “Configuring a BCC Administrator User,” on page 45.

---

10 Verify that Novell Cluster Services and BCC 1.1 SP1 appear to be functioning correctly by performing the following tests:

10a Create a new BCC enabled pool resource, BCC migrate it between peer clusters, and verify that it migrates correctly.

10b Make a load script change (for example, add a space character and save the change) to an existing BCC resource that was created with 1.0, allow the revised load script to synchronize, then verify that the load script was updated at all eDirectory locations.

10c BCC migrate the pool resource that you modified in Step 10b between peer clusters, and verify that it migrates correctly.

10d Check all SAN scripts ensure that they will perform the desired functions.

11 Continue with Section 5.3.2, “Upgrading the BCC Cluster from 1.1 SP1 to 1.1 SP2,” on page 58.

5.3.2 Upgrading the BCC Cluster from 1.1 SP1 to 1.1 SP2

In the second phase of the upgrade from BCC 1.0, you perform a rolling cluster upgrade for each peer cluster from BCC 1.1 SP1 to BCC 1.1 SP2 to achieve the final configuration as follows:

Intermediate Configuration	Final Configuration
NetWare 6.5 SP6	NetWare 6.5 SP8
BCC 1.1 SP1 for NetWare	BCC 1.1 SP2 for NetWare
Identity Manager 2.*x*	Identity Manager 3.5.1

1 Stop the Identity Manager (formerly DirXML) drivers from running in the Identity Manager node of every peer cluster in the business continuity cluster.

2 Delete the Identity Manager 2.*x* drivers and driver sets on the Identity Manager node of every peer cluster.

3 Perform a rolling cluster upgrade from NetWare 6.5 SP6 to SP8 for all nodes in a peer cluster:

3a Issue the `cluster leave` command on one node in the cluster.

3b Upgrade NetWare 6.5 from version SP6 to SP8, and apply the latest patches for NetWare and Novell Cluster Services.

You can perform a script upgrade by using NWCONFIG.

For information, see the *NW65 SP8: Installation Guide*.

3c If the node is running iManager, update the storage related plug-ins for iManager.

After upgrading to NetWare 6.5 SP8, iManager no longer displays the roles for Clusters, Storage, and DirXML. The storage related plug-ins require special handling because some storage features were reorganized in the NetWare 6.5 SP8 release. For information, see "Storage Related Plug-Ins Must Be Uninstalled" (http://www.novell.com/documentation/nw65/nw65_readme/data/bi59826.html#biv1r9v) in the *NW65 SP8: Readme* (http://www.novell.com/documentation/nw65/nw65_readme/data/readme.html).

NOTE: The DirXML plug-ins will be replaced by Identity Manager 3.5.1 plug-ins in Step 5 on page 59.

3c1 In iManager, uninstall the old iManager plug-ins for Novell Cluster Services (`ncsmgmt.npm`), NSS (`nssmgmt.npm`), Archive and Version Services (`arkmgmt.npm`), and Storage Management, (`storagemgmt.npm`).

3c2 In iManager, install the new set of storage-related plug-ins for iManager 2.7.2 from the NetWare 6.5 SP8 installation media:

Novell AFP (`afpmgmt.npm`)
Novell Archive and Version Services (`arkmgmt.npm`)
Novell CIFS (`cifsmgmt.npm`)
Novell Cluster Services (`ncsmgmt.npm`)
Novell Distributed File Services (`dfsmgmt.npm`)
Novell Storage Services (`nssmgmt.npm`)
Storage Management, (`storagemgmt.npm`)

3c3 Stop Tomcat 5 by issuing `tc5stop` command.

3c4 Reset Apache 2 by using the `ap2webrs` command.

3c5 Start Tomcat 5 by using the `tomcat` command.

3c6 Wait at least 5 minutes for the iManager changes to take effect, then restart iManager.

3c7 In iManager, verify that BCC tasks are available in the *Clusters* role.

3d Repeat Step 3a through Step 3c7 for each node in the cluster.

4 On each node in the cluster, upgrade Business Continuity Clustering from version 1.1 SP1 to 1.1 SP2 by installing BCC 1.1 SP2.

IMPORTANT: Do not start BCC at this time.

For instructions, see Chapter 4, "Installing Business Continuity Clustering," on page 35.

5 On the Identity Manager node in the cluster, upgrade Identity Manager from version 2.*x* to 3.51.

For information, see "Upgrading" (http://www.novell.com/documentation/idm35/install/data/ampxjxi.html) in the *Identity Manager 3.5.1 Installation Guide*.

5a Before you begin, make sure that the Identity Manager node meets the 3.5.1 upgrade requirements.

5b Upgrade Identity Manager from version 2.*x* to 3.51.

IMPORTANT: Do not start the drivers yet.

The upgrade updates the Identity Manager software and its plug-ins for iManager 2.7.2 on the same node.

5c Stop Tomcat 5 by issuing `tc5stop` command.

5d Reset Apache 2 by using the `ap2webrs` command.

5e Start Tomcat 5 by using the `tomcat` command.

5f Wait at least 5 minutes for the iManager changes to take effect, then restart iManager.

5g In iManager, verify that *Identity Manager* is available.

IMPORTANT: In iManager, the Identity Manager plug-ins are displayed as version 3.6. The 3.6 version functions properly with Identity Manager 3.5.1.

---

6 Issue the `cluster join` command on each node in the cluster.

7 Issue the `ldbcc.ncf` command on each node to have them join the BCC.

8 Repeat Step 3 through Step 7 for each peer cluster.

9 Verify that Novell Cluster Services and Business Continuity Clustering appear to be functioning correctly by migrating a BCC enabled resource between peer clusters.

10 Create the new BCC driver sets for Identity Manager.

11 On the Identity Manager node in every peer cluster, create the new Identity Manager drivers with the new BCC 1.1 SP2 templates.

IMPORTANT: Do not create a new partition when creating the driver set.

As a guide, you can use the information from Step 1 and Step 5b in Section 5.3.1, “Upgrading the BCC Cluster from 1.0 to 1.1 SP1,” on page 56.

12 Start the new Identity Manager drivers for each peer cluster.

13 Verify that Novell Cluster Services and BCC 1.1 SP2 appear to be functioning correctly by performing the following tests:

13a Create a new BCC enabled pool resource, BCC migrate it between peer clusters, and verify that it migrates correctly.

13b Make a load script change (for example, add a space character and save the change) to an existing BCC resource that was created with 1.0, allow the revised load script to synchronize, then verify that the load script was updated at all eDirectory locations.

13c BCC migrate the pool resource that you modified in Step 13b between peer clusters, and verify that it migrates correctly and that the Identity Manager drivers are synchronizing.

13d Check all SAN scripts ensure that they will perform the desired functions.

5.4 Upgrading Clusters from BCC 1.1 SP1 to SP2 for NetWare

Use the procedure in this section to upgrade from BCC version SP1 to SP2:

Initial Configuration	Final Configuration
NetWare 6.5 SP6	NetWare 6.5 SP8
BCC 1.1 SP1 for NetWare	BCC 1.1 SP2 for NetWare
Identity Manager 2.x or 3.0.*x*	Identity Manager 3.5.1

You must perform the upgrade sequence on every node in each peer cluster that is part of your existing business continuity cluster.

IMPORTANT: Before you begin, review Section 4.1, “Requirements for BCC 1.1 SP2 for NetWare,” on page 35.

To upgrade BCC 1.1 from SP1 to SP2, perform the following tasks:

5.4.1 Upgrading NetWare and BCC on the Clusters

Perform a rolling cluster upgrade from NetWare 6.5 SP6 to SP8 for all nodes in a peer cluster:

1. Stop the Identity Manager drivers from running in the Identity Manager node of every peer cluster in the business continuity cluster.
2. Issue the `cluster leave` command on one node in the cluster.
3. Upgrade NetWare 6.5 from version SP6 to SP8, and apply the latest patches for NetWare and Novell Cluster Services.

   You can perform a script upgrade by using NWCONFIG.

   For information, see the *NW65 SP8: Installation Guide*.
4. If the node is running iManager, update the storage related plug-ins for iManager.

   After upgrading to NetWare 6.5 SP8, iManager no longer displays the roles for Clusters, Storage, and DirXML. The storage related plug-ins require special handling because some storage features were reorganized in the NetWare 6.5 SP8 release. For information, see “Storage Related Plug-Ins Must Be Uninstalled” (http://www.novell.com/documentation/nw65/nw65_readme/data/bi59826.html#biv1r9v) in the *NW65 SP8: Readme* (http://www.novell.com/documentation/nw65/nw65_readme/data/readme.html).

   NOTE: The DirXML plug-ins will be replaced by Identity Manager 3.5.1 plug-ins in Section 5.4.3, “Upgrading Identity Manager,” on page 62.

   4a In iManager, uninstall the old iManager plug-ins for Novell Cluster Services (`ncsmgmt.npm`), NSS (`nssmgmt.npm`), Archive and Version Services (`arkmgmt.npm`), and Storage Management, (`storagemgmt.npm`).

   4b In iManager, install the new set of storage-related plug-ins for iManager 2.7.2 from the NetWare 6.5 SP8 installation media:

   Novell AFP (`afpmgmt.npm`)
   Novell Archive and Version Services (`arkmgmt.npm`)
   Novell CIFS (`cifsmgmt.npm`)
   Novell Cluster Services (`ncsmgmt.npm`)
   Novell Distributed File Services (`dfsmgmt.npm`)
   Novell Storage Services (`nssmgmt.npm`)

Storage Management, (`storagemgmt.npm`)

4c Stop Tomcat 5 by issuing `tc5stop` command.

4d Reset Apache 2 by using the `ap2webrs` command.

4e Start Tomcat 5 by using the `tomcat` command.

4f Wait at least 5 minutes for the iManager changes to take effect, then restart iManager.

4g In iManager, verify that BCC tasks are available in the *Clusters* role.

5 Repeat Step 2 through Step 4g for each node in the cluster.

6 Issue the `cluster join` command on each node in the cluster.

7 On each node in the cluster, upgrade Business Continuity Clustering from version 1.1 SP1 to 1.1 SP2 by installing BCC 1.1 SP2.

IMPORTANT: Do not start BCC at this time.

For instructions, see Chapter 4, “Installing Business Continuity Clustering,” on page 35.

8 Issue the `ldbcc.ncf` command on each node to have them join the BCC.

9 Repeat Step 2 through Step 8 for each peer cluster in your business continuity cluster.

5.4.2 Authorizing the BCC Administrator User

The BCC Administrator user must be a trustee of the Cluster objects in your BCC, and have at least read and write rights to the all attributes rights property.

1 Start your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2 Specify your username and password, specify the tree where you want to log in, then click *Login*.

3 In the left column, click *Rights*, then click the *Modify Trustees* link.

4 Specify the Cluster object name, or browse and select it, then click *OK*.

5 If the BCC Administrator user is not listed as a trustee, click the *Add Trustee* button, browse and select the User object, then click *OK*.

6 Click *Assigned Rights* for the BCC Administrator user, then ensure that the Read and Write check boxes are selected for the All Attributes Rights property.

7 Click *Done* to save your changes.

8 Repeat Step 3 through Step 7 for the other Cluster objects in for every cluster in the business continuity cluster.

5.4.3 Upgrading Identity Manager

On the Identity Manager node in the cluster, upgrade to Identity Manager 3.5.1.

Before you begin:

- Review the BCC 1.1 SP2 configuration requirements in Section 4.1.7, “Identity Manager 3.5.1 Bundle Edition,” on page 39.
- Make sure that the Identity Manager node meets the 3.5.1 upgrade requirements. For information, see “Upgrading” (http://www.novell.com/documentation/idm35/install/data/ampxjxi.html) in the *Identity Manager 3.5.1 Installation Guide*.

For information about installing or upgrading Identity Manager, see the *Identity Manager 3.51 Installation Guide* (http://www.novell.com/documentation/idm35/install/data/front.html).

1. Upgrade Identity Manager to version 3.51.

   IMPORTANT: Do not start the drivers yet.

   The upgrade updates the Identity Manager software and its plug-ins for iManager 2.7.2 on the same node.
2. Stop Tomcat 5 by issuing `tc5stop` command.
3. Reset Apache 2 by using the `ap2webrs` command.
4. Start Tomcat 5 by using the `tomcat` command.
5. Wait at least 5 minutes for the iManager changes to take effect, then restart iManager.
6. In iManager, verify that *Identity Manager* is available.

   IMPORTANT: In iManager, the Identity Manager plug-ins are displayed as version 3.6. The 3.6 version functions properly with Identity Manager 3.5.1.

5.4.4 Deleting and Re-Creating the BCC Driver Sets and Drivers

After completing the upgrade procedures for every node in all clusters of the business continuity cluster, you must delete and re-create the Identity Manager driver sets and drivers with the BCC 1.1 SP2 templates.

On the Identity Manager node in every peer cluster, do the following:

1. Delete the Identity Manager 2.*x* or 3.0.*x* drivers and driver sets on the Identity Manager node of every peer cluster.
2. Create the new BCC driver sets for Identity Manager.
3. Create the new Identity Manager drivers with the new BCC 1.1 SP2 templates.
4. Start the new Identity Manager drivers for each peer cluster.

5.4.5 Verifying the BCC Upgrade

Verify that Novell Cluster Services and BCC 1.1 SP2 appear to be functioning correctly by performing the following tests:

1. Create a new BCC enabled pool resource, BCC migrate it between peer clusters, and verify that it migrates correctly.

2. Make a load script change (for example, add a space character and save the change) to an existing BCC resource that was created with 1.0, allow the revised load script to synchronize, then verify that the load script was updated at all eDirectory locations.
3. BCC migrate the pool resource that you modified in Step 13b between peer clusters, and verify that it migrates correctly and that the Identity Manager drivers are synchronizing.
4. Check all SAN scripts ensure that they will perform the desired functions.

Configuring Business Continuity Clustering Software

6

After you have installed and configured Identity Manager and the Novell® Business Continuity Clustering software, and you have configured file system mirroring, you need to set up the BCC software.

- Section 6.1, “Configuring Identity Manager Drivers for the Business Continuity Cluster,” on page 65
- Section 6.2, “Configuring Clusters for Business Continuity,” on page 71
- Section 6.3, “BCC-Enabling Cluster Resources,” on page 76

6.1 Configuring Identity Manager Drivers for the Business Continuity Cluster

The Identity Manager preconfigured templates for iManager that were installed when you ran the Novell Business Continuity Clustering installation must be configured so you can properly manage your business continuity cluster. The preconfigured templates include the following:

- **Cluster Resource Synchronization:** A set of policies, filters, and objects that synchronize cluster resource information between any two of the peer clusters. This template must always be configured, even in a single-tree business continuity cluster.
- **User Object Synchronization:** A set of policies, filters, and objects that synchronize User objects between any any two trees (or partitions) that contain the clusters in the business continuity cluster. Typically, this template is used to configure drivers when the clusters in your business continuity cluster are in different eDirectory™ trees. You might also need to set up User Object Synchronization drivers between clusters if you put User objects in a different eDirectory partition than is used for the Cluster objects; however, this is not a recommended configuration. See Appendix B, “Implementing a Multiple-Tree BCC,” on page 125 for more information about implementing BCC between two trees.

The Identity Manager engine and eDirectory driver must be installed on one node in each cluster. The node where Identity Manager is installed must have an eDirectory full replica with at least read/write access to all eDirectory objects that will be synchronized between clusters. For information about the full replica requirements, see Section 4.1.5, “Novell eDirectory 8.8,” on page 38.

Identity Manager requires a credential that allows you to use drivers beyond an evaluation period. The credential can be found in the BCC license. In the Identity Manager interface in iManager, enter the credential for each driver that you create for BCC. You must also enter the credential for the matching driver that is installed in a peer cluster. You can enter the credential, or put the credential in a file that you point to.

- Section 6.1.1, “Configuring the Identity Manager Drivers and Templates,” on page 66
- Section 6.1.2, “Creating SSL Certificates,” on page 68
- Section 6.1.3, “Synchronizing Identity Manager Drivers,” on page 68
- Section 6.1.4, “Preventing Identity Manager Synchronization Loops,” on page 69

6.1.1 Configuring the Identity Manager Drivers and Templates

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2. Specify your username and password, specify the tree where you want to log in, then click *Login*.

3. In the left column, click *Identity Manager Utilities*, then click the *New Driver* link.

4. Choose to place the new driver in a new driver set, then click *Next*.

   Both the *User Object Synchronization* driver and the *Cluster Resource Synchronization* driver can be added to the same driver set.

5. Specify the driver set name, context, and the server that the driver set will be associated with.

   The server is the same server where you installed the Identity Manager engine and eDirectory driver.

6. Choose to *not* create a new partition for the driver set, then click *Next*.

7. Choose to import a preconfigured driver from the server, select the Identity Manager preconfigured template for cluster resource synchronization, then click *Next*.

   The template name is `BCCClusterResourceSynchronization.XML`.

8. Fill in the values on the wizard page as prompted, then click *Next*.

   Each field contains an example of the type of information that should go into the field. Descriptions of the information required are also included with each field.

   - **Driver name:** Specify a unique name for this driver to identify its function. For example, *Cluster1SyncCluster2*. If you use both preconfigured templates, you must specify different driver names for each driver template.
   - **Name of SSL Certificate:** If you do not have an SSL certificate, leave this value set to the default. The certificate is created later in the configuration process. See "Creating SSL Certificates" on page 68 for instructions on creating SSL certificates.

     In a single tree configuration, if you specify the SSL CertificateDNS certificate that was created when you installed NetWare on the Identity Manager node, you do not need to create an additional SSL certificate later.
   - **DNS name of other IDM node:** Specify the DNS name or IP address of the Identity Manager server in the other cluster.
   - **Port number for this driver:** If you have a business continuity cluster that consists of three or four clusters, you must specify unique port numbers for each driver template set. The default port number is 2002.

     You must specify the same port number for the same template in the other cluster. For example, if you specify 2003 as the port number for the resource synchronization template, you must specify 2003 as the port number for the resource synchronization template in the peer driver for the other cluster.
   - **Full Distinguished Name (DN) of the cluster this driver services:** For example, Cluster1.siteA.Novell.

- **Fully Distinguished Name (DN) of the landing zone container:** Specify the context of the container where the cluster pool and volume objects in the other cluster are placed when they are synchronized to this cluster.

  This container is referred to as the landing zone. The NCP™ server objects for the virtual server of a BCC enabled resource are also placed in the landing zone.

  IMPORTANT: The context must already exist and must be specified using dot format without the tree name. For example, siteA.Novell.

  Prior to performing this step, you could create a separate container in eDirectory specifically for these cluster pool and volume objects. You would then specify the context of the new container in this step.

The IDM Driver object must have sufficient rights to any object it reads or writes in the following containers:

- The Identity Manager driver set container.
- The container where the Cluster object resides.
- The container where the Server objects reside.

  If server objects reside in multiple containers, this must be a container high enough in the tree to be above all containers that contain server objects. The best practice is to have all server objects in one container.

- The container where the cluster pool and volume objects are placed when they are synchronized to this cluster.

  This container is referred to as the landing zone. The NCP server objects for the virtual server of a BCC enabled resource are also placed in the landing zone.

You can do this by making the IDM Driver object security equivalent to another User object with those rights. See Step 9.

IMPORTANT: If you choose to include User object synchronization, exclude the Admin User object from being synchronized. See Step 7 in Section B.5, “Synchronizing the BCC-Specific Identity Manager Drivers,” on page 128 for information about synchronizing User objects when adding new clusters to the business continuity cluster.

---

9 Make the IDM Driver object security equivalent to an existing User object:

9a Click *Define Security Equivalences*, then click *Add*.

9b Browse to and select the desired User object, then click *OK*.

9c Click *Next*, then click *Finish*.

10 Repeat Step 1 through Step 9 above on the other clusters in your business continuity cluster.

This includes creating a new driver and driver set for each cluster.

IMPORTANT: If you have upgraded to Identity Manager 3 and click either the cluster resource synchronization driver or the user object synchronization driver, a message is displayed prompting you to convert the driver to a new architecture. Click *OK* to convert the driver.

6.1.2 Creating SSL Certificates

It is recommended that you create an SSL certificate for the Cluster Resource Synchronization driver. Creating one certificate creates the certificate for a driver pair. For example, creating an SSL certificate for the Cluster Resource Synchronization driver also creates the certificate for the Cluster Resource Synchronization drivers on the other clusters.

To create an SSL certificate:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager Utilities*, then click *NDS-to-NDS Driver Certificates*.
4. Specify the requested driver information for this cluster, then click *Next*.

   You must specify the driver name (including the context) you supplied in Step 8 on page 66 for this cluster. Use the following format when specifying the driver name:

   `DriverName.DriverSet.OrganizationalUnit.OrganizationName`

   Ensure that there are no spaces (beginning or end) in the specified context, and do not use the `cn=DriverName.ou=OrganizationalUnitName.o=OrganizationName` format.
5. Specify the requested driver information for the driver in the other cluster.

   Use the same format specified in Step 4.
6. Click *Next*, then click *Finish*.

6.1.3 Synchronizing Identity Manager Drivers

If you are adding a new cluster to an existing business continuity cluster, you must synchronize the BCC-specific Identity Manager drivers after you have created the BCC-specific Identity Manager drivers and SSL certificates. If the BCC-specific Identity Manager drivers are not synchronized, clusters cannot be enabled for business continuity. Synchronizing the Identity Manager drivers is only necessary when you are adding a new cluster to an existing business continuity cluster.

To synchronize the BCC-specific Identity Manager drivers:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager*, then click the *Identity Manager Overview* link.
4. Search for and find the BCC driver set.
5. Click the red *Cluster Sync* icon for the driver you want to synchronize, then click the *Migrate from eDirectory* button.

6 Click *Add*, browse to and select the Cluster object for the new cluster you are adding to the business continuity cluster, then click *OK*.

Selecting the Cluster object causes the BCC-specific Identity Manager drivers to synchronize.

If you have multiple eDirectory trees in your BCC, see Section B.5, “Synchronizing the BCC-Specific Identity Manager Drivers,” on page 128.

6.1.4 Preventing Identity Manager Synchronization Loops

If you have three or more clusters in your business continuity cluster, you should set up synchronization for the User objects and Cluster Resource objects in a manner that prevents Identity Manager synchronization loops. Identity Manager synchronization loops can cause excessive network traffic and slow server communication and performance.

For example, in a three-cluster business continuity cluster, an Identity Manager synchronization loop occurs when Cluster One is configured to synchronize with Cluster Two, Cluster Two is configured to synchronize with Cluster Three, and Cluster Three is configured to synchronize back to Cluster One. This is illustrated in Figure 6-1 below.

Figure 6-1 *Three-Cluster Identity Manager Synchronization Loop*

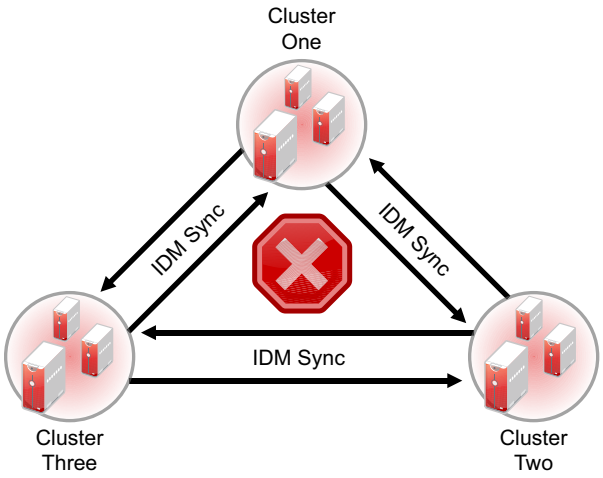

A preferred method is to make Cluster One an Identity Manager synchronization master in which Cluster One synchronizes with Cluster Two, and Cluster Two and Cluster Three both synchronize with Cluster One. This is illustrated in Figure 6-2 below.

Figure 6-2 *Three-Cluster Identity Manager Synchronization Master*

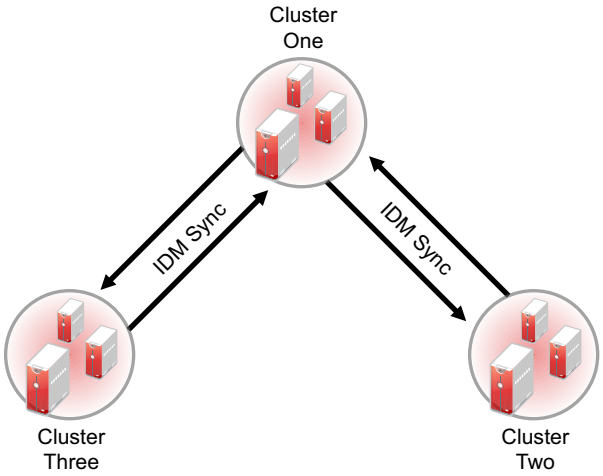

You could also have Cluster One synchronize with Cluster Two, Cluster Two synchronize with Cluster Three, and Cluster Three synchronize back to Cluster Two as illustrated in Figure 6-3.

Figure 6-3 *Alternate Three-Cluster Identity Manager Synchronization Scenario*

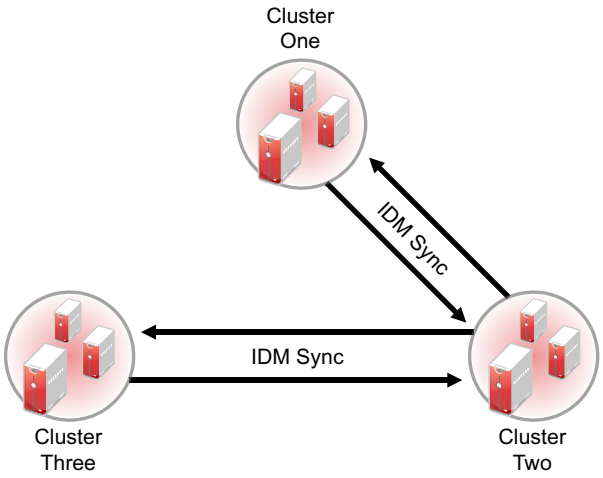

To change your BCC synchronization scenario:

1. In the Connections section of the Business Continuity Cluster Properties page, select one or more peer clusters that you want a cluster to synchronize to, then click *Edit*.

In order for a cluster to appear in the list of possible peer clusters, that cluster must have the following:

- Business Continuity Clustering software installed.
- Identity Manager installed.
- The BCC-specific Identity Manager drivers configured and running.
- Be enabled for business continuity.

6.2 Configuring Clusters for Business Continuity

The following tasks must be performed on each separate Novell Cluster Services cluster that you want to be part of the business continuity cluster:

NOTE: Identity Manager must be configured and running before configuring clusters for business continuity.

6.2.1 Enabling Clusters for Business Continuity

If you want to enable a cluster to fail over selected resources or all cluster resources to another cluster, you must enable business continuity on that cluster.

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed. This server should be in the same eDirectory tree as the cluster you are enabling for business continuity.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. Ensure that the BCC-specific Identity Manager drivers are running:
   - **3a** In the left column, click *Identity Manager,* and then click the *Identity Manager Overview* link.
   - **3b** Search the eDirectory Container or tree for the BCC-specific Identity Manager drivers.
   - **3c** For each driver, click the upper right corner of the driver icon to see if a driver is started or stopped.
   - **3d** If the driver is stopped, start it by selecting *Start*.
4. In the left column, click *Clusters*, then click the *Cluster Options* link.
5. Specify a cluster name, or browse and select one.
6. Click the *Properties* button, then click the *Business Continuity* tab.
7. Ensure that the *Enable Business Continuity Features* check box is selected.
8. Repeat Step 1 through Step 7 for the other cluster that this cluster will migrate resources to.
9. Continue with Adding Cluster Peer Credentials.

6.2.2 Adding Cluster Peer Credentials

In order for one cluster to connect to a second cluster, the first cluster must be able to authenticate to the second cluster. To make this possible, you must add the username and password of the user that the selected cluster will use to connect to the selected peer cluster.

IMPORTANT: In order to add or change cluster peer credentials, you must access iManager on a server that is in the same eDirectory tree as the cluster you are adding or changing peer credentials for.

---

1. In the *Connections* section of the Business Continuity Cluster Properties page, select the peer cluster, then click *Edit*.

   In order for a cluster to appear in the list of possible peer clusters, the cluster must have the following:

   - Business Continuity Clustering software installed.
   - Identity Manager installed and running.
   - The BCC-specific Identity Manager drivers configured and running.
   - Be enabled for business continuity.

2. Add the administrator username and password that the selected cluster will use to connect to the selected peer cluster.

   When adding the administrator username, do not include the context for the user. For example, use `bccadmin` instead of `bccadmin.prv.novell`.

   Rather than using the Admin user to administer your BCC, you should consider creating another user with sufficient rights to the appropriate contexts in your eDirectory tree to manage your BCC. For information, see Section 4.3, “Configuring a BCC Administrator User,” on page 45.

3. Repeat Step 1 and Step 2 for the other cluster that this cluster will migrate resources to.

4. Continue with Adding Search-and-Replace Values to the Resource Replacement Script.

6.2.3 Adding Search-and-Replace Values to the Resource Replacement Script

To enable a resource for business continuity, certain values (such as IP addresses) specified in resource load and unload scripts need to be changed in corresponding resources in the other clusters. You need to add the search-and-replace strings that are used to transform cluster resource load and unload scripts from this cluster to another cluster. Replacement scripts are for inbound changes to scripts for objects being synchronized from other clusters, not outbound.

TIP: You can see the IP addresses that are currently assigned to resources by entering the `display secondary ipaddress` command at the NetWare server console of cluster servers.

The search-and-replace data is cluster-specific, and it is not synchronized via Identity Manager between the clusters in the business continuity cluster.

To add resource script search-and-replace values:

1. In iManager, click *Clusters* > *Cluster Options*, select the Cluster object, click *Properties*, then select the *Busines Continuity*.
2. In the *Resource Script Replacements* section of the Business Continuity Cluster Properties page, click *New*.
3. Add the desired search-and-replace values.

   The search-and-replace values you specify here apply to all resources in the cluster that have been enabled for business continuity.

   For example, if you specify 10.1.1.1 as the search value and 192.168.1.1 as the replace value, the resource with the 10.1.1.1 IP address in its scripts is searched for in the primary cluster and, if found, the 192.168.1.1 IP address is assigned to the corresponding resource in the secondary cluster.

   You can also specify global search-and-replace addresses for multiple resources in one line. This can be used only if the last digits in the IP addresses are the same in both clusters. For example, if you specify 10.1.1. as the search value and 192.168.1. as the replace value, the software finds the 10.1.1.1, 10.1.1.2, 10.1.1.3 and 10.1.1.4 addresses, and replaces them with the 192.168.1.1, 192.168.1.2, 192.168.1.3, and 192.168.1.4 addresses, respectively.

   IMPORTANT: Make sure to use a trailing dot in the search-and-replace value. If a trailing dot is not used, 10.1.1 could be replaced with an IP value such as 192.168.100 instead of 192.168.1.

   ---

4. (Optional) Select the *Use Regular Expressions* check box to use wildcard characters in your search-and-replace values. The following links provide information on regular expressions and wildcard characters:

   - Regular Expressions (http://www.opengroup.org/onlinepubs/007908799/xbd/re.html)
   - Regular-Expressions.info (http://www.regular-expressions.info/)
   - Wikipedia (http://en.wikipedia.org/wiki/Regular_expression)
   - oreilly.com (http://www.oreilly.com/catalog/regex/)

   You can find additional information on regular expressions and wildcard characters by searching the Web.
5. Click *Apply* to save your changes.

   Clicking *OK* does not apply the changes to the directory.
6. Verify that the change has been synchronized with the peer clusters by the Identity Vault.
7. Continue with Section 6.2.4, “Adding SAN Management Configuration Information,” on page 73.

6.2.4 Adding SAN Management Configuration Information

You can create scripts and add commands that are specific to your SAN hardware. These scripts and commands might be needed to promote mirrored LUNs to primary on the cluster where the pool resource is being migrated to, or demote mirrored LUNs to secondary on the cluster where the pool resource is being migrated from.

You can also add commands and Perl scripts to the resource scripts to call other scripts. Any command that can be run at the NetWare server console can be used. The scripts or commands you add are stored in eDirectory. If you add commands to call outside scripts, those scripts must exist on every server in the cluster.

IMPORTANT: Scripts are not synchronized by Identity Manager.

To add SAN management configuration information:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Clusters*, then click the *Cluster Options* link.
4. Specify a cluster name, or browse and select one.
5. Under *Cluster Objects*, select a cluster resource that is enabled for business continuity, then click *Details*.

   Cluster resources that are enabled for business continuity have the BCC label on the resource type icon.
6. Click the *Business Continuity* tab, then click *SAN Management*.
7. Create BCC SAN management load and unload scripts:

   7a Under *BCC Load Scripts*, click *New* to bring up a page that lets you create a script to promote mirrored LUNs on a cluster.

   You can also delete a script, edit a script by clicking *Details*, or change the order that load scripts execute by clicking the *Move Up* and *Move Down* links.

   7b Specify the values on the SAN Management Script Details page.

   Descriptions of the information required for the page fields and options include:

   - **Name and Description:** Specify a name, and if desired, a description of the script you are creating.
   - **CIMOM IP/DNS:** If you are not using a template and if you selected the *CIM Client* check box on the previous page, specify the IP address or DNS name for your SAN. This is the IP address or DNS name that is used for SAN management.
   - **Namespace:** If you selected the *CIM Client* check box on the previous page, accept the default namespace, or specify a different namespace for your SAN.

     Namespace determines which models and classes are used with your SAN. Consult your SAN documentation to determine which namespace is required for your SAN.
   - **Username and Password:** If you selected the *CIM Client* check box on the previous page, specify the username and password that is used to connect to and manage your SAN.
   - **Port:** If you selected the *CIM Client* check box on the previous page, accept the default port number or specify a different port number. This is the port number that CIMOM (your SAN manager) uses. Consult your SAN documentation to determine which port number you should use.

- **Secure:** If you selected the *CIM Client* check box on the previous page, select or deselect the *Secure* check box depending whether you want SAN management communication to be secure (HTTPS) or non-secure (HTTP).
- **Script Parameters:** If desired, specify variables and values for the variables that are used in the SAN management script.

  To specify a variable, click *New*, then provide the variable name and value in the fields provided. Click *OK* to save your entries. You can specify additional variables by clicking *New* again and providing variable names and values. You can also edit and delete existing script parameters by clicking the applicable link.
- **Script Parameters Text Box:** Use this text box to add script commands to the script you are creating.

  These script commands are specific to your SAN hardware. You can add a Perl script, or any commands that can be run on Linux or NetWare (depending on your platform). If you add commands to call outside scripts, those scripts must exist on every server in the cluster.
- **CIM Enabled:** Select this box if your SAN supports SMI-S and you did not select the *CIM Client* check box on the previous page. This causes the CIM-specific fields to become active on this page.
- **Synchronous:** If this check box is not selected, multiple scripts can be run concurrently. Selecting the box causes scripts to run individually, one after another. Most SAN vendors do not support running multiple scripts concurrently.
- **Edit Flags:** This is an advanced feature, and should not be used except under the direction of Novell Support.

7c Click *Apply* and *OK* on the Script Details page, then click *OK* on the Resource Properties page to save your script changes.

IMPORTANT: After clicking *Apply* and *OK* on the Script Details page, you are returned to the Resource Properties page (with the *Business Continuity* tab selected). If you do not click *OK* on the Resource Properties page, your script changes are not saved.

IMPORTANT: The CIMOM daemons on all nodes in the business continuity cluster should be configured to bind to all IP addresses on the server.

Business Continuity Clustering connects to the CIMOM by using the master IP address for the cluster. Because the master IP address moves to other nodes during a failover or migration, the CIMOM must be configured to bind to all IP addresses (secondary and primary), rather than just the primary IP address of the host.

You can do this by editing the `openwbem.conf` file. See "Changing the OpenWBEM CIMOM Configuration" in the *NW 6.5 SP8: OpenWBEM Services Administration Guide*.

6.2.5 Verifying BCC Administrator User Trustee Rights and Credentials

You must ensure that the user who manages your BCC (BCC Administrator user) is a trustee of the Cluster objects and has at least Read and Write eDirectory rights to the All Attributes Rights property. For instructions, see "Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects" on page 46.

You must also ensure that the BCC Administrator user has file system rights to the `_ADMIN:\Novell\Cluster` directory of all nodes in your BCC. This is necessary because the_ADMIN volume is virtual, and is created each time the server starts. For this reason, you cannot assign eDirectory trustee rights to the `_ADMIN` volume. For instructions, see "Assigning Trustee Rights for the BCC Administrator User to the _ADMIN Volume" on page 46.

You must also ensure that the BCC Administrator user has Read, Write, Create, Erase, Modify, and File Scan access rights to the `sys:/tmp` directory on every node in your clusters. For instructions, see "Assigning Trustee Rights for the BCC Administrator User to the sys:\tmp Directory" on page 47.

6.3 BCC-Enabling Cluster Resources

Cluster resources can be configured for business continuity after they are created. Configuring a resource for business continuity consists of enabling that resource for business continuity, adding load and unload scripts search-and-replace data specific to the resource, and selecting peer clusters for the resource.

IMPORTANT: In a business continuity cluster, you should have only one NSS pool for each LUN that could be failed over to another cluster. This is necessary because in a business continuity cluster, entire LUNs fail over to other clusters, rather than individual pools, which fail over to other nodes within a cluster.

A cluster-enabled NSS pool must contain at least one volume before its cluster resource can be enabled for business continuity. You get an error message if you attempt to enable the resource for business continuity if its NSS pool does not contain a volume.

Also, if you have encrypted NSS volumes in your BCC, then all clusters in that BCC must be in the same eDirectory tree. If not, then the clusters in the other eDirectory tree cannot decrypt the NSS volumes. This rule applies to both NetWare and Linux BCCs.

6.3.1 Enabling a Cluster Resource for Business Continuity

Cluster resources must be enabled for business continuity on the primary cluster before they can be synchronized and appear as resources in the other clusters in the business continuity cluster. Enabling a cluster resource makes it possible for that cluster resource or cluster pool resource to be migrated to another cluster.

IMPORTANT: Although you can add search-and-replace data that is resource-specific after you enable a resource for business continuity, we recommend adding the search-and-replace data for the entire cluster before you enable resources for business continuity. See "Adding Search-and-Replace Values to the Resource Replacement Script" on page 72 for instructions on adding search-and-replace data for the entire cluster.

When you enable a resource for business continuity and that resource has been synchronized and appears in the other clusters, the preferred nodes for the other clusters are by default set to all nodes in the cluster. If you want to change the resource's preferred nodes for other clusters in your BCC, you must manually do it. Changes to the preferred nodes list in the primary cluster do not automatically replicate to the preferred nodes lists for other clusters in your BCC.

---

1. (Conditional) If you are creating a new cluster resource or cluster pool resource, follow the instructions for creating a cluster resource or cluster pool resource using iManager in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*, then continue with Step 2.
2. Enable a cluster resource or cluster pool resource for business continuity:
   - **2a** Start your Internet browser and enter the URL for iManager.

     The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
   - **2b** Specify your username and password, specify the tree where you want to log in, then click *Login*.
   - **2c** In the left column, click *Clusters*, then click the *Cluster Options* link.
   - **2d** Specify a cluster name, or browse and select one.
   - **2e** Select the desired cluster resource from the list of Cluster objects.
   - **2f** Click the *Details* link, then click the *Business Continuity* tab.
3. Ensure that the *Enable Business Continuity Features* check box is selected.
4. Continue with Step 1 in the Adding Resource Script Search-and-Replace Values section.

6.3.2 Adding Resource Script Search-and-Replace Values

If you did not previously add search-and-replace data specific to the entire cluster, you must now add it for this resource.

IMPORTANT: Adding resource script search-and-replace values for the entire cluster is recommended rather than adding those values for individual cluster resources. You should contact Novell Support prior to adding search-and-replace values for individual cluster resources.

To enable a resource for business continuity, certain values (such as IP addresses, DNS names, and tree names) specified in resource load and unload scripts need to be changed in corresponding resources in the other clusters. You need to add the search-and-replace strings that are used to transform cluster resource load and unload scripts from this cluster to another cluster.

The search-and-replace data you add is resource-specific, and it is not synchronized via Identity Manager between the clusters in the business continuity cluster.

To add resource script search-and-replace values specific to this resource:

1. In the *Resource Replacement Script* section of the page, click *New*.

   If a resource has already been configured for business continuity, you can click *Edit* to change existing search-and-replace values or click *Delete* to delete them.
2. Add the desired search-and-replace values, then click *OK*.

The search-and-replace values you specify here apply to only to the resource you are enabling for business continuity. If you want the search-and-replace values to apply to any or all cluster resources, add them to the entire cluster instead of just to a specific resource.

See "Adding Search-and-Replace Values to the Resource Replacement Script" on page 72 for more information on resource script search-and-replace values and adding those values to the entire cluster.

3 Do one of the following:

- If this is an existing cluster resource, continue with Step 1 in the Selecting Peer Clusters for the Resource section below.
- If you are creating a new cluster resource, click *Next*, then continue with Step 1 in the Selecting Peer Clusters for the Resource section.

You can select the Use Regular Expressions check box to use wildcard characters in your search-and-replace values. The following links provide information on regular expressions and wildcard characters:

- Regular Expressions (http://www.opengroup.org/onlinepubs/007908799/xbd/re.html)
- Regular-Expressions.info (http://www.regular-expressions.info/)
- Wikipedia (http://en.wikipedia.org/wiki/Regular_expression)
- oreilly.com (http://www.oreilly.com/catalog/regex/)

You can find additional information on regular expressions and wildcard characters by searching the Web.

IMPORTANT: If you change the resource-specific search-and-replace data after initially adding it, you must update the resource load or unload scripts in one of the other clusters by editing it and adding a space or a comment. This causes the script to be updated with the new search-and-replace data.

You could also update the IP address on the cluster protocols page in iManager to cause IP address search-and-replace values to be updated for both load and unload scripts. This might require you to go back and change the IP addresses specified in the resource load and unload scripts in the source cluster to their original values.

6.3.3 Selecting Peer Clusters for the Resource

Peer clusters are the other clusters that this cluster resource can be migrated to. The cluster or clusters that you select determine where the resource can be manually migrated. If you decide to migrate this resource to another cluster, you must migrate it to one of the clusters that has been selected.

1 Select the other clusters that this resource can be migrated to.

2 Do one of the following:

- If you are creating a new non-pool cluster resource that contains a Reiser or Ext3 file system, click *Finish*.
- If this is an existing non-pool cluster resource that contains a Reiser or Ext3 file system, click *Apply*.

- If you are creating a new cluster pool resource, click *Next*, then add the SAN management configuration information. For information, see “Adding SAN Management Configuration Information” on page 73.
- If this is an existing cluster pool resource, add the SAN management configuration information. For information, see “Adding SAN Management Configuration Information” on page 73.

6.3.4 Adding SAN Array Mapping Information

For information on adding SAN array mapping information, see “Adding SAN Management Configuration Information” on page 73.

Managing a Business Contiuity Cluster

7

This section can help you effectively manage a business continuity cluster with the Novell® Business Continuity Clustering software. It describes how to migrate cluster resources from one Novell Cluster Services™ cluster to another, to modify peer credentials for existing clusters, and to generate reports of the cluster configuration and status.

For information about using console commands to manage your business continuity cluster, see Appendix A, “Console Commands for BCC,” on page 121.

7.1 Migrating a Cluster Resource to a Peer Cluster

Although Novell Business Continuity Clustering provides an automatic failover feature that fails over resources between peer clusters, we recommend that you manually migrate cluster resources between the peer clusters instead. For information about configuring and using automatic failover for a business continuity cluster, Appendix C, “Setting Up Auto-Failover,” on page 131.

7.1.1 Understanding BCC Resource Migration

If the node where a resource is running fails, if the entire cluster fails, or if you just want to migrate the resource to another cluster, you can manually start the cluster resource on another cluster in the business continuity cluster. If the source cluster site fails, you must go to the destination cluster site to manually migrate or bring up resources at that site. Each resource starts on its preferred node on the destination cluster.

Migrating a pool resource to another cluster causes the following to happen:

1. If the source cluster can be contacted, the state of the resource is changed to offline.
2. The resource changes from primary to secondary on the source cluster.
3. Any SAN script that is associated with the pool resource is run.
4. On the destination cluster, the resource changes from secondary to primary so that it can be brought online.

A custom Perl script can be created for disk mapping on Fibre Channel SANs. The purpose of this script is to make the LUNs in the SAN available to the destination cluster. A reverse script is also created for testing purposes so pool resources can be migrated back to the source cluster.

5. The `cluster scan for new devices` command is executed on the destination cluster so that the cluster is aware of LUNs that are now available.
6. Resources are brought online and load on the most preferred node in the cluster.

TIP: You can use the `cluster migrate` command to start resources on nodes other than the preferred node on the destination cluster.

---

7. Resources appear as running and primary on the cluster where you have migrated them.

7.1.2 Migrating Cluster Resources between Clusters

WARNING: Do not migrate resources for a test failover if the peer (LAN) connection between the source and destination cluster is down. Possible disk problems and data corruption could occur. This warning does not apply if resources are migrated during an actual cluster site failure.

To manually migrate cluster resources from one cluster to another:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Clusters*, then click the *BCC Manager* link.
4. Specify a cluster name, or browse and select one.
5. Select one or more cluster resources, then click *BCC Migrate*.
6. Select the cluster where you want to migrate the selected resources, then click *OK*.

   The resources migrate to their preferred node on the destination cluster. If you select *Any Configured Peer* as the destination cluster, the Business Continuity Clustering software chooses a destination cluster for you. The destination cluster that is chosen is the first cluster that is up in the peer clusters list for this resource.

7.2 Changing Cluster Peer Credentials

You can change the credentials that are used by a one peer cluster to connect to another peer cluster. You might need to do this if the administrator username or password changes for any clusters in the business continuity cluster. To do this, you change the username and password for the administrative user who the selected cluster uses to connect to another selected peer cluster.

IMPORTANT: Make sure the new administrator username meets the requirements specified in Section 4.3, “Configuring a BCC Administrator User,” on page 45.

---

1. Start your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

IMPORTANT: In order to add or change cluster peer credentials, you must access iManager on a server that is in the same eDirectory™ tree as the cluster you are adding or changing peer credentials for.

---

2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Cluster Administration*, then click the *Management* link.
4. Specify a cluster name, or browse and select one.
5. Click *Connections* and select a peer cluster.
6. Edit the administrator username and password that the selected cluster will use to connect to the selected peer cluster, then click *OK*.

   When specifying a username, you do not include the Novell eDirectory context for the user name.

NOTE: If the business continuity cluster has clusters in multiple eDirectory trees, and you specify a common username and password, each eDirectory tree in the business continuity cluster must have the same username and password.

7.3 Viewing the Current Status of a Business Continuity Cluster

You can view the current status of your business continuity cluster by using either iManager or the server console of a cluster in the business continuity cluster.

- Section 7.3.1, “Using iManager to View the Cluster Status,” on page 83
- Section 7.3.2, “Using Console Commands to View the Cluster Status,” on page 84

7.3.1 Using iManager to View the Cluster Status

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Clusters*, then click the *BCC Manager* link.
4. Specify a cluster name, or browse and select one.
5. Use this page to see if all cluster peer connections are up or if one or more peer connections are down. You can also see the status of the BCC resources in the business continuity cluster.

7.3.2 Using Console Commands to View the Cluster Status

At the server console of a server in the business continuity cluster, enter the following commands to get different kinds of status information:

```
cluster view
cluster status
cluster connections
```

7.4 Generating a Cluster Report

You can generate a report for each cluster in the business continuity cluster to list information on a specific cluster, such as current cluster configuration, cluster nodes, and cluster resources. You can print or save the report by using your browser.

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Clusters*, then click the *Cluster Manager* link.
4. Specify a cluster name, or browse and select one.
5. Click the *Run Report* button.

7.5 Disabling Business Continuity Cluster Resources

After enabling a resource for business continuity, it is possible to disable it. You might want to disable BCC for a cluster resource in any of the following cases:

- You accidentally enabled the resource for business continuity.
- You no longer want the cluster resource to be able to fail over between peer clusters.
- You plan to delete the cluster resource.
- You plan to remove the peer cluster from the business continuity cluster. In this case, you must disable BCC for each cluster resource before you disable BCC for the cluster.

IMPORTANT: If you disable Business Continuity Clustering for a cluster by using either iManager or the `cluster disable` console command, the cluster resources in that cluster that that have been enabled for business continuity are automatically disabled for business continuity. If you re-enable Business Continuity Clustering for the cluster, you must again re-enable each of its cluster resources that you want to be enabled for business continuity.

This can be a time-consuming process if you have many cluster resources that are enabled for business continuity. For this reason, you should use caution when disabling Business Continuity Clustering for an entire cluster.

If BCC enabled resources need to be BCC disabled, remove the secondary peer clusters from the resource's assigned list, then disable BCC only from the primary cluster, either by using iManager or command line. Do not BCC disable the same resource from multiple peer clusters.

To disable BCC for a cluster resource:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
4. Specify the cluster name, or browse and select it.
5. Select the desired cluster resource from the list of Cluster objects
6. Click the *Details* link.
7. On the *Preferred Nodes* tab, remove the secondary peer clusters from the *Assigned* list, then disable BCC for the resource on the primary peer cluster.
   - **7a** Click the *Preferred Nodes* tab.
   - **7b** From the *Assigned Nodes* list, select the servers in the peer clusters you want to unassign from the resource, then click the left-arrow button to move the selected servers to the *Unassigned Nodes* list.

     The primary peer cluster and the node where the resource is running cannot be moved from the *Assigned* list to the *Unassigned* list.
   - **7c** Click *Apply* to save node assignment changes.
8. On the *Details* page, click the *Business Continuity* tab, deselect the *Enable Business Continuity Features* check box, then click *Apply*.
9. Wait for Identity Manager to synchronize the changes.

   This could take from 30 seconds to one minute, depending on your configuration.
10. Delete the Cluster Resource object on the clusters where you no longer want the resource to run.

7.6 Resolving Business Continuity Cluster Failures

There are several failure types associated with a business continuity cluster that you should be aware of. Understanding the failure types and knowing how to respond to each can help you more quickly recover a cluster. Some of the failure types and responses differ, depending on whether you have implemented SAN-based mirroring or host-based mirroring. Promoting or demoting LUNs is sometimes necessary when responding to certain types of failures.

NOTE: The terms *promote* and *demote* are used here in describing the process of changing LUNs to a state of primary or secondary, but your SAN vendor documentation might use different terms such as *mask* and *unmask*.

7.6.1 SAN-Based Mirroring Failure Types and Responses

SAN-based mirroring failure types and responses are described in the following sections:

Primary Cluster Fails but Primary SAN Does Not

This type of failure can be temporary (transient) or long-term. There should be an initial response and then a long-term response based on whether the failure is transient or long-term. The initial response is to restore the cluster to normal operations. The long-term response is total recovery from the failure.

Promote the secondary LUN to primary. Cluster resources load (and become primary on the second cluster). If the former primary SAN has not been demoted to secondary, you might need to demote it manually. The former primary SAN must be demoted to secondary before bringing cluster servers back up. Consult your SAN hardware documentation for instructions on demoting and promoting SANs. You can use the `cluster resetresources` console command to change resource states to offline and secondary.

Prior to bringing up the cluster servers, you must ensure that the SAN is in a state in which the cluster resources cannot come online and cause a divergence in data. Divergence in data occurs when connectivity between SANs has been lost and both clusters assert that they have ownership of their respective disks.

Primary Cluster and Primary SAN Both Fail

Bring the primary SAN back up and follow your SAN vendor’s instructions to remirror and, if necessary, promote the former primary SAN back to primary. Then bring up the former primary cluster servers and fail back the cluster resources.

Secondary Cluster Fails but Secondary SAN Does Not

No additional response is necessary for this failure other than recovering the secondary cluster. When you bring the secondary cluster back up, the LUNs are still in a secondary state to the primary SAN.

Secondary Cluster and Secondary SAN Both Fail

Bring the secondary SAN back up and follow your SAN vendor's instructions to remirror. When you bring the secondary cluster back up, the LUNs are still in a secondary state to the primary SAN.

Primary SAN Fails but Primary Cluster Does Not

When the primary SAN fails, the primary cluster also fails. Bring the primary SAN back up and follow your SAN vendor's instructions to remirror and, if necessary, promote the former primary SAN back to primary. You might need to demote the LUNs and resources to secondary on the primary SAN before bringing them back up. You can use the `cluster resetresources` console command to change resource states to offline and secondary. Bring up the former primary cluster servers and fail back resources.

Secondary SAN Fails but Secondary Cluster Does Not

When the secondary SAN fails, the secondary cluster also fails. Bring the secondary SAN back up and follow your SAN vendor's instructions to remirror. Then bring the secondary cluster back up. When you bring the secondary SAN and cluster back up, resources are still in a secondary state.

Intersite SAN Connectivity Is Lost

Recover your SANs first, then remirror from the good side to the bad side.

Intersite LAN Connectivity Is Lost

Users might not be able to access servers in the primary cluster but can possibly access servers in the secondary cluster. If both clusters are up, nothing additional is required. An error is displayed. Wait for connectivity to resume.

If you have configured the automatic failover feature, see Appendix C, "Setting Up Auto-Failover," on page 131.

7.6.2 Host-Based Mirroring Failure Types and Responses

Primary Cluster Fails but Primary SAN Does Not

Response for this failure is the same as for SAN-based mirroring described in Primary Cluster Fails but Primary SAN Does Not in Section 7.6.1, “SAN-Based Mirroring Failure Types and Responses,” on page 86. Do not disable MSAP (Multiple Server Activation Prevention), which is enabled by default.

Primary Cluster and Primary SAN Both Fail

Bring up your primary SAN or iSCSI target before bringing up your cluster servers. Then run the `Cluster Scan For New Devices` command from a secondary cluster server. Ensure that remirroring completes before bringing downed cluster servers back up.

If necessary, promote the former primary SAN back to primary. Then bring up the former primary cluster servers and fail back the cluster resources.

Secondary Cluster Fails but Secondary SAN Does Not

No additional response is necessary for this failure other than recovering the secondary cluster. When you bring the secondary cluster back up, the LUNs are still in a secondary state to the primary SAN.

Secondary Cluster and Secondary SAN Both Fail

Bring up your secondary SAN or iSCSI target before bringing up your cluster servers. Then run the `Cluster Scan For New Devices` command on a primary cluster server to ensure that remirroring takes place. When you bring the secondary cluster back up, the LUNs are still in a secondary state to the primary SAN.

Primary SAN Fails but Primary Cluster Does Not

If your primary SAN fails, all nodes in your primary cluster also fail. Bring up your primary SAN or iSCSI target and then bring up your cluster servers. Then run the `Cluster Scan For New Devices` command from a secondary cluster server. Ensure that remirroring completes before bringing downed cluster servers back up.

If necessary, promote the former primary SAN back to primary. You might need to demote the LUNs and resources to secondary on the primary SAN before bringing them back up. You can use the `cluster resetresources` console command to change resource states to offline and secondary. Bring up the former primary cluster servers and fail back resources.

Secondary SAN Fails but Secondary Cluster Does Not

Bring up your secondary SAN or iSCSI target before bringing up your cluster servers. Then run the `Cluster Scan For New Devices` command on a primary cluster server to ensure remirroring takes place. Then bring the secondary cluster back up. When you bring the secondary SAN and cluster back up, resources are still in a secondary state.

Intersite SAN Connectivity Is Lost

You must run the `Cluster Scan For New Devices` command on both clusters to ensure that remirroring takes place. Recover your SANs first, then remirror from the good side to the bad side.

Intersite LAN Connectivity Is Lost

Users might not be able to access servers in the primary cluster but can possibly access servers in the secondary cluster. If both clusters are up, nothing additional is required. An error is displayed. Wait for connectivity to resume.

If you have configured the automatic failover feature, see Appendix C, “Setting Up Auto-Failover,” on page 131.

Virtual IP Addresses

8

With the release of NetWare® 6.5, Novell® enhanced the TCP/IP stack to support virtual IP addresses. This feature is another high-availability offering that enables administrators to easily manage the name-to-IP address associations of business services. It complements the existing load balancing and fault tolerance features of the TCP/IP stack and enhances the availability of servers that reside on multiple subnets.

A virtual IP address is an IP address that is bound to a virtual Network Interface Card (NIC) and is driven by a new virtual driver named `vnic.lan`. As the name suggests, this virtual NIC is a purely virtual entity that has no physical hardware counterpart. A virtual NIC can be thought of as a conventional TCP/IP loopback interface with added external visibility. Virtual IP addresses can also be thought of as conventional loopback addresses with the 127.0.0.0 IP network constraint relaxed. A server with a virtual NIC and a virtual IP address acts as an interface to a virtual internal IP network that contains the server as the one and only host.

Regardless of their virtual nature, virtual IP addresses and virtual NICs behave like physical IP addresses and physical NICs, and they are similarly configured by using either the INETCFG server-based utility or the Novell Remote Manager Web-based utility.

8.1 Virtual IP Address Definitions and Characteristics

8.1.1 Definitions

Virtual driver: The `vnic.lan` driver provided by Novell.

Virtual board (NIC): Any board configured to use the virtual driver.

Virtual IP address: Any IP address that is bound to a virtual board.

Virtual IP network: The IP network that the virtual IP address is a part of. This is defined by the virtual IP address together with the IP network mask that it is configured with.

Host mask: The IP network mask consisting of all 1s - FF.FF.FF.FF (255.255.255.255).

Physical IP address: Any IP address that is not a virtual IP address. It is an IP address that is configured over a physical hardware NIC.

Physical IP network: An IP network that a physical IP address is a part of. A physical IP network identifies a logical IP network that is configured over a physical hardware wire.

8.1.2 Characteristics

Virtual IP addresses are unique in that they are bound to a virtual “ether” medium instead of to a “physical” network medium such as Ethernet. In other words, the virtual IP address space is different than the physical IP address space. As a result, virtual IP network numbers need to be different from physical IP network numbers. However, this mutual exclusivity of the IP address space for the physical and virtual networks doesn’t preclude the possibility of configuring multiple virtual IP networks in a single network domain.

8.2 Virtual IP Address Benefits

In spite of their simplicity, virtual IP addresses offer two main advantages over their physical counterparts:

- Section 8.2.1, “High Availability,” on page 92
- Section 8.2.2, “Unlimited Mobility,” on page 95
- Section 8.2.3, “Support for Host Mask,” on page 95
- Section 8.2.4, “Source Address Selection for Outbound Connections,” on page 95

These advantages exist because virtual IP addresses are purely virtual and are not bound to a physical network wire.

8.2.1 High Availability

If a virtual IP address is defined on a multihomed server with more than one physical NIC, a virtual IP address is a highly reachable IP address on the server when compared to any of the physical IP addresses. This is especially true in the event of server NIC failures. This assumes that the server is running a routing protocol and is advertising its “internal” virtual IP network—which only it knows about and can reach—to other network nodes.

Physical IP addresses might not be reachable because:

- TCP/IP protocols use link-based (network-based) addressing to identify network nodes. As a result, the routing protocols preferentially deliver a packet to the server through the network that the target IP address is part of.
- Dynamic routing protocols are extremely resilient to intermediate link and router failures, but they do not adapt well to failures of links at the last hop that ultimately delivers a packet to its destination.

  This is because the last hop link is typically a stub link that does not carry any routing heartbeats. Therefore, if one of the physical cards in a server fails, the server can become inaccessible, along with any service that it hosts on the corresponding physical IP address. This can occur in spite of the fact that the server is still up and running and can be reached through the other network card.

The virtual IP address feature circumvents this problem by creating a virtual IP network different from any of the existing physical IP networks. As a result, any packet that is destined for the virtual IP address is forced to use a virtual link as its last hop link. Because it is purely virtual, this last hop link can be expected to always be up. Also, because all other real links are forcibly made to act as intermediate links, their failures are easily handled by the dynamic routing protocols.

The following figure illustrates a multihomed server with all nodes running a dynamic routing protocol.

Figure 8-1 *Multihomed Server Running a Dynamic Routing Protocol*

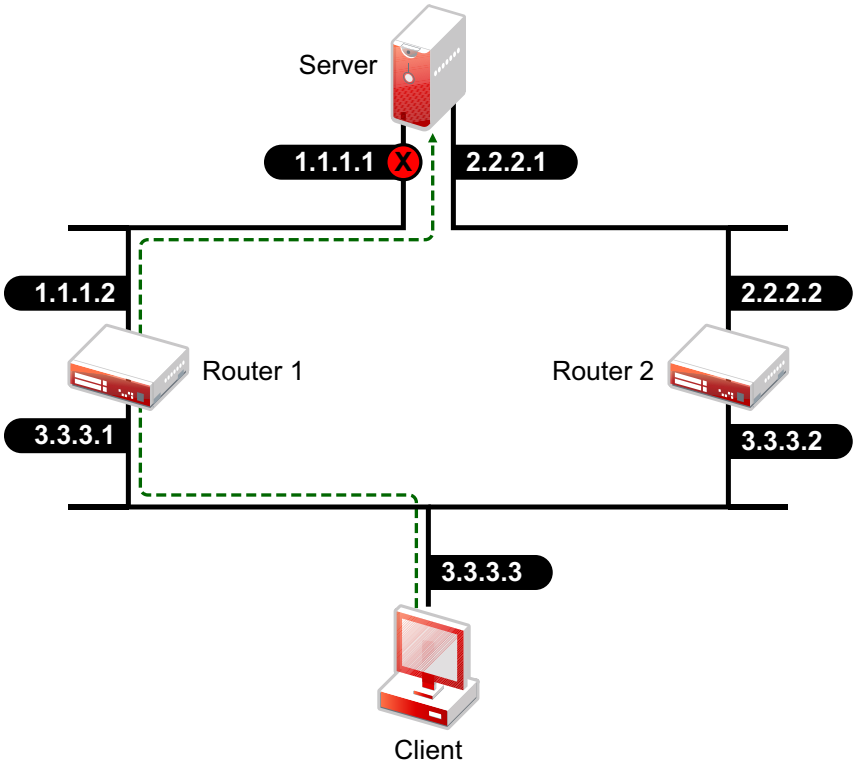

In this network, the server is a multihomed server hosting a critical network service. For simplicity, assume that all nodes are running some dynamic routing protocol.

If the client attempts to communicate with the server with the 1.1.1.1 IP address, it tries to reach the server through the nearest router, which is Router 1. If the 1.1.1.1 interface were to fail, Router 1 would continue to advertise reachability to the 1.0.0.0/FF.0.0.0 network and the client would continue to forward packets to Router 1. Thes undeliverable packets would ultimately be dropped by Router 1. Therefore, in spite of the fact that the service is still up and running and can be reached through the other active interface, it is rendered unreachable. In this scenario, a recovery would involve the ability of the client application to retry the alternate IP address 2.2.2.1 returned by the name server.

Consider the same scenario with the server configured with a virtual IP address and the client communicating with the virtual IP address instead of one of the server's real IP addresses, as shown in the following figure.

Figure 8-2 *Multihomed Server Using Virtual IP Addresses*

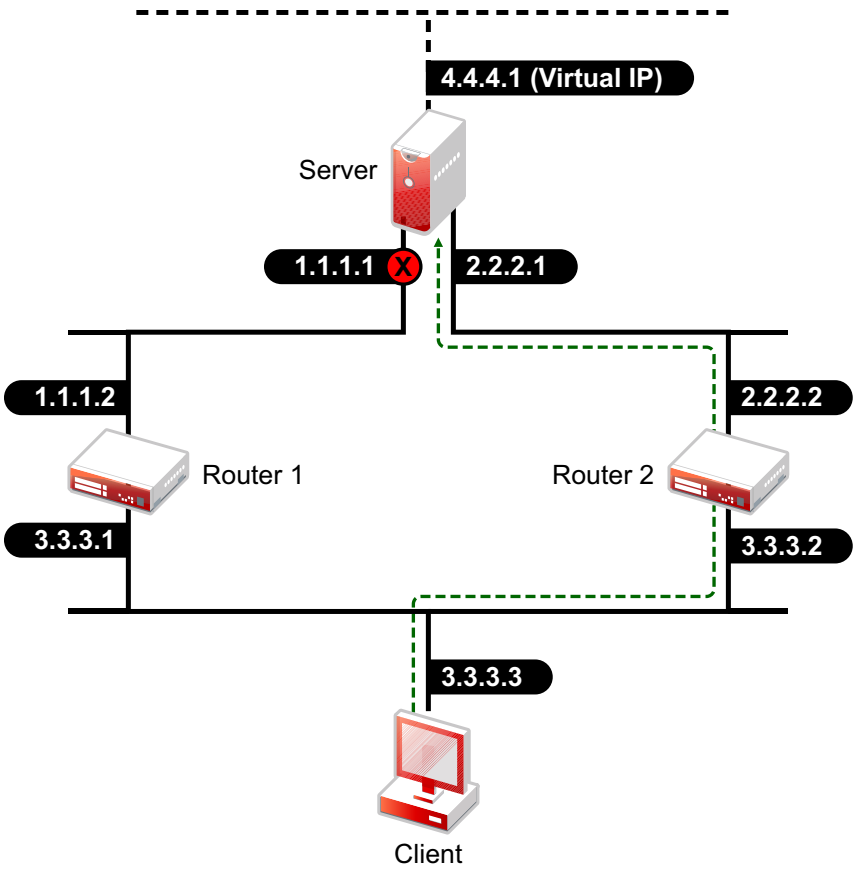

In this configuration, if the 1.1.1.1 interface were to fail, the client would ultimately learn the new route through Router 2 and would correctly forward packets to Router 2 instead of Router 1. Thus, despite physical interface failures, a virtual IP address on a multihomed server acts as an always-reachable IP address for the server.

Generally speaking, if a connection between two machines is established by using a virtual IP address as the end-point address at either end, the connection is resilient to interface failures at either end.

There are two important side effects that directly follow from the highly reachable nature of virtual IP addresses:

- They completely and uniquely identify a multihomed server

  A multihomed server with a virtual IP address no longer needs to carry multiple DNS entries for its name in the naming system.

- They significantly enhance the LAN redundancy inherent in a multihomed server

  If one of the subnets that a server interfaces to fails completely or is taken out of service for maintenance, the routing protocols reroute the packets addressed to the virtual IP address through one of the other active subnets.

The resilience against interface failures provided by virtual IP addresses depends on the fault resilience provided by the dynamic routing protocols, as well as on fault recovery features such as retransmissions built into the application logic.

8.2.2 Unlimited Mobility

Unlike physical IP addresses, which are limited in their mobility, virtual IP addresses are highly mobile. The degree of mobility is determined by the number of servers that an IP address on a specific server could be moved to. In other words, if you choose a physical IP address as an IP address of a network resource, you are limiting the set of potential servers to which this resource could transparently fail over to.

If you choose a virtual IP address, the set of servers that the resource could be transparently moved to is potentially unlimited. This is because of the nature of virtual IP addresses; they are not bound to a physical wire and, as a result, they carry their virtual network to wherever they are moved. Again, there is an implicit assumption that the location of a virtual IP address is advertised to the owning server through some routing protocol. The ability to move an IP address across different machines becomes particularly important when you need to transparently move or fail over a network resource that is identified by an IP address (which could be a shared volume or a mission-critical service) to another server on another network.

This unlimited mobility of virtual IP addresses is an advantage to network administrators, offering more ease of manageability and greatly minimizing network reorganization overhead. For network administrators, shuffling services between different IP networks is the rule rather than the exception. The need often arises to move a machine hosting a particular service to some other IP network, or to move a service hosted on a particular machine to be rehosted on some other machine connected to a different IP network. If the service is hosted on a physical IP address, accommodating these changes involves rehosting the service on a different IP address pulled out from the new network, and appropriately changing the DNS entry for the service to point to the new IP address. However, if the service is hosted on a virtual IP address, the necessity of changing the DNS entries for the service is eliminated.

8.2.3 Support for Host Mask

Virtual boards support configuring virtual IP addresses with a host mask. This results in a single address being used rather than an entire subnet. See Section 8.3, “Reducing the Consumption of Additional IP Addresses,” on page 96.

8.2.4 Source Address Selection for Outbound Connections

Full resilience of connections to interface failures can be ensured only when the connections are established between machines through using virtual IP addresses as end point addresses. This means an application that initiates outbound connections to a virtual IP address should also use a virtual IP address as its local end point address.

This isn’t difficult if the application binds its local socket end point address with a virtual IP address. However, there are some legacy applications that bind their sockets to a wildcard address (such as 0.0.0.0). When these applications initiate an outbound connection to other machines, TCP/IP chooses the outbound interface’s IP address as the local socket end point address. In order for these legacy applications to take advantage of the fault resilience provided by the virtual IP address feature, the default source address selection behavior of TCP/IP has been enhanced to accommodate the use of a virtual IP address as the source IP address. As a result, whenever a TCP or UDP application initiates an outbound connection with a wildcard source IP address, TCP/IP chooses the first bound virtual IP address as the source IP address for the connection.

This enhanced source address selection feature can be enabled or disabled globally as well as on a per-interface basis. This feature is enabled by default on all interfaces.

8.3 Reducing the Consumption of Additional IP Addresses

In any network environment, one of the first obstacles is how clients locate and connect to the services. A business continuity cluster can exacerbate this problem because services can migrate to nodes on a completely different network segment. Although there are many potential solutions to this problem, such as DNS and SLP, none of them offers the simplicity and elegance of virtual IP addresses. With virtual IP addresses, the IP address of the service can follow the service from node to node in a single cluster, as well as from node to node in separate, distinct clusters. This makes the client reconnection problem trivial; the client just waits for the new route information to be propagated to the routers on the network. No manual steps are required, such as modifying a DNS server.

The only drawback in using virtual IP addresses is the consumption of additional IP addresses. This constraint stems from the requirement that virtual IP network addresses must be different from all other real IP network addresses. Although this constraint is not particularly severe in enterprises that use private addressing (where the IP address space is potentially large), it could become limiting in organizations that do not use private addresses.

To use a virtual IP address in a business continuity cluster, we recommend using a host mask. To understand why, consider the fact that each service in a clustered environment must have its own unique IP address or, a unique virtual IP address. Furthermore, consider that each virtual IP address belongs to a virtual IP network whose route is being advertised by a single node within a cluster. Because Novell Cluster Services™ can migrate a service and its virtual IP address from one node to another, the virtual IP network must migrate to the same node as the service. If multiple virtual IP addresses belong to a given virtual IP network, one of two events must occur:

- All services associated with the virtual IP addresses on a given virtual IP network must fail over together.
- The virtual IP addresses on a given virtual IP network must go unused, thereby wasting a portion of the available address space.

Neither of these situations is desirable. Fortunately, the use of host masks remedies both.

In enterprises that use fixed-length subnetting together with a dynamic routing protocol like RIP-I, each virtual IP address could consume a large number of host IP addresses. One way to circumvent this problem is to configure a virtual IP address with a host mask of all 1s (that is, FF.FF.FF.FF), thereby consuming only one host IP address. Of course, the viability of this option depends on the ability of the RIP-I routers on the network to recognize and honor the advertised host routes.

In autonomous systems that use variable-length subnet masking (VLSM) together with routing protocols like RIP-II or OSPF, the consumption of additional IP addresses is not a major problem. You could simply configure a virtual IP address with an IP network mask as large as possible (including a host mask of all 1s), thereby limiting the number of addresses consumed by the virtual IP address space.

8.4 Configuring Virtual IP Addresses

The routers in a virtual IP address configuration must be running the RIP I or RIP II protocols. For a business continuity cluster, RIP II is the preferred protocol and should be used whenever possible. In NetWare, this can be accomplished by configuring the NetWare RIP Bind Options to use RIP I and RIP II, or RIP II only. Also, the command `SET RIP2 AGGREGATION OVERRIDE=ON` must be added to the `autoexec.ncf` file of any NetWare routers in this configuration.

After the appropriate virtual IP addresses and host masks have been determined, you can enable virtual IP addresses in a business continuity cluster by using the following process:

1. The `autoexec.ncf` file on each node in both clusters must be modified to add the following two lines. The first line loads the virtual driver and creates a virtual board named VNIC. The second line disables RIP 2 route aggregation on the cluster nodes.

   ```
   LOAD VNIC NAME=VNIC
   SET RIP2 AGGREGATION OVERRIDE=ON
   ```

2. The command to bind a virtual IP address for the service must be added to the cluster resource load script.

   The following is an example of a cluster resource load script for a standard NetWare volume called `HOMES`. This example uses host masks and assumes the virtual board has been named VNIC. Notice that the command to add a secondary IP address has been replaced with the `BIND IP VNIC Mask=255.255.255.255 Address=10.1.1.1` command, which binds the virtual IP address 10.1.1.1 to the virtual board.

   ```
   nss /poolactivate=HOMES
   mount HOMES VOLID=254
   CLUSTER CVSBIND ADD BCC_HOMES_SERVER 10.1.1.1
   NUDP ADD BCC_HOMES_SERVER 10.1.1.1
   BIND IP VNIC Mask=255.255.255.255 Address=10.1.1.1
   ```

3. The command to unbind the virtual IP address must be added to the cluster resource unload script.

   The following is the matching cluster resource unload script for the same NetWare volume discussed above. Notice the command to delete the secondary IP address has been replaced with the `UNBIND IP VNIC Address=10.1.1.1` command, which unbinds the virtual IP address 10.1.1.1 from the virtual board.

   ```
   UNBIND IP VNIC Address=10.1.1.1
   CLUSTER CVSBIND DEL BCC_HOMES_SERVER 10.1.1.1
   NUDP DEL BCC_HOMES_SERVER 10.1.1.1
   nss /pooldeactivate=HOMES /overridetype=question
   ```

4. If the cluster resource is a clustered-enabled pool or volume, the IP address of that resource needs to be changed to the virtual IP address. You can do this by using ConsoleOne®, Novell Remote Manager, or iManager. This change is not needed for any non-volume cluster resources like DHCP.

8.4.1 Displaying Bound Virtual IP Addresses

To verify that a virtual IP address is bound, enter `display secondary ipaddress` at the server console of the cluster server where the virtual IP address is assigned. This displays all bound virtual IP addresses. A maximum of 256 virtual IP addresses can be bound.

Troubleshooting Business Continuity Clustering 1.1

9

This section contains the following topics to help you troubleshoot Novell® Business Continuity Clustering 1.1.

9.1 Cluster Connection States

The following table identifies the different cluster connection states and gives descriptions and possible actions for each state.

Table 9-1 *BCC Connection States*

BCC Connection State	Number	Description	Possible Actions
Normal	0	The connections between clusters are functioning normally.	None required.
Authenticating	1	BCC is in the process of authenticating to a peer cluster.	Wait until the authentication process is finished.
Invalid Credentials	2	You entered the wrong username or password for the selected peer cluster.	Enter the correct username and password that this cluster will use to connect to the selected peer cluster.
Cannot Connect	3	This cluster cannot connect to the selected peer cluster.	Ping the peer cluster to see if it is up and reachable. Ensure that BCC is running on the peer cluster and that Novell Cluster Services™ is running on the servers in the peer cluster. Ensure that OpenWBEM is running on the peer cluster. Ensure that a firewall is not preventing access on OpenWBEM ports 5988 and 5989. Ensure that the Admin file system is running. To do this, see if the `_admin` volume is mounted on NetWare.
Not Authorized	4	The connected user does not have sufficient rights for permissions.	Assign the appropriate trustee rights to the user who will manage your BCC. For information, see "Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects" on page 46.
Connection Unknown	5	The connection state between clusters is unknown.	This connection state might be caused by any number of problems, including a severed cable or link problems between geographic sites.

The connection state numbers are recorded in a log file that you can use to view connection and status changes for BCC.

The default path to the log file on Linux is `/var/log/messages`. The administrator might have changed this path from the default. Search for BCCD to view BCC related messages and entries in the log file.

To view the log file on NetWare®:

1. At the NetWare server console, enter `log +copy syslog`.

2 Using an editor, open the file that is referenced in the message that appears.

   You can get additional information on how to use the log file by entering `help log` at the NetWare server console.

9.2 Driver Ports

If your Identity Manager driver or drivers won't start, check for a port number conflict. Identity Manager driver port numbers must not be the same as other driver port numbers in the cluster or ports being used by other services such as Apache.

To check driver port numbers:

1 Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2 Specify your username and password, specify the tree where you want to log in, then click *Login*.

3 In the left column, click *Identity Manager*, then click the *Identity Manager Overview* link.

4 Select *Search Entire Tree*, then click *Search*.

5 Select the driver you want to check by clicking the red *Cluster Sync* icon or the blue *User Sync* icon.

6 Click the red or blue icon again, then click the *Identity Manager* tab (if it is not already selected).

7 In the *Authentication context* field, view and if necessary change the port numbers next to the IP address.

   For example, the *Authentication context* field might contain a value similar to 10.1.1.12:2003:2003. In this example, the first port number (2003) is the port number for the corresponding Identity Manager driver on the cluster this cluster is synchronizing with. The second port number (2003) is the port number for the Identity Manager driver on this cluster.

   These port numbers should be the same, but should not be the same as the port numbers for other Identity Manager drivers on either this or the remote cluster.

8 If you change the port numbers, restart the driver by clicking the upper right corner of the *Cluster Sync* or *User Sync* icon (whichever you have chosen), then click *Restart driver*.

9 If you changed the port number in Step 7 above, change the port numbers to be the same for the corresponding driver in the other cluster.

   You can do this by repeating Step 1 through Step 8 for the Identity Manager driver on the other cluster.

9.3 Excluded Users

If certain users do not synchronize between clusters, it is possible that those users are in the excluded users list.

NOTE: The eDirectory™ Admin user should never be synchronized between clusters.

To see the excluded users list:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager*, then click the *Identity Manager Overview* link.
4. Select *Search Entire Tree*, then click *Search*.
5. Select the user synchronization driver you want to check by clicking the blue *User Sync* icon.

   This is not necessary for the cluster synchronization driver.
6. Click the blue icon again, then click the *Identity Manager* tab if it is not already selected.
7. Click *Excluded Users*, and view, add, or remove users as desired.

9.4 Security Equivalent User

If resources or peers don't appear in other clusters in your BCC, it is possible that either a cluster or user synchronization driver is not security equivalent to a user with administrative rights to the cluster.

NOTE: Rather than using the eDirectory Admin user to administer your BCC, you should consider creating another user with sufficient rights to the appropriate contexts in your eDirectory tree to manage your BCC.

The IDM Driver object must have sufficient rights to create, modify, and delete objects and attributes in the following containers:

- The Identity Manager driver set container.
- The container where the Cluster object resides.
- The container where the Server objects reside.

  If server objects reside in multiple containers, this must be a container high enough in the tree to be above all containers that contain server objects. The best practice is to have all server objects in one container.
- The container where the cluster pool and volume objects are placed when they are synchronized to this cluster. This container is sometimes referred to as the landing zone. The NCP™ server objects for the virtual server of a business-continuity-enabled resource are also placed in the landing zone.

To make the Cluster Resource Synchronization driver or User Object Synchronization driver security equivalent to a user with administrative rights:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2 Specify your username and password, specify the tree where you want to log in, then click *Login*.

3 In the left column, click *Identity Manager*, then click the *Identity Manager Overview* link.

4 Choose *Search Entire Tree*, then click *Search*.

5 Select the driver you want to check by clicking the red *Cluster Sync* icon or the blue *User Sync* icon.

6 Click the red or blue icon again, then click the *Identity Manager* tab if it is not already selected.

7 Click *Security Equals*, and view or add a security equivalent user as needed.

8 Repeat Step 5 through Step 7 for the other drivers in your BCC.

You must also ensure that the BCC Administrator user has Read, Write, Create, Erase, Modify, and File Scan access rights to the `sys:/tmp` directory on every node in your NetWare clusters.

For Linux, ensure that the BCC Administrator user is a LUM-enabled user. To LUM-enable a user, see "Managing User and Group Objects in eDirectory" (http://www.novell.com/documentation/oes/lumadgd/data/aeucqum.html) in the *Novell Linux User Management Technology Guide*.

NOTE: For NetWare, if you are concerned about denial of service attacks with the BCC Administrator user, you can set a quota of 5 MB for that user. This can prevent the BCC Administrator user from filling the `sys:` volume by copying an excessive number of files to the `sys:/tmp` directory

---

9.5 Certificates

If SSL certificates are not present or have not been created, Identity Manager drivers might not start or function properly. Novell recommends using SSL certificates for encryption and security.

NOTE: You should create or use a different certificate than the default (dummy) certificate (BCC Cluster Sync KMO) that is included with BCC.

See "Creating SSL Certificates" on page 68 for more information on creating SSL certificates for BCC.

9.6 Clusters Cannot Communicate

If the clusters in your BCC cannot communicate with each other, it is possible that the User object you are using to administer your BCC does not have sufficient rights to the Cluster objects for each cluster. To resolve this problem, ensure that the BCC Administrator user is a trustee of the Cluster objects and has at least Read and Write rights to the All Attributes Rights property.

1 Start your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2 Specify your username and password, specify the tree where you want to log in, then click *Login*.

3 In the left column, click *Rights*, then click the *Modify Trustees* link.

4 Specify the Cluster object name, or browse and select it, then click *OK*.

5 If the BCC Administrator user is not listed as a trustee, click the *Add Trustee* button, browse and select the User object, then click OK.

6 Click *Assigned Rights* for the BCC Administrator user, then ensure that the *Read* and *Write* check boxes are selected for the *All Attributes Rights* property.

7 Click *Done* to save your changes.

8 Repeat Step 3 through Step 7 for the other Cluster objects in your BCC.

9.7 BCC Startup Flags

There are three optional flags that can be used by BCC during startup. On NetWare, edit the `sys:\system\ldbcc.ncf` file and change the `load bccd.nlm` line to `load bccd.nlm -flags`. Replace *flags* with any combination of v and/or t. This could include v, vt, t, or tv.

Table 9-2 *Optional Startup Flags*

Startup Flag	Description
d	On Linux, this flag keeps the bccd from forking (keeps the process in foreground) and log messages are printed to the running terminal screen (stdout) along with the normal syslog. This flag is not used on NetWare.
v	On both Linux and NetWare, this flag turns on more verbose logging.
t	On both Linux and NetWare, this flag turns on tracing. With tracing turned on, certain sections of code that fail will report a message containing the condition that failed along with a file and line number in the code indicating where the condition failed. This is helpful for reporting problems to Novell Support.

9.8 Problems with Installing BCC on NetWare

Occasional problems might exist when installing BCC software on NetWare cluster servers. For example, you might experience problems expanding an Organizational Unit (OU) when browsing the eDirectory tree during the BCC installation.

To resolve this and similar problems, rename the `c:\program files\common files\novell` directory on the Windows machine where the installation is being run and restart the Business Continuity Clustering 1.1 installation program.

9.9 Identity Manager Drivers for Cluster Synchronization Do Not Start

If the Identity Manager drivers for cluster synchronization do not start, the problem might be caused by one of the following conditions:

- A certificate has not been created. For information, see “Creating SSL Certificates” on page 68.
- The ports used by the driver are not unique and available.

Each eDirectory driver must listen on a different port number. To specify the port number, access the driver properties in iManager and specify the appropriate port number in the IP address field. See Section 6.1, “Configuring Identity Manager Drivers for the Business Continuity Cluster,” on page 65 for more information.

The format for specifying the port number in the IP address field is *remote IP:remote port:local port*. For example, you could specify something similar to 10.1.1.1:2001:2001. If it is not being used for other drivers, port 2001 can be used for the User object driver and port 2002 for the Cluster object driver.

- The driver has been disabled.

  Click the red icon for the driver on the Identity Manager Driver Overview page. You can enable the driver by using the radio buttons in the Driver Startup section of the page that displays.

  Selecting the Auto Start option is recommended.

- Unknown communications problems.

  See Section 9.11, “Tracing Identity Manager Communications,” on page 106.

- -670 errors in the DSTrace logs. See TID # 10090395 (http://support.novell.com/docs/Tids/Solutions/10090395.html)

  This is commonly caused by `rconag6.nlm` loading before SAS, which causes problems when Identity Manager tries to load. The load order is sometimes changed by the installation of a product that changes the order of the lines within the `autoexec.ncf` file.

9.10 Identity Manager Drivers Do Not Synchronize Objects from One Cluster to Another

If objects are not synchronizing between clusters, the problem might be caused by one of the following conditions:

- The drivers are not running.
- A driver is not security equivalent to an object with the necessary rights in the tree.
- You have underscores and spaces in object names.

  eDirectory interprets underscores and spaces as the same character. For example, if you have a cluster template named iFolder Server and you try to synchronize a resource named iFolder_Server, the synchronization fails. This is because the underscore character is mapped to a space. eDirectory returns an error that the entry already exists.

- The eDirectory partition on the Identity Manager node is incorrect.

  This partition must contain the cluster container, the DriverSet, the Landing Zone OU, and the server containers (Virtual NCP™ Servers, Volumes, and Pools).

- The drivers are not communicating on the same port.

  For example, if the driver on Cluster A is listening on port 2002, the driver on Cluster B must bind to port 2002 on Cluster A in order for the driver communication to work properly.

  The format for specifying the port number in the IP address field is *remote IP:remote port:local port*. For example, you could specify something similar to 10.1.1.1:2001:2001. If it is not being used for other drivers, port 2001 can be used for the User object driver and port 2002 for the Cluster object driver.

- See Tracing Identity Manager Communications below.

9.11 Tracing Identity Manager Communications

DSTrace is used to trace Identity Manager communications. In a BCC, it is generally best to trace both sides of the communication channel (both drivers).

To trace the communications for the BCC-specific Identity Manager drivers on a NetWare BCC:

1. Modify two attributes on both DriverSet objects:
   - **1a** Start your Internet browser and enter the URL for iManager.

     The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
   - **1b** Specify your username and password, specify the tree where you want to log in, then click *Login*.
   - **1c** Click the *View Objects* button at the top of the iManager page.
   - **1d** In the left column, browse to and right-click the desired DriverSet object, then select *Modify Object*.
   - **1e** Click the *General* tab, and in the list of valued attributes, click *Identity Manager-DriverTraceLevel*, then click *Edit*.
   - **1f** Ensure that the Trace Level is set to 4, then click *OK*.
   - **1g** Repeat Step 1e and Step 1f for the Identity Manager-XSLTraceLevel attribute, also setting the trace level to 4.
   - **1h** Repeat Step 1d through Step 1g for the other driver sets you want to trace.
2. At the NetWare server console, load DSTrace by entering `dstrace`.
3. Configure DSTrace by entering `dstrace inline -all +dvrs +dxml` at the NetWare server console.
4. Enable DSTrace by again entering `dstrace` at the NetWare server console.
5. Run the desired actions for the information you want traced.
6. Disable DSTrace by entering `dstrace off` at the NetWare server console.

The trace file is located at `sys:/system/dstrace.log`. You might want to delete this file before starting a trace so the events logged in the file are specific to the actions you are tracing.

9.12 Peer Cluster Communication Not Working

If BCC communication between peer clusters is not functioning, the problem might be caused by one of the following conditions:

- The credentials for the remote cluster have not been set.

  You cannot use iManager on a server in one tree to set credentials for a BCC cluster in another tree. This is because BCC and iManager use the tree key to encrypt the credentials. Setting credentials by using iManager in a different tree uses an invalid tree encryption.
- LIBC has not been updated. See OES SP2, NW6.5 SP5 Update 1: TID # 2974185 (http://support.novell.com/docs/Readmes/InfoDocument/2974185.html).
- A firewall is blocking port 5988 or 5989 (CIM).

9.13 Administration of Peer Clusters Not Functional

This problem is normally caused by the BCC Administrator user not having file system rights to the cluster administration files. See "Verifying BCC Administrator User Trustee Rights and Credentials" on page 75.

9.14 A Resource Does Not Migrate to a Peer Cluster

If you cannot migrate a resource from one cluster to a peer cluster, the problem might be caused by one of the following conditions:

- The resource has not been BCC-enabled.
- Remote clusters cannot communicate.

  See Section 9.12, "Peer Cluster Communication Not Working," on page 106.
- Syslog shows error 1019 (NSMI error).

  There could be a partial script in `sys:/tmp`. The scripts are named `NSMIT-######.tmp`. If there is a partial script, NSMI is experiencing an error and not communicating with the SAN to make disks visible. One reason for this is that the script might have a single % character in it that needs to be an escape (%% instead of %). For example, hashes in Perl need to have escape characters.

9.15 A Resource Cannot Be Brought Online

If you cannot bring a BCC-enabled resource online, it is possible the resource might be set as secondary. If the NCS:BCC State attribute is equal to 1, the resource is set to secondary and cannot be brought online.

On the resource object, change the NCS:BCC State attribute to 0 to set the resource to the primary state. Also, increment the NCS:Revision attribute one number so that Novell Cluster Services recognizes that the resource properties have been updated. See Step 1 on page 106 for an example of how to modify object attributes.

9.16 Dumping Syslog on NetWare

The command `log +dump syslog` sends the output of the syslog to the console screen. This command has limited use because only the last few entries of the log can be viewed.

You can use the `log +copy syslog` command to copy the syslog to a file and then use the NetWare Edit utility to view it. The output file that syslog is copied to is displayed after the command is executed.

Enter `log help` at the NetWare server console to get additional information on syslog.

9.17 Slow Failovers

Failovers might be slow because the resource is slow to go offline on the source cluster. This can happen if the failover occurs during a time when file I/O is taking place on the cluster.

To resolve this problem, edit the resource unload script and change the `NUDP DEL name ip` line to `NUDP ODEL name ip`. Unlike the `DEL` command, the `ODEL` command does not wait for all NCP connections to close. This makes it much faster.

NOTE: The cluster resource must be brought offline and then back online for changes to the unload script to take effect. Client data might be lost if clients are accessing the resource when it is brought offline.

9.18 Resource Script Search-and-Replace Functions Do Not Work

If resource script search-and-replace functions are not working, the problem might be caused by one of the following conditions:

- You did not click the *Apply* button on the Properties page. Clicking *OK* when entering the scripts does not apply the changes to the directory.
- You added the search-and-replace values to the resource instead of to the cluster.

  The search-and-replace values to apply to a specific resource instead of all resources in the cluster.
- If you are testing search-and-replace functionality, you might have made the changes too rapidly.

  Identity Manager merges all changes into one, so if you quickly add a change and then delete it, Identity Manager might view it as no change. You should make a change and verify that the change has synchronized with the other cluster before deleting it.

9.19 Virtual NCP Server IP Addresses Won't Change

If the IP address for a virtual (NCP) server does not change properly, the problem might be caused by one of the following conditions:

- The IP address has been changed only on the load and unload script pages.

  You must also change the IP address on the protocols page for the virtual server. Changing the IP address on the protocols page causes the load and unload scripts to automatically update.
- The virtual server has an extra IP address.

  IP address changes should always be made on the Protocols page of the iManager cluster plug-in, not in load and unload scripts.

  This is the only way to change the IP address on the virtual NCP server object in eDirectory.

  On NetWare, you can also use the General tab in ConsoleOne® to view the IP addresses for the virtual server. If there are extra IP addresses, do the following:

  1. In ConsoleOne, click *Page Options* and disable the *General* tab.
  2. Click *Cancel* to exit the properties dialog box.
  3. Access the properties dialog box again.
  4. Click the *Other* tab and delete the extra IP addresses.

The attribute is Network Address. Do not delete the entire attribute, just the values for the extra IP addresses.

5. Click *Page Options* and enable the *General* tab.
6. Click *Apply* to save your changes.

9.20 The IP Address, Virtual Server DN, or Pool Name Does Not Appear on the iManager Cluster Configuration Page

You might see a DSML read error if you select properties for a Cluster object and then click the *BCC* tab.

The eDirectory pointers for the cluster resource are either missing or are invalid. The following list shows the required attributes. The format for the list is:

Object in the directory (Object Class Name)
Attribute Name -->The object in the directory the attribute points to
Attribute Name -->The object in the directory the attribute points to

Clustered Volume (Object Class “NCS:Volume Resource”)
NCS:NCP Server -->Virtual NCP Servers
NCS:Volumes-->All Volumes

Virtual NCP Server (Object Class “NCP Server”)
Resource-->Cluster Resource
NCS:NetWare Cluster-->Cluster Object
NCS:Volumes-->All Volumes

Volume (Object Class “Volume”)
nssfsPool-->Pool Object
Host Server-->Virtual NCP Server

Pool (Object Class “nssfsPool”)
Host Server-->Virtual NCP Server

9.21 Blank Error String iManager Error Appears While Bringing a Resource Online

If you get an error in iManager with a blank error string (no text appears with the error message) while attempting to bring a resource online, it is possible that Novell Cluster Services views the resource as secondary even though BCC has changed the resource to primary and iManager shows the resource as primary.

To resolve this problem, make a change to the cluster properties to cause the NCS:Revision attribute to increment. You could add a comment to the resource load script to cause this to happen.

9.22 Mapping Drives in Login Scripts

Consider the following when mapping drives in login scripts in a BCC:

- Using an FDN (such as map s:=BCCP_Cluster_HOMES.servers.:shared) to a cluster-enabled volume has been tested and does not work.

When the resource fails over to the secondary cluster, the DN does not resolve to a server/volume that is online. This causes the `map` command to fail.

- Using the *`SLP Server Name`*`/VOL:shared` syntax has been tested and works.

  SLP Server Name is the name being advertised in SLP as specified in the resource load script. This method requires a client reboot.

- See TID 10057730 (http://support.novell.com/docs/Tids/Solutions/10057730.html) for information on modifying the server cache Time To Live (TTL) value on the Novell Client™.

9.23 Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable

Consider the following when mapping drives in login scripts in a BCC:

- Using the *%HOME_DIRECTORY* variable (such as map u:=%HOME_DIRECTORY) has been tested and does not work.

  When you fail the resource over to the secondary cluster, the %HOME_DIRECTORY variable still resolves to the old volume object, and the `map` command fails.

- Using a temporary environment variable has been tested and does not work.

  For example:

  ```
  set FOO=%HOME_DIRECTORY
  MAP u:=%FOO
  ```

- Using a false volume object along with ICE and LDIF has been tested and works.
  - Create a new volume object that references the real volume object.

    The Host Server attribute must point to the virtual NCP server in the primary cluster, and the Host Resource Name attribute must specify the name of the volume. This new volume object can be referred to as a volume reference.

  - All User objects must be modified to have their Home Directory attribute reference the new volume object (volume reference).
  - Use LDIF and ICE in the NSMI script (SAN Array Mapping Information) area.

    This script modifies the new volume reference object and updates the Host Server attribute to point to the virtual NCP server in the secondary cluster.

    NOTE: Using LDIF/ICE prevents you from using the NSMI script to control the SAN. If you want to use LDIF/ICE and the NSMI script, you must have two NCF files: one for the SAN, and one for LDIF/ICE. The NSMI script must then call each NCF file separately.

    See TID 10057730 (http://support.novell.com/docs/Tids/Solutions/10057730.html) for information on modifying the server cache Time To Live (TTL) value on the Novell Client.

  A sample NSMI script is included below:

```
#!ICE -b -D LDAP -d cn=root,ou=servers,o=lab -w novell -S LDIF -f

#@ -s0 -w20

version: 1

dn: cn=HOMES_REF, ou=servers,o=lab

changetype: modify

replace: hostServer

hostServer:

cn=BCC_CLUSTER_HOMES_SERVER,ou=From_BCCP,ou=servers,o=lab
```

The first line in the sample script instructs NSMI to run the ICE utility.

- The -b parameter automatically closes the ICE window.
- The -d parameter is the administrator DN that is used to modify eDirectory.
- The -w parameter is the password.
- There must be a trailing space after the -f parameter.

The second line in the sample script includes NSMI-specific options.

- -s0 causes NSMI to not wait for a signal file.
- -w20 causes NSMI to wait 20 seconds before proceeding

  Failure to add the wait causes the temporary LDIF file to be deleted before ICE can read it. This causes ICE to fail.

9.24 BCC Error Codes

The following table lists the different BCC error codes by number and gives a brief description for each error code.

Table 9-3 *BCC Error Codes and Messages*

Error Code Number	Message
1000	Unknown error.
1001	Received XML is invalid.
1002	The object pointers in eDirectory for the given cluster resource are invalid.
1003	The referenced object is not a valid NCS/BCC object.
1004	The referenced cluster resource is in an invalid state.
1005	The specified resource or cluster is not enabled for Business Continuity.
1006	An invalid parameter was passed to the BCC API.
1007	Attempt to allocate memory failed.
1008	Attempt to communicate with the BCC VFS system failed.

Error Code Number	Message
1009	The size of the specified buffer is not large enough.
1010	Error performing a DSML read.
1011	Error performing a DSML modify.
1012	Operation not supported.
1013	Error obtaining lock on synchronization object.
1014	Invalid credentials.
1015	Error returned from the NICI API.
1016	Cannot find peer cluster data.
1017	Invalid BCC API version.
1018	Could not find a pool for the specified cluster resource.
1019	Error managing the SAN via the Novell SAN Management Interface.
1020	CIM Client error.
1021	Error creating a system resource (mutex, semaphore, etc.).
1022	File IO error.
1023	No data.
1024	Not a member of the cluster.
1025	Invalid token in the script.
1026	Invalid or unknown cluster.
1027	The NSMI script is too long. It must be less than 64 KB.
1028	The cluster-enabled pool resource does not contain a volume.
1029	An operation has timed out.
1030	The specified resource is already busy.

The error code numbers are recorded in a log file that you can use to view status changes for BCC.

To view the log file on NetWare:

1. At the NetWare server console, enter `log +copy syslog`.
2. Open the file that is referenced in the message that appears.

   You can get additional information on how to use the log file by entering `help log` at the NetWare server console.

Security Considerations

10

This section contains specific instructions on how to configure and maintain a business continuity cluster in the most secure way possible. It contains the following subsections:

10.1 Security Features

The following table contains a summary of the security features of Novell® Business Continuity Clustering 1.2:

Table 10-1 *Business Continuity Clustering 1.2 Security Features*

Feature	Yes/No	Details
Users are authenticated	Yes	Administrative users are authenticated via eDirectory™.
Users are authorized	Yes	Users are authorized via eDirectory trustees.
Access to configuration information is controlled	Yes	Access to the administrative interface is restricted to valid users who have write rights to the configuration files.
Roles are used to control access	Yes	Configurable through iManager.
Logging and/or security auditing is done	Yes	Fake syslog.
Data on the wire is encrypted by default	Yes	The following data is encrypted on the wire: ◆ Inter-cluster communications ◆ Identity Manager data can be encrypted
Data stored is encrypted	No	
Passwords, keys, and any other authentication materials are stored encrypted	Yes	Inter-cluster communications for usernames and passwords are encrypted. Cluster credentials are stored encrypted in eDirectory.
Security is on by default	Yes	

10.2 Security Configuration

The following subsections provide a summary of security-related configuration settings for Business Continuity Clustering 1.1:

- Section 10.2.1, “BCC Configuration Settings,” on page 114
- Section 10.4, “Security Information for Dependent Products,” on page 119

10.2.1 BCC Configuration Settings

The following table lists the BCC configuration settings that are security-related or impact the security of BCC:

Table 10-2 *BCC Security Configuration Settings*

Configuration Setting	Possible Values	Default Value	Recommended Value for Best Security
Inter-cluster communications scheme	HTTP (port 5988) HTTPS (port 5989)	HTTPS	HTTPS
Identity Manager communications	Secure/Non-secure	This is the certificate in the Identity Manager driver setup. If you create the driver with an SSL certificate, it is secure. If not, it is in the clear.	Secure SSL certificates are mandatory for User Synchronization drivers to synchronize between trees.
BCC Administrator user quota in `sys:tmp`	0 MB to unlimited MB	Unlimited MB	5 MB Set this by using the NSS user quotas feature.

Configuration Setting	Possible Values	Default Value	Recommended Value for Best Security
BCC Administrator user	Any LUM-enabled eDirectory User	This is the user you specify when you are setting the BCC credentials. The BCC Administrator user is not automatically assigned the rights necessary to manage all aspects of each peer cluster. When managing individual clusters, you must log in as the Cluster Administrator user. You can manually assign the Cluster Administrator rights to the BCC Administrator user for each of the peer clusters if you want the BCC Administrator user to have all rights.	Unique BCC Administrator user (not the Admin user and not the Cluster Admin user)
BCC Administrator group	Any LUM-enabled eDirectory group	bccgroup	Unique group used for BCC administration. See Section 4.3, “Configuring a BCC Administrator User,” on page 45.
Peer cluster CIMOM URL (same as the Inter-cluster communication scheme)	http://*cluster_ip_address* *cluster_ip_address* where https:// is assumed.	*cluster_ip_address* https:// is assumed.	Default value

10.2.2 Changing the NCS: BCC Setings Attribute in the BCC XML Configuration

WARNING: You should not change the configuration settings for the NCS:BCC Settings attribute unless instructed to do so by Novell Support. Doing so can have adverse affects on your cluster nodes and BCC.

The following XML for the NCS:BCC Settings attribute is saved on the local Cluster object in eDirectory. The BCC must be restarted for changes to these settings to take effect. These are advanced settings that are intentionally not exposed in the BCC plug-in for iManager.

```
<bccSettings>
        <adminGroupName>bccgroup</adminGroupName>
        <authorizationCacheTTL>300</authorizationCacheTTL>
        <cimConnectTimeout>15</cimConnectTimeout>
        <cimReceiveTimeout>30</cimReceiveTimeout>
        <cimSendTimeout>30</cimSendTimeout>
        <idlePriorityThreshold>3</idlePriorityThreshold>
        <initialNormalThreads>3</initialNormalThreads>
        <initialPriorityThreads>2</initialPriorityThreads>
        <ipcResponseTimeout>15</ipcResponseTimeout>
        <maximumPriorityThreads>20</maximumPriorityThreads>
        <minimumPriorityThreads>2</minimumPriorityThreads>
        <resourceOfflineTimeout>300</resourceOfflineTimeout>
        <resourceOnlineTimeout>300</resourceOnlineTimeout>
        <scanForNewDevicesDelay>5</scanForNewDevicesDelay>
</bccSettings>
```

Table 10-3 provides additional information on each setting:

Table 10-3 *BCC XML Settings*

Setting	Description	Default Value
`<adminGroupName>`	The name of the LUM-enabled group that BCC uses on Linux.	bccgroup
`<authorizationCacheTTL>`	The number of seconds the authorization rights are cached in the BCC OpenWBEM provider.	300 seconds This is supported in BCC 1.1 and later.
`<cimConnectTimeout>`	BCC CIM client connect timeout in seconds.	15 seconds
`<cimReceiveTimeout>`	BCC CIM client receive timeout in seconds.	30 seconds
`<cimSendTimeout>`	BCC CIM client send timeout in seconds.	30 seconds
`<idlePriorityThreshold>`	The number of idle high priority threads before BCC starts killing priority threads.	3
`<initialNormalThreads>`	The number of normal threads created by BCC.	3
`<initialPriorityThreads>`	The number of high priority threads created by BCC at startup.	2
`<ipcResponseTimeout>`	The timeout in seconds that BCC waits for an IPC response.	15
`<maximumPriorityThreads>`	The maximum number of high priority threads BCC creates.	20
`<minimumPriorityThreads>`	The minimum number of high priority threads BCC keeps after killing idle high priority threads.	2
`<resourceOfflineTimeout>`	The number of seconds BCC waits for a resource to go offline during a BCC migrate.	300

Setting	Description	Default Value
<resourceOnlineTimeout>	The number of seconds BCC waits for a resource to go online during a BCC migrate.	300
<scanForNewDevicesDelay>	The number of seconds BCC sleeps after performing a scan for new devices during a BCC migration of a resource.	5

10.2.3 Disabling SSL for Inter-Cluster Communication

Disabling SSL for inter-cluster communication should only be done for debugging purposes, and should not be done in a production environment or for an extended period of time.

To turn off SSL for inter-cluster communication, or to specify a different communication port, you need to modify the Novell Cluster Services™ Cluster object that is stored in eDirectory by using an eDirectory management tool such as iManager or ConsoleOne®. See the *Novell iManager 2.7 Administration Guide* for information on using iManager.

Disabling SSL communication to a specific peer cluster requires changing the BCC management address to the peer cluster. The address is contained in the NCS:BCC Peers attribute that is stored on the NCS Cluster object.

For example, a default NCS:BCC Peers attribute could appear similar to the following example where https:// is assumed and is never specified explicitly:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES-TREE</tree>
  <address>10.1.1.10</address>
</peer>
```

To disable SSL for inter-cluster communication, you would change the <address> attribute to specify http:// with the IP address, as shown in the following example:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES-TREE</tree>
  <address>http://10.1.1.10</address>
</peer>
```

The BCC management address of `chicago_cluster` now specifies non-secure HTTP communication.

The BCC management port can also be changed by modifying the NCS:BCC Peers attribute values.

The default ports for secure and non-secure inter-cluster communication are 5989 and 5988 respectively.

For example, if you want to change the secure port on which OpenWBEM listens from port 5989 to port 1234, you would change the `<address>` attribute value in the above examples to:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES-TREE</tree>
  <address>10.1.1.10:1234</address>
</peer>
```

The attribute now specifies that inter-cluster communication uses HTTPS over port number 1234.

The NCS:BCC Peers attribute has a value for each peer cluster in the BCC. Attribute values are synchronized among peer clusters by the BCC-specific Identity Manager driver, so a change to an attribute value on one cluster causes that attribute value to be synchronized to each peer cluster in the BCC.

The changes do not take effect until either a reboot of each cluster node, or a restart of the Business Continuity Clustering software on each cluster node.

The following table provides an example of possible combinations of scheme and port specifier for the `<address>` tag for values of the NCS:BCC Peers attribute:

Table 10-4 *Example of Scheme and Port Specifier Values for the NCS:BCC Peers Attribute*

Value	Protocol Used	Port Used
10.1.1.10	HTTPS	5989
10.1.1.10:1234	HTTPS	1234
http://10.1.1.10	HTTP	5988
http://10.1.1.10:1234	HTTP	1234

10.3 General Security Guidelines

- Servers should be kept in a physically secure location with access by authorized personnel only.
- The corporate network should be physically secured against eavesdropping or packet sniffing. Any packets associated with the administration of BCC should be the most secured.
- Access to BCC configuration settings and logs should be restricted. This includes file system access rights, FTP access, access via Web utilities, SSH, and any other type of access to these files.
- Services that are used to send BCC data to other servers or e-mail accounts or that protect BCC data should be examined periodically to ensure that they have not been tampered with.
- When synchronizing cluster or user information between servers outside the corporate firewall, the HTTPS protocol should be employed. Because resource script information is passed between clusters, strong security precautions should be taken.
- When a BCC is administered by users outside of the corporate firewall, the HTTPS protocol should be used. A VPN should also be employed.
- If a server is accessible from outside the corporate network, a local server firewall should be employed to prevent direct access by a would-be intruder.
- Audit logs should be kept and analyzed periodically.

10.4 Security Information for Dependent Products

The following table provides links to security-related information for other products that impact the security of BCC:

Table 10-5 *Security Information for Other Products*

Product Name	Links to Security Information
eDirectory	Security for eDirectory is provided by NICI. See the *NICI 2.7x Administration Guide (http://www.novell.com/documentation/nici27x/nici_admin_guide/data/a20gkue.html)*.
Identity Manager (IDM)	"Security: Best Practices" (http://www.novell.com/documentation/idm35/admin/data/b1bsw73.html) in the *Identity Manager 3.5x Administration Guide*.
iSCSI	See "Configuring Access Control to iSCSI Targets" and "Enabling and Configuring iSCSI Initiator Security" in the *NW 6.5 SP8: iSCSI 1.1.3 Administration Guide*.
Novell Cluster Services for NetWare	In the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*, see "Configuration Requirements".
Novell Storage Services for NetWare	In the *NW 6.5 SP8: NSS File System Administration Guide*, see: ◆ "File Access for Users" ◆ "Securing Access to NSS Volumes, Directories, and Files" ◆ "Security Considerations"
OpenWBEM	OpenWBEM should be configured on each node to allow only the necessary users. OpenWBEM by default allows all users. For information, see "Changing the Authentication Configuration" in the *NW 6.5 SP8: OpenWBEM Services Administration Guide*.

Console Commands for BCC

A

Novell® Business Continuity Clustering (BCC) provides server console commands to help you perform certain BCC management tasks. Some of the commands can also be used to manage Novell Cluster Services™ clusters.

Table A-1 lists the server console commands for managing a business continuity cluster and gives a brief description of each command. For other cluster console commands, see "Console Commands for Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

To execute a cluster console command, enter `cluster` followed by the command. For example, if this cluster is a member of a business continuity cluster, and you want to see this cluster's peer clusters, enter `cluster view` at the server console. You can also enter `cluster help` at the console prompt to get information on the commands and their functions.

Table A-1 *Console Commands for Novell Business Continuity Clustering*

Console Command	Description
`cluster credentials [`*`cluster`*`]`	Lets you change the administrator username and password that this cluster uses to connect to the specified peer cluster. The cluster you specify must be a member of a cluster that has already been enabled for Business Continuity Clustering.
`cluster disable [`*`resource`*`]`	Disables Business Continuity Clustering for the specified resource. The resource you specify must be a member of a cluster that has already been enabled for Business Continuity Clustering. If no resource is specified, the entire cluster is disabled for Business Continuity Clustering. If you disable Business Continuity Clustering for a cluster by using the `cluster disable` console command, it is also disabled for those cluster resources that have been enabled for Business Continuity Clustering. If you re-enable Business Continuity Clustering for the cluster, you must re-enable each individual cluster resource that you want to be enabled for business continuity. This can be a time-consuming process if you have many cluster resources that are enabled for business continuity. For this reason, you should use caution when disabling Business Continuity Clustering for an entire cluster.

Console Command	Description
`cluster enable [resource]`	Enables Business Continuity Clustering for the specified resource. The resource you specify must be a member of a cluster that has already been enabled for Business Continuity Clustering. If no resource is specified, the entire cluster is enabled for Business Continuity Clustering. When enabling a resource for business continuity, previous versions of the CLI would not set the peer clusters where the resource was assigned to run. The only way to set peer clusters for a resource was through iManager. The new version of the cluster CLI automatically sets all the clusters in the BCC on a resource. Assigning a resource to specific clusters still must be done through iManager.
`cluster migrate [source/resource] [destination/nodename]`	Migrates the specified resource from the specified source cluster to the specified target (destination) cluster. Specifying * for the resource name migrates all BCC-enabled resources. Specifying * for the node name brings the resource online at the most preferred node.
`cluster resetresources`	Changes the state of all resources on this cluster to offline and secondary. This is a recovery procedure that should be run when a cluster in a BCC is brought back into service. You should run this command when only one node is a member of the cluster. 1. After a failure, bring up one node in the cluster. All other nodes should remain powered off. 2. Run the `cluster resetresources` command. 3. Bring up the remaining nodes in the cluster.
`cluster view`	Displays the node name, cluster epoch number, master node name, a list of nodes that are currently members of the cluster, and peer clusters if this cluster is a member of a business continuity cluster.
`cluster resources [resource]`	Lets you view the state and location of cluster resources and whether resources are primary or secondary. You can optionally specify a specific resource name.
`cluster status`	Lets you view the state and location of cluster resources and whether resources are primary or secondary. If the resource state is primary, the node where the resource is running is displayed. If the resource state is secondary, the cluster where the resource is located is displayed.
`cluster connections [-a]`	Displays the connection status of the cluster. Specifying -a attempts to show the connection status of all clusters in the BCC.

Console Command	Description
`cluster refresh`	This command should not be used except under the direction of Novell Support.

Implementing a Multiple-Tree BCC

B

Although the information contained in the other sections of this document describes an eDirectory™ single-tree implementation of Novell® Business Continuity Clustering (BCC) 1.1 Support Pack (SP) 2 for NetWare® 6.5 SP8, the information is also useful for configuring and managing BCC in a multiple-tree environment. This section contains additional instructions and information for multiple-tree BCC implementations.

B.1 Planning a Multiple-Tree Solution

Consider the following guidelines when creating a multiple-tree business continuity cluster.

B.1.1 Cluster Synchronization

Creating the Cluster Synchronization driver pair and its associated SSL Certificate is required in order for BCC to work between the two eDirectory trees. Typically, the cluster synchronization is between one cluster in each of the trees, then those clusters synchronize the information with other peer clusters in the same tree.

B.1.2 User Synchronization

If user-based access control is required for the cluster resources in the business continuity cluster, you must synchronize the user identities between the two eDirectory trees.To do this, create a single User Synchronization driver for the business continuity cluster.

If user-based access control is not needed for your cluster resources, the User Synchronization driver is optional.

Only one User Synchronization driver is required per eDirectory tree, even if you have multiple business continuity clusters set up. It is okay to have multiple User Syncrhonization drivers per tree.

The User object container should be in the same eDirectory partition as the Identity Manager node and cluster that you are using to create the User Synchronization driver. If the User object container is not in the same partition, you can create a partition for the User container, then add a read/write

replica of the partition on the Identity Manager node in the cluster that you are using to create the User Synchronization driver. If you create multiple User Synchronization drivers, each of the clusters involved must have a read/write replica of that User object container. An alternative approach when using a single User Synchronization driver is to make the eDirectory master server be a node in the cluster, install Identity Manager on that same node, then use that cluster when creating a User Synchronization driver. In this case, you do not need to create the User object container and to add server replicas.

The BCCAdmin user needs administrator rights in the container where the User objects reside so that User objects can be synchronized between eDirectory trees. For information, see Section 4.1.5, “Novell eDirectory 8.8,” on page 38.

B.1.3 SSL Certificates for Drivers

In a multiple-tree business continuity cluster, you should create separate SSL certificates for the Cluster Resource Synchronization driver and for the User Object Synchronization driver. We recommend that you create SSL certificates for your business continuity cluster to support secure data transfers between eDirectory trees. BCC works without the SSL certificates, but there is a security consideration.

You create one certificate for each of the driver pairs if the data flow is unidirectional. Two certificates are required if the data flow for the driver is bidirectional (one certificate for each direction). For example, create one SSL certificate for data flowing from TreeA to TreeB and a second SSL certificate for data flowing from TreeB to TreeA.

For security considerations, you should create or use a different certificate than the default (dummy) certificate (BCC Cluster Sync KMO) that is included with BCC.

B.2 Using Identity Manager to Copy User Objects to Another eDirectory Tree

The procedures explained in this section are normally performed after meeting “Requirements for BCC 1.1 SP2 for NetWare” on page 35.

The Identity Manager eDirectory driver has a synchronization feature that copies objects that exist in one tree to another tree where they don’t exist. For business continuity clusters, this feature can be used to copy User objects from one cluster to another cluster in a separate eDirectory tree. For example, if you have one tree that has 10,000 users and a second new tree that does not yet have users defined, you can use Identity Manager to quickly copy the 10,000 users to the new tree.

For more information on copying User objects by using Identity Manager, see “Migrating or Copying User Objects” (http://www.novell.com/documentation/idmdrivers/edirectory/data/brj81j4.html) in the *Identity Manager Driver for eDirectory Implementation Guide*.

B.3 Configuring User Object Synchronization

If the clusters in your business continuity cluster are in separate eDirectory trees and you require user-based access control, then User object synchronization is required.

To configure the Identity Manager driver for User object synchronization:

1. Start your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager Utilities*, then click the *New Driver* link.
4. Choose to either place the new driver in a new driver set, or add the driver to the driver set you created for cluster resource synchronization, then click *Next*.

   Both the User Object Synchronization driver and the Cluster Resource Synchronization driver can be added to the same driver set.
5. Specify the driver set name, context, and the server that the driver set will be associated with.

   The server is the same server where you installed the Identity Manager engine and eDirectory driver.
6. Choose to *not* create a new partition for the driver set, then click *Next*.
7. Choose to import a preconfigured driver from the server, select the Identity Manager preconfigured template for User object synchronization, then click *Next*.

   The template name is `BCCUserObjectSynchronization.XML`.
8. Fill in the values on the wizard page as prompted, then click *Next*.

   Each field contains an example of the type of information that should go into the field. Descriptions of the information required are also included with each field.

   Additional information for the wizard page fields can be found in "Importing the Sample Driver Configuration" (http://www.novell.com/documentation/dirxmldrivers/edirectory/data/bozojif.html) in the *Identity Manager Driver for eDirectory Implementation Guide*.
9. In the left column of the iManager page, click *Identity Manager*, then click *Identity Manager Overview*.
10. Search the eDirectory tree for the Identity Manager driver sets by clicking *Search*.
11. Click the *User Sync* driver icon, then click *Migrate from eDirectory*.
12. Click *Add*, browse to and select the context that contains the User objects, then click *OK*.
13. (Optional) Exclude the Admin User object from being synchronized:

    13a Click the *Exclude Administrative Roles* button, then click *Add*.

    13b Browse to and select the Admin User object, then click *OK*.
14. Perform Step 1 through Step 13 for each cluster that is in a separate tree.

B.4 Creating SSL Certificates

In a multiple-tree BCC, you must create an SSL certificate for the Cluster Resource Synchronization Driver, and an SSL certificate for the User Object Synchronization driver. Creating one certificate creates the certificate for a driver pair. For example, creating an SSL certificate for the Cluster Resource Synchronization driver creates the certificate for the Cluster Resource Synchronization drivers on both clusters.

To create an SSL certificate:

1. Start your Internet browser and enter the URL for iManager.

The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.

2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager Utilities*, then click *NDS-to-NDS Driver Certificates*.
4. Specify the requested driver information for both eDirectory trees.

   You must specify the driver name (including the context) you supplied in Step 8 on page 66 for the current tree. Use the following format when specifying the driver name:

   DriverName.DriverSet.OrganizationalUnit.OrganizationName

   Ensure that there are no spaces (beginning or end) in the specified context, and do not use the `cn=DriverName.ou=OrganizationalUnitName.o=OrganizationName` format.

B.5 Synchronizing the BCC-Specific Identity Manager Drivers

After creating the BCC-specific Identity Manager drivers and SSL certificates, if you are adding a new cluster to an existing business continuity cluster in a multiple-tree BCC, you must synchronize the BCC-specific Identity Manager drivers for the Cluster and User objects. If the BCC-specific Identity Manager drivers are not synchronized, clusters cannot be enabled for business continuity. This is not necessary unless you are adding a new cluster to an existing business continuity cluster.

To synchronize the BCC-specific Identity Manager drivers:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Identity Manager*, then click the *Identity Manager Overview* link.
4. Search for and select the BCC driver set.
5. Click the red *Cluster Sync* icon for the driver you want to sync, then click the *Migrate from eDirectory* button.
6. Click *Add*, browse to and select the Cluster object for the new cluster you are adding, then click *OK*.

   Selecting the Cluster object causes the BCC-specific Identity Manager drivers to synchronize.
7. If you chose to include User object synchronization, repeat the above steps, and in Step 5, click the *User Sync* icon.

B.6 Preventing Identity Manager Synchronization Loops

If you have three or more clusters each in separate eDirectory trees in your business continuity cluster, you should set up IDM User object and Cluster Resource object synchronization in a manner that prevents Identity Manager synchronization loops. Identity Manager synchronization loops can cause excessive network traffic and slow server communication and performance.

For example, in a three-cluster business continuity cluster, an Identity Manager synchronization loop occurs when Cluster One is configured to synchronize with Cluster Two, Cluster Two is configured to synchronize with Cluster Three, and Cluster Three is configured to synchronize back to Cluster One. This is illustrated in Figure B-1.

Figure B-1 *Three-Cluster Identity Manager Synchronization Loop*

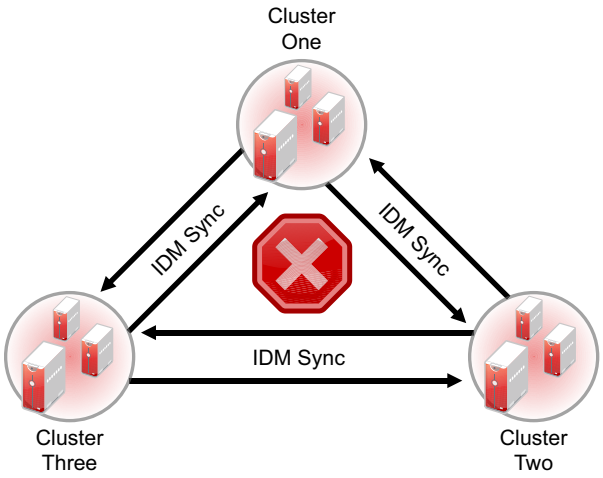

A preferred method is to make Cluster One an Identity Manager synchronization master in which Cluster One synchronizes with Cluster Two, and Cluster Two and Cluster Three both synchronize with Cluster One. This is illustrated in Figure B-2.

Figure B-2 *Three-Cluster Identity Manager Synchronization Master*

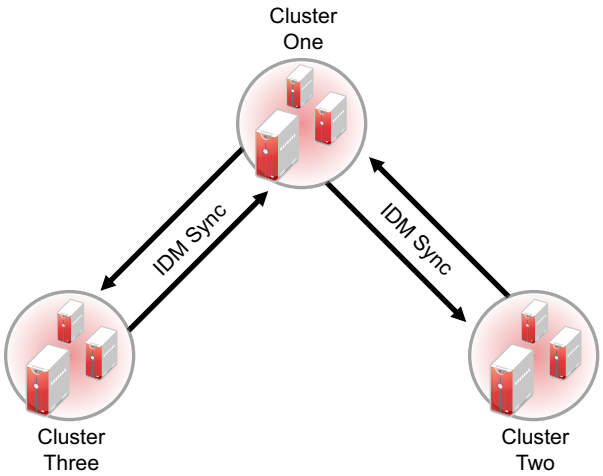

You could also have Cluster One synchronize with Cluster Two, Cluster Two synchronize with Cluster Three, and Cluster Three synchronize back to Cluster Two as illustrated in Figure B-3.

Figure B-3 *Alternate Three-Cluster Identity Manager Synchronization Scenario*

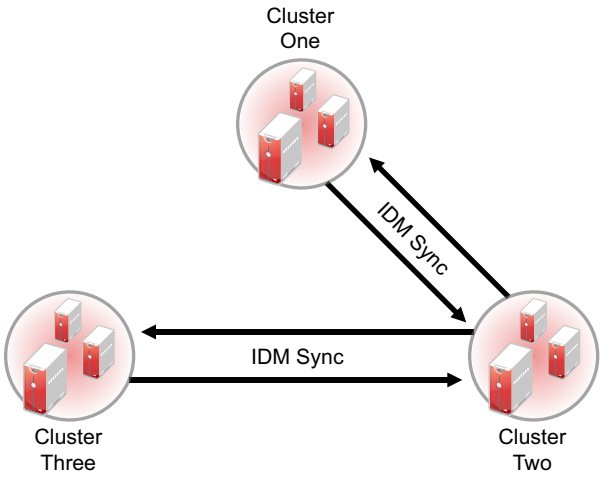

To change your BCC synchronization scenario:

1. In the Connections section of the Business Continuity Cluster Properties page, select one or more peer clusters that you want a cluster to synchronize to, then click *Edit*.

In order for a cluster to appear in the list of possible peer clusters, that cluster must have the following:

- Business Continuity Clustering software installed.
- Identity Manager installed.
- BCC Identity Manager drivers configured and running.
- Be enabled for business continuity.

B.7 Migrating Resources to Another Cluster

IMPORTANT: If you are migrating a pool to a cluster in another tree and you want to maintain that pool's volume trustee assignments, you must migrate the pool to a server with an eDirectory replica. The replica must be at least read-only and must contain all users. After migrating the pool to a server with an eDirectory replica, enter the following console command on that server for each volume in the pool:

```
NSS/ResetObjectIDStore=volumename
```

This command updates all volume trustee assignments and should be run at night, on a weekend, or during a period of low network utilization. Trustee assignments become effective immediately, but might take a few hours to display correctly in management utilities.

If you migrate the pool to a server in another tree without an eDirectory replica, you must, within 90 days, migrate that pool to a server with an eDirectory replica and then run the command for each volume.

Setting Up Auto-Failover

C

Auto-failover is available beginning in Novell® Business Continuity Clustering 1.1. To set up the auto-failover feature, you must enable it, then configure the auto-failover settings.

WARNING: Auto-failover is disabled by default and is not recommended. It should only be enabled after a thorough examination and review of your network and geographic site infrastructure. You should seriously consider the adverse conditions that might occur as a result of enabling this feature.

These conditions may include but are not limited to:

- Data loss at one or more geographic sites
- Data corruption at one or more geographic sites
- Data divergence at one or more geographic sites

For example, if there is a loss of communication between two clusters and auto-failover has been enabled and configured, each cluster will assert ownership of BCC-enabled cluster resources. These resources then automatically load on both clusters.

When communication between cluster has been restored, some of the data on each cluster is different. This is called data divergence. Also, the mirroring or synchronization process either fails, or attempts to overwrite any changed data on one cluster. This causes either data loss or data corruption.

C.1 Enabling Auto-Failover

To enable auto-failover for all Business Continuity Cluster resources in a cluster:

1. Start your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/iManager.html. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
2. Specify your username and password, specify the tree where you want to log in, then click *Login*.
3. In the left column, click *Clusters*, then click the *Cluster Options* link.
4. Specify a cluster name, or browse and select the Cluster object.
5. Click the *Properties* button and then click the *Business Continuity* tab.
6. Click the *Main* link.
7. Click the *Auto-Failover* button.

8 Click the *Auto-Failover* link just under the tabs.

9 Select the *Enable Automatic Failover of Business Continuity Cluster Resources* check box, then click *Apply*.

10 Continue with Section C.2, “Creating an Auto-Failover Policy,” on page 132 to create a failover policy.

Auto-failover is not completely enabled until you create an auto-failover policy.

C.2 Creating an Auto-Failover Policy

By default, no auto-failover policy exists for BCC. You must create an auto-failover policy for each cluster in your BCC where you want auto-failover enabled. This is required to automatically fail over resources from one cluster to another.

1 In iManager, under *Cluster Membership Monitoring Settings*, select a cluster and click the *Edit* link.

2 Under *Membership Threshold*, select the *Enable* check box, select either *Percent Fail* or *Nodes Fail*, and specify either the percentage of failed nodes or the number of failed nodes.

The node failure number or percentage you specify must be met for the selected cluster before resources automatically fail over to another cluster.

IMPORTANT: Do not use a membership condition of total node failure (either 100 percent or the total number of nodes); the condition cannot be satisfied because the cluster will not be up to report this state.

If a cluster has been totally downed, you must bring up the master node in the downed cluster, then run the `cluster resetresources` command on that node before you begin manually migrating the BCC-enabled resources to peer clusters.

---

3 Under *Communication Timeout*, select the *Enable* check box and specify the number of minutes that must elapse without any communication between clusters before resources automatically fail over to another cluster.

4 Click *OK* to finish editing the policy.

5 Click the *Apply* button to save your settings.

C.3 Refining the Auto-Failover Policy

You can further refine auto-failover policies to give you more control over if or when an auto-failover occurs. To do this, click the *Advanced* button to display additional fields for specifying auto-failover criteria and adding monitoring information.

The policy for automatic failover is configured by creating rules. Each row in the Failover Policy Configuration table represents a rule that applies to a single cluster, or to all clusters in the business continuity cluster. Each rule contains a set of conditions. Each condition tests one of the following criteria:

- The value of an indication reported by a monitor
- The amount of time the connection to a cluster has been down
- If the connection to a cluster is up

These conditions can be combined in any order to construct a more robust rule that helps to avoid an undesired failover. For failover to occur, each condition of only one rule must be satisfied for the specified cluster or clusters.

For rules with monitor conditions automatically created by using the Cluster Membership Monitoring Settings table, you can add a condition that tests if the connection to the peer cluster is up. Adding this condition changes the behavior of the rule. With this rule, a graceful automatic failover of resources can happen when the connection to the peer cluster is up.

You can also specify or change the criteria for percent or number of nodes that are used to determine if an automatic failover can occur.

IMPORTANT: Do not use a membership condition of total node failure (either 100 percent or the total number of nodes); the condition cannot be satisfied because the cluster will not be up to report this state.

If a cluster has been totally downed, you must bring up the master node in the downed cluster, then run the `cluster resetresources` command on that node before you begin manually migrating the BCC-enabled resources to peer clusters.

You should create a separate rule with a connection down condition. Adding a connection down condition to an existing rule with a condition that tests cluster membership is not recommended. It is highly unlikely that cluster membership information for a specific cluster will be reported to peer clusters when the connection to that specific cluster is down.

For example, a rule might contain only one condition that tests whether a connection to a specific cluster has been down for five or more minutes. Failover occurs when peer clusters agree that the connection to the cluster specified in the rule has been down for five or more minutes. If the peer clusters do not agree about the connection being down (that is, one cluster has a valid connection to the specified cluster), failover does not occur. More complex rules can be constructed that contain multiple conditions.

If previously configured, the fields under *Failover Policy Configuration* should already contain information on the policies that were created in the *Cluster Membership Monitoring Settings* section of the page.

1. Under *Failover Policy Configuration*, select a policy and click *Edit* to further refine a rule. Click *Delete* to remove the rule, or click *New* to create a new rule that you can add the additional failover conditions to.
2. Select the cluster that you want the rule to apply to, or select *All* to apply the policy to all clusters.
3. Under *Conditions*, choose the type of condition and the appropriate values. To add multiple conditions to the rule, click the *Add* button below the condition.

   You can use the default setting of *Monitor* if you don't want to apply the cluster up or cluster down criteria to this policy. You can also specify or change the percent or number of nodes criteria that are used to determine if an auto failover can occur.
4. Click *Apply* to save your settings.

C.4 Adding or Editing Monitor Configurations

Clicking the *Advanced* button also displays an additional section on this page called *Health Monitor Configuration*. Monitors are an important part of the automatic failover feature, and are separate processes that perform a specialized task to analyze the health of a specific cluster or all clusters in the BCC. These monitors report an indication of health to BCC. BCC in turn uses the reported information to analyze the failover policy to determine if resources should be migrated from a specific cluster. Business Continuity Clustering 1.1 ships with two monitors (nodecnt and node pnt) that report an indication of health that represents either the percentage or number of nodes that are not a member of a specific cluster.

If they are configured by using the Cluster Membership Monitoring Settings table, the fields under *Health Monitor Configuration* should already contain information for the health monitor (nodepnt or nodecnt) included with Business Continuity Clustering 1.1. Although default values have already been set, you can customize some of the monitor settings for the cluster membership monitors. If you have created your own custom monitor, you can click *New* to add configuration settings to your monitor.

1. In iManager, under *Monitor Name* in the *Health Monitor Configuration* section, select a monitor and click *Edit*.
2. Under *Clusters*, select the cluster or clusters that you want this monitor to apply to.
3. Specify the maximum health indication that the monitor will report.

   This value is used when creating a failover policy to validate the rules. This is the maximum value that can be used for the threshold type when you create a failover policy. For example, if you specified percent fail membership monitoring, the maximum values would be 100 (for 100 percent) for the nodepnt monitor. If you specified nodes fail membership monitoring, the maximum value for the nodecnt monitor is the maximum number of nodes permitted in a cluster, which is 32. If you created your own custom monitor, the values could be different.

   For the nodepnt and nodecnt monitors, the *Maximum Health Indication* value is for information only, and should not be changed.
4. Under *Short Polling Interval*, specify the number of seconds the monitor will wait each time it contacts the cluster or clusters to get health information.

   The *Long Polling Interval* is not used with the default nodepnt and nodecnt monitors. This value might be used for some custom monitors.
5. Specify NetWare® as the platform that you want to be monitored by the health monitor and whether you want the monitor enabled for the selected clusters.

   The *Optional Parameter* field specifies a monitor-specific string value that is passed to the monitor as a startup parameter.

   The nodepnt and nodecnt monitors do not support optional parameters.
6. Click *Apply* to save your settings.

NOTE: See the BCC NDK documentation (http://developer.novell.com/documentation/cluster/ncss_enu/data/bktitle.html) for more information on creating custom failover policies.

Configuring Host-Based File System Mirroring for NSS Pools

D

Several methods and scenarios exist for mirroring data between geographically separate sites. Each method has its own strengths and weaknesses. For a Novell® Business Continuity Clustering system, you need to choose either host-based mirroring or SAN-based mirroring (also called array-based mirroring) and whether you want the mirroring to be synchronous or asynchronous.

Figure D-1 *Synchronous Mirroring*

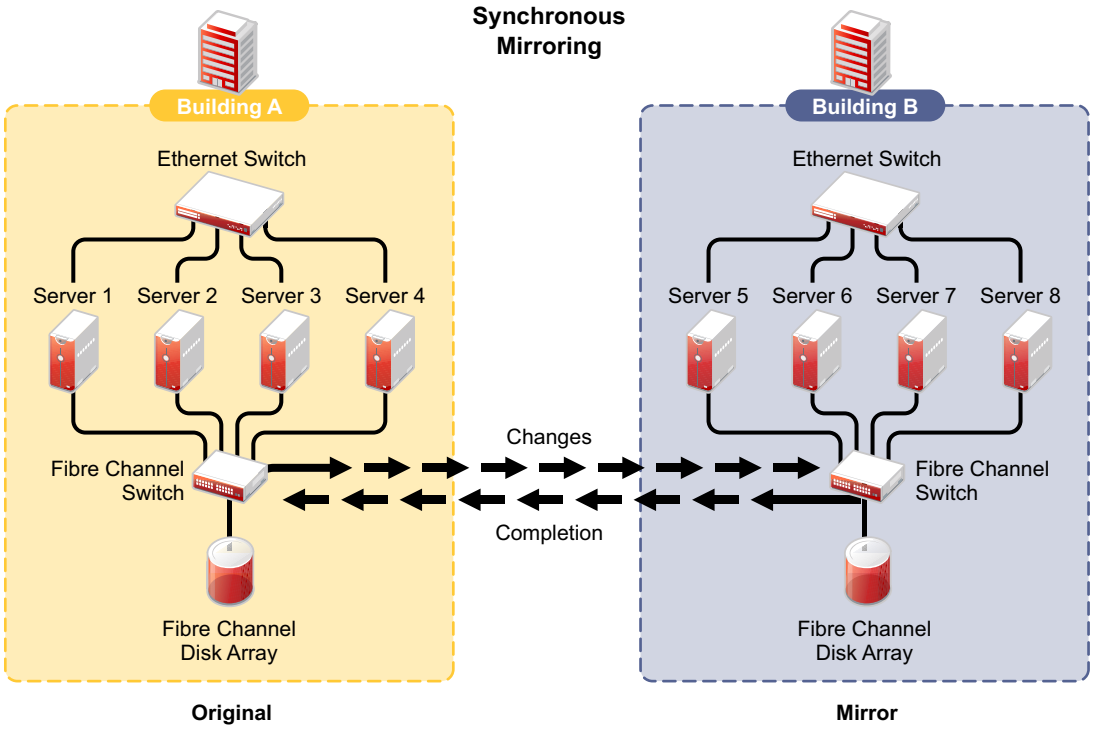

IMPORTANT: The Business Continuity Clustering product does not perform data mirroring. You must separately configure either SAN-based mirroring or host-based mirroring.

Storage-based synchronous mirroring is preferred and is provided by storage hardware manufacturers. For information about storage-based mirroring, consult your storage system vendor or see the storage system vendor documentation.

Host-based synchronous mirroring functionality is included with the Novell Storage Services™ (NSS) file system (NSS mirroring) that is part of NetWare 6.5. NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. It is an alternative to SAN-based synchronous replication options.

IMPORTANT: NSS pool snapshot technology does not work in a business continuity cluster.

- Section D.1, “Creating and Mirroring NSS Partitions on Shared Storage,” on page 136

D.1 Creating and Mirroring NSS Partitions on Shared Storage

NSS provides a software RAID 1 configuration option that mirrors NSS pool partitions. The partitions are automatically created by the NSS management tools when you create a pool and when you mirror the pool. NSS supports two to four segments in a software RAID 1. To ensure disaster recovery, the device you select to mirror should be in another storage array in the other data center.

For example, you create the original pool in one cluster, then create mirrors for that pool in the other peer clusters, for a total two to four segments.

NSS partitions must be mirrored after they are created. If you have an existing partition that you want to mirror, you can either create another partition of equal size on another device to mirror the first partition to, or let the mirroring software automatically create another partition of equal size on another device.

Prior to creating and mirroring NSS partitions on shared storage, ensure that you have the following:

- All servers in the cluster connected to a shared storage system
- One or more drive arrays configured on the shared storage system
- The drives on the shared storage system marked as shared.
- NSS is installed and running. For information, see “Installing and Configuring Novell Storage Services” in the *NW 6.5 SP8: NSS File System Administration Guide*.
- You can add CIFS or AFP as advertising protocols for the NSS pool resource when you create the resource, or add them later by using the Clusters plug-in for iManager.
- You need a static IP address for the pool resource. It must be in the same subnet as the cluster master IP address.

To create and mirror NSS pools:

1. Start NSSMU by entering `NSSMU` at the server console of a cluster server.
2. Select *Partitions* from the NSSMU main menu.
3. Press the Insert key and select the device on your shared storage system where you want to create a partition.

   With a device marked as sharable for clustering, all partitions on that device are automatically sharable.

   Device names are not changeable and might be labeled something like 0x2 or 0x1.
4. Select *NSS* as the partition type, then specify the partition size and, if desired, an NSS pool name and label.

   If you specify a pool name, a pool by that name is automatically created on the partition. If no pool name is specified, you need to create a pool on the partition later.

   4a If you chose to create a pool, choose whether you want the pool to be activated and cluster-enabled when it is created.

The *Activate on Creation* option is enabled by default. This causes the pool to be activated as soon as it is created. If you choose not to activate the pool, you need to manually activate it later before it can be used.

The *Cluster Enable on Creation* option is also enabled by default. If you want to cluster-enable the pool at the same time it is created, accept the default entry (*Yes*) and continue with Step 4b. If you want to cluster-enable the pool at a later date, see the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide* for more information.

4b Specify the *Virtual Server Name*, *IP Address*, *Advertising Protocols* and, if necessary, the *CIFS Server Name*.

When you cluster-enable a pool, a virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is hosting the pool.

You can select one or all of the advertising protocols. NCP™ is the protocol used by Novell clients, CIFS is the protocol used by Microsoft clients, and AFP is the protocol used by Macintosh clients. Selecting any of the protocols causes lines to be added to the pool resource load and unload scripts to activate the selected protocols on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to all your clients.

If you select CIFS as one of the protocols, a *CIFS Server Name* is also required. This is the server name CIFS clients see when they browse the network. A default server name is listed, but you can change the server name by editing the text in the field.

4c Select *Create* to create and cluster-enable the pool.

5 Select the partition you want to mirror (this might be the partition you just created) and press the F3 key.

6 Select the device with free space or the partition you want to mirror to, then select *YES* to mirror the partition.

To ensure disaster recovery, the device you select to mirror should be in another storage array in the other data center.

D.2 Creating NSS Volumes

After you create the pool and its mirror, you must create an NSS volume on each of the pool resources so that it works properly as a cluster resource. To create an NSS volume on a shared pool, follow the instructions in "Creating an NSS Volume on a Shared Pool" in the *What's New for Novell Cluster Services for NetWare*.

You must create at least one shared volume in a cluster-enabled pool. Typically, all volumes are created when you initially set up the cluster resource and before you need to cluster migrate or fail over the resource to other servers in the cluster.

You can add volumes to the pool later by cluster migrating the pool cluster resource back to the original server node in the cluster where the pool was created. Otherwise, you get an eDirectory™ error because the tools only look for the Pool object under its current server node, and not under the original node where it it was created.

To create or modify home directories, Distributed File Services junctions, or any other elements that are managed through eDirectory objects, you must cluster migrate the pool resource back to the node where it was created before you perform those management tasks. This restriction also applies to management tasks like renaming a pool or volume that change information in the eDirectory objects for the shared pool or volume.

D.3 Novell Cluster Services Configuration and Setup

After configuring NSS mirroring and creating a pool and volume on the mirrored NSS partition, if you did not cluster-enable the NSS pool on the mirrored partition when you created it, do so by following the instructions in "Configuring a Cluster Resource for (Cluster Enabling) an NSS Pool and Its Volumes" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.

When you cluster-enable a shared disk pool, the commands to start and stop the pool resource are automatically added to the resource load and unload scripts.

D.4 Checking NSS Volume Mirror Status

After you have configured NSS mirroring with Novell Cluster Services™, you should check to ensure that it is working properly in a cluster environment.

1. Ensure that the volumes on the cluster-enabled pool are mounted on the assigned server by entering `volumes` at the server console.
2. Check the mirror status of the mirrored partition by entering `mirror status` at the server console of the server where the NSS pool on the mirrored partition is active.

   After entering `mirror status`, you should see a message indicating that mirror status is 100 percent or a message indicating that the mirrored object is fully synchronized.
3. Migrate the pool to another server in the cluster, and check again to ensure that the volumes on the pool are mounted by entering `volumes` at the server console.
4. Check the mirror status of the partition again by entering `mirror status` at the server console.

IMPORTANT: If you create or delete a pool or partition on shared storage that is part of a business continuity cluster, you must run the `cluster scan for new devices` command on a server in each of the other clusters that belong to the business continuity cluster.

Documentation Updates

E

This section contains information about documentation content changes made to the *Novell® Business Continuity Clustering 1.1 SP2 Administration Guide for NetWare® 6.5 SP8* since the initial 1.1 SP2 release. If you are an existing user, review the change entries to readily identify modified content. If you are a new user, simply read the guide in its current state.

Refer to the publication date, which appears on the title page, to determine the release date of this guide. For the most recent version of the *Novell Business Continuity Clustering 1.1 for NetWare Administration Guide*, see the Business Continuity Clustering Documentation Web site (http://www.novell.com/documentation/bcc/).

In this section, content changes appear in reverse chronological order, according to the publication date. Within a dated entry, changes are grouped by section and sequenced alphabetically. Each change entry provides a link to the related topic and a brief description of the change.

This document was updated on the following dates:

E.1 August 14, 2009

Updates were made to the following sections. The changes are explained below.

E.1.1 Console Commands for BCC

Location	Change
Appendix A, “Console Commands for BCC,” on page 121	Updated to add a link to Novell Cluster Services console commands.

E.1.2 Installing Business Continuity Clustering

Location	Change
Section 4.1.17, “Web Browser,” on page 44	This section is new.

E.1.3 Upgrading Business Continuity Clustering for NetWare

Location	Change
Section 5.2, “Disabling BCC 1.0, Upgrading Servers to NetWare 6.5 SP8, then Enabling BCC 1.1 SP2,” on page 54	This section is new.

E.2 April 28, 2009

Updates were made to the following sections. The changes are explained below.

- Section E.2.1, “Converting BCC Clusters from NetWare to Linux,” on page 140

E.2.1 Converting BCC Clusters from NetWare to Linux

Location	Change
Chapter 5, “Upgrading Business Continuity Clustering for NetWare,” on page 53	Updated for clarity.
Chapter 5, “Upgrading Business Continuity Clustering for NetWare,” on page 53	Updated for clarity.