

Business Continuity Clustering 2.6

Administration Guide for Open Enterprise Server 2018 SP1

January 2019

Legal Notices

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.microfocus.com/about/legal/>.

Copyright © 2019 Micro Focus Software, Inc. All Rights Reserved.

Contents

About This Guide	9
1 Overview of Business Continuity Clustering	11
1.1 Disaster Recovery Implications	11
1.2 Disaster Recovery Implementations	12
1.2.1 LAN-Based versus Internet-Based Applications	12
1.2.2 Host-Based versus Storage-Based Data Mirroring	13
1.2.3 Stretch Clusters versus Cluster of Clusters	13
1.3 Business Continuity Clustering	19
1.4 BCC Deployment Scenarios	20
1.4.1 Two-Site Business Continuity Cluster Solution	20
1.4.2 Multiple-Site Business Continuity Cluster Solution	22
1.4.3 Low-Cost Business Continuity Cluster Solution	23
1.5 Key Concepts	23
1.5.1 Business Continuity Clusters	23
1.5.2 Cluster Resources	23
1.5.3 Landing Zone	23
1.5.4 BCC Drivers for Identity Manager	23
2 What's New or Changed for BCC	25
2.1 BCC Engine	25
2.2 What's Next	26
3 Planning a Business Continuity Cluster	27
3.1 Determining Design Criteria	27
3.2 Best Practices	27
3.3 LAN Connectivity Guidelines	28
3.3.1 VLAN	28
3.3.2 Channel Bonding	28
3.3.3 IP Addresses	29
3.3.4 Name Resolution and SLP	30
3.4 SAN Connectivity Guidelines	30
3.5 Storage Design Guidelines	30
3.6 eDirectory Design Guidelines	31
3.6.1 Object Location	31
3.6.2 Cluster Context	31
3.6.3 Partitioning and Replication	31
3.6.4 Objects Created by the BCC Drivers for Identity Manager	32
3.6.5 Landing Zone	32
3.6.6 Naming Conventions for BCC-Enabled Resources	32
3.7 Cluster Design Guidelines	33
4 Installation Requirements for BCC	35
4.1 Getting a Business Continuity Clustering License	35
4.2 Downloading Business Continuity Clustering Software	35
4.3 Minimum System Requirements	36

4.4	Open Enterprise Server 2018 SP1	37
4.5	SLP	38
4.6	Small Footprint CIM Broker and CIMOM	38
4.7	Novell Cluster Services for Linux	39
4.7.1	Cluster Names	39
4.7.2	Cluster Resource Names	39
4.7.3	Storage	40
4.7.4	Volume IDs	40
4.7.5	Cluster Containers	40
4.7.6	Peer Cluster Credentials	40
4.8	eDirectory 9.1	41
4.8.1	eDirectory Containers for Clusters	41
4.8.2	eDirectory Read/Write Replica on Each IDM Node	41
4.8.3	Rights Needed for Installing BCC	42
4.8.4	Rights Needed for BCC Management	42
4.8.5	Rights Needed by the BCC Administrator to Manage Peer Clusters	42
4.8.6	Rights Needed by BCC Drivers	43
4.8.7	eDirectory Requirements for Identity Manager	43
4.9	Identity Manager 4.7	43
4.9.1	Downloading the Identity Manager and Patches	44
4.9.2	Credential for Drivers	44
4.9.3	IDM Components	44
4.9.4	BCC Cluster Resource Synchronization Template for the eDirectory Driver	44
4.10	iManager 3.1	44
4.11	Storage-Related Plug-Ins for iManager 3.1	45
4.12	Shared Disk Systems	46
4.13	Mirroring Shared Disk Systems between Peer Clusters	46
4.14	LUN Masking for Shared Devices	46
4.15	Link Speeds	47
4.16	Ports	47
4.17	Web Browser	47
4.17.1	Web Browser Language Setting	47
4.17.2	Web Browser Character Encoding Setting	48
4.18	What's Next	48
5	Installing Business Continuity Clustering	49
5.1	Overview of the BCC Installation and Configuration	49
5.2	Installing iManager and Identity Manager on One Node in Each Peer Cluster	50
5.2.1	Installing iManager and the Plug-Ins Used by BCC	50
5.2.2	Installing and Configuring the Identity Manager Components Used by BCC	51
5.2.3	Installing the Identity Manager Plug-in for iManager	54
5.2.4	Adding Peer Clusters to Your My Clusters List	54
5.3	Configuring a BCC Administrator User and Group	55
5.3.1	Accessing iManager	55
5.3.2	Creating the BCC Group and Administrator User	56
5.3.3	Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects	57
5.3.4	Adding the BCC Administrator User to the ncsgroup on Each Cluster Node	58
5.4	Installing and Configuring the Business Continuity Clustering Software	59
5.4.1	Installing the Business Continuity Clustering	59
5.4.2	Configuring BCC Software	61
5.5	Installing the BCC Cluster Resource Template	62
5.6	Configuring the Device Rescan for BCC Migration of Resources	62
5.7	Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software	63
5.7.1	Creating a YaST Auto-Configuration Profile	63

5.7.2	Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media	64
5.7.3	Installing and Configuring Business Continuity Clustering on Each Cluster Node	65
5.7.4	Removing the NFS Share from Your Server	65
5.7.5	Cleaning Up the Business Continuity Clustering Installation Source	66
5.8	What's Next	66

6 Configuring the Identity Manager Drivers for BCC 67

6.1	Understanding the BCC Drivers	67
6.1.1	IDM Driver Set and Driver Guidelines	68
6.1.2	Preventing Synchronization Loops for Identity Manager Drivers	68
6.1.3	Example 1: Two Peer Clusters	71
6.1.4	Example 2: Three Peer Clusters	71
6.1.5	Example 3: Four Peer Clusters	72
6.2	Prerequisites for Configuring the BCC Drivers for Identity Manager	72
6.2.1	Identity Manager	73
6.2.2	eDirectory	73
6.2.3	Landing Zone Container	73
6.2.4	BCC Admin User and Group	73
6.3	Creating the BCC Driver Set	73
6.4	Creating a BCC Driver for a BCC Driver Set	75
6.5	Enabling or Disabling the Synchronization of Email Settings	79
6.6	Synchronizing Identity Manager Drivers	79
6.7	Changing the Identity Manager Synchronization Drivers	79
6.8	What's Next	79

7 Configuring BCC for Peer Clusters 81

7.1	Enabling Clusters for Business Continuity	81
7.2	Adding Peer Cluster Credentials	83
7.3	Adding Search-and-Replace Values to the Resource Replacement Script	87
7.4	Enabling Linux POSIX File Systems to Run on Secondary Clusters	88
7.5	Verifying BCC Administrator User Trustee Rights and Credentials	89
7.6	Disabling BCC for a Peer Cluster	89
7.7	What's Next	90

8 Configuring BCC for Cluster Resources 91

8.1	Requirements for Cluster Resources	91
8.1.1	LUNs for Cluster Pool Resources	91
8.1.2	Volumes for Cluster Pool Resources	91
8.1.3	Shared Disk Cluster Resources	92
8.1.4	Cluster Resources for OES 2018 SP1 Services	92
8.2	BCC-Enabling Cluster Resources	93
8.3	Configuring Search-and-Replace Values for an Individual Cluster Resource	94
8.4	Adding BCC Load and Unload Scripts	95
8.5	Assigning Preferred Peer Clusters for the Resource	97
8.6	Assigning Preferred Nodes in Peer Clusters	98
8.7	Disabling BCC for a Cluster Resource	99
8.8	Deleting or Unsharing a BCC-Enabled Shared NSS Pool Resource	100
8.9	Permanently Relocating a Cluster Resource to a Peer Cluster	102
8.10	Renaming a BCC-Enabled Pool or Volume	103

9	Managing a Business Continuity Cluster	105
9.1	Migrating a Cluster Resource to a Peer Cluster	105
9.1.1	Understanding BCC Resource Migration	105
9.1.2	Migrating Cluster Resources between Clusters	106
9.2	Bringing a Downed Cluster Back in Service	107
9.3	Changing Peer Cluster Credentials	107
9.4	Viewing the Current Status of a Business Continuity Cluster	109
9.4.1	Using iManager to View the Cluster Status	109
9.4.2	Using Console Commands to View the Cluster Status	109
9.5	Generating a Cluster Report	109
9.6	Resolving Business Continuity Cluster Failures	110
9.6.1	Storage-Based Mirroring Failure Types and Responses	110
9.6.2	Host-based Mirroring Failure Types and Responses	112
10	Troubleshooting Business Continuity Clustering	115
10.1	Identity Manager	115
10.1.1	Identity Manager Plug-Ins Do Not Appear in iManager	115
10.1.2	Identity Manager Drivers for Cluster Synchronization Do Not Start	115
10.1.3	Identity Manager Drivers Do Not Synchronize Objects from One Cluster to Another	116
10.1.4	Tracing Identity Manager Communications	116
10.1.5	SSL Certificates Are Missing	117
10.2	Peer Clusters	117
10.2.1	Incorrect Status in the Cluster View Report After Enabling BCC on Peer Clusters	118
10.2.2	Administration of Peer Clusters Is Not Functional	118
10.2.3	Peer Cluster Communication Is Not Working	118
10.2.4	Cluster Connection States	118
10.2.5	Driver Port Number Conflicts	119
10.2.6	Security Equivalent User	120
10.2.7	Clusters Cannot Communicate	121
10.3	BCC-Enabled Cluster Resources	121
10.3.1	Resource Cannot Be Brought Online	121
10.3.2	Resource Does Not Migrate to a Peer Cluster	122
10.3.3	Blank Error String iManager Error Appears While Bringing a Resource Online	122
10.3.4	Clustered Pool Is Stuck in an eDirectory Synchronization State	122
10.3.5	Mapping Drives in Login Scripts	122
10.3.6	Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable	123
10.3.7	Resource Script Search-and-Replace Functions Do Not Work	124
10.3.8	Virtual NCP Server IP Addresses Won't Change	124
10.3.9	The IP Address, Virtual Server DN, or Pool Name Does Not Appear on the iManager Cluster Configuration Page	125
10.4	BCC Startup Flags	125
10.5	BCC Error Codes	126
11	Security Considerations	129
11.1	Security Features	129
11.2	Security Configuration	130
11.2.1	BCC Configuration Settings	130
11.2.2	Changing the NCS: BCC Settings Attributes in the BCC XML Configuration	131
11.2.3	Disabling SSL for Inter-Cluster Communication	132
11.2.4	Restricting the Network Address for Administration	134
11.3	General Security Guidelines	134
11.4	Security Information for Dependent Products	134

A	Console Commands for BCC	137
A.1	Using Console Commands	137
A.2	Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh	140
B	Configuration Worksheet for the BCC Drivers for Identity Manager	141
B.1	Cluster1 to Cluster2	141
B.2	Cluster2 to Cluster1	142
C	Setting Up Auto-Failover	145
C.1	Enabling Auto-Failover	145
C.2	Creating an Auto-Failover Policy	146
C.3	Refining the Auto-Failover Policy	146
C.4	Adding or Editing Monitor Configurations	147
D	Configuring Host-Based File System Mirroring for NSS Pools	149
D.1	Creating and Mirroring NSS Pools on Shared Storage	150
D.2	Creating NSS Volumes	152
D.3	Novell Cluster Services Configuration and Setup	153
D.4	Checking NSS Volume Mirror Status	153
E	Using Dynamic DNS with BCC	155
E.1	Requirements and Assumptions	155
E.1.1	Software Requirements	155
E.1.2	DNS Server	155
E.1.3	TSIG Keys	156
E.1.4	DNS Record Time-to-Live Values	156
E.2	Configuring the DNS Server for Dynamic DNS	156
E.2.1	Creating the TSIG Keys for DNS Server Authentication	156
E.2.2	Configuring the DNS Server with the Public Key	157
E.2.3	Configuring the DNS Server Zones	158
E.2.4	Testing the DNS Server	158
E.3	Configuring the Cluster Resources for Dynamic DNS	162
E.3.1	Modifying the BCC Load Script	163
E.3.2	Public and Private Keys	165
E.3.3	Testing the Perl Wrapper Script	166
E.4	Testing the Dynamic DNS Solution	166
F	Using Virtual IP Addresses with BCC	169
F.1	Understanding Internal Virtual IP Networks	169
F.1.1	Virtual Adapter	169
F.1.2	Host Mask	169
F.1.3	Internal Router	170
F.2	Virtual IP Address Benefits	170
F.2.1	High Availability	170
F.2.2	Unlimited Mobility	171
F.2.3	Automatic Name Resolution	171
F.3	Planning a Virtual IP Network Implementation	172
F.3.1	Routing Protocol	172
F.3.2	LAN Routers	172
F.3.3	Internal Router	172
F.3.4	IP Addresses for BCC-Enabled Cluster Resources	172

F.3.5	Host Mask	173
F.4	Configuring a Virtual Router with OSPF	173
F.5	Configuring Virtual IP Addresses	174
G	Upgrading to Identity Manager 4.7	177
G.1	Requirements for Using Identity Manager	177
G.2	Upgrading to Identity Manager 4.7	177
G.3	Patching IDM 4.7 for OES 2018 SP1	178
H	Upgrading to BCC 2.6 on OES 2018 SP1	179
H.1	System Requirements for the BCC	179
H.2	Upgrading to BCC 2.6	179
I	Documentation Updates	181
I.1	January, 2019	181
I.2	September 21, 2013	181

About This Guide

This guide describes how to install, configure, and manage Business Continuity Clustering 2.6 for Open Enterprise Server (OES) 2018 Support Pack (SP1) servers in combination with Novell Cluster Services.

- ♦ [Chapter 1, “Overview of Business Continuity Clustering,”](#) on page 11
- ♦ [Chapter 2, “What’s New or Changed for BCC,”](#) on page 25
- ♦ [Chapter 3, “Planning a Business Continuity Cluster,”](#) on page 27
- ♦ [Chapter 4, “Installation Requirements for BCC,”](#) on page 35
- ♦ [Chapter 5, “Installing Business Continuity Clustering,”](#) on page 49
- ♦ [Chapter 6, “Configuring the Identity Manager Drivers for BCC,”](#) on page 67
- ♦ [Chapter 7, “Configuring BCC for Peer Clusters,”](#) on page 81
- ♦ [Chapter 8, “Configuring BCC for Cluster Resources,”](#) on page 91
- ♦ [Chapter 9, “Managing a Business Continuity Cluster,”](#) on page 105
- ♦ [Chapter 10, “Troubleshooting Business Continuity Clustering,”](#) on page 115
- ♦ [Chapter 11, “Security Considerations,”](#) on page 129
- ♦ [Appendix A, “Console Commands for BCC,”](#) on page 137
- ♦ [Appendix B, “Configuration Worksheet for the BCC Drivers for Identity Manager,”](#) on page 141
- ♦ [Appendix C, “Setting Up Auto-Failover,”](#) on page 145
- ♦ [Appendix D, “Configuring Host-Based File System Mirroring for NSS Pools,”](#) on page 149
- ♦ [Appendix E, “Using Dynamic DNS with BCC,”](#) on page 155
- ♦ [Appendix F, “Using Virtual IP Addresses with BCC,”](#) on page 169
- ♦ [Appendix G, “Upgrading to Identity Manager 4.7,”](#) on page 177
- ♦ [Appendix H, “Upgrading to BCC 2.6 on OES 2018 SP1,”](#) on page 179
- ♦ [Appendix I, “Documentation Updates,”](#) on page 181

Audience

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

The [Security Considerations](#) section provides information of interest for security administrators or anyone who is responsible for the security of the system.

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 and enter your comments there.

Documentation Updates

The latest version of this *Business Continuity Clustering 2.6 Administration Guide for OES 2018 SP1* is available under “BCC for OES 2018 SP1” on the [Business Continuity Clustering Documentation Web site \(http://www.novell.com/documentation/bcc/\)](http://www.novell.com/documentation/bcc/).

Additional Documentation

For information about Open Enterprise Server 2018 Support Pack 1, see the [OES 2018 SP1 \(http://www.novell.com/documentation/open-enterprise-server-2018/\)](http://www.novell.com/documentation/open-enterprise-server-2018/) documentation Web site on Novell.com.

For information about Novell Cluster Services, see the [OES 2018 SP1: Novell Cluster Services for Linux Administration Guide](#).

For information about eDirectory 9.1, see the [eDirectory 9.1 \(https://www.netiq.com/documentation/edirectory-91/\)](https://www.netiq.com/documentation/edirectory-91/) documentation Web site on NetIQ.com.

For information about iManager 3.1, see the [iManager 3.1 \(https://www.netiq.com/documentation/imanager-31/\)](https://www.netiq.com/documentation/imanager-31/) documentation Web site on NetIQ.com.

For information about Identity Manager 4.7, see the [Identity Manager 4.7 \(https://www.netiq.com/documentation/identity-manager-47/\)](https://www.netiq.com/documentation/identity-manager-47/) documentation Web site on NetIQ.com.

1 Overview of Business Continuity Clustering

As corporations become more international, fueled in part by the reach of the Internet, the requirement for service availability has increased. 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 challenge 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. Make one mistake and the redundant site is no longer able to effectively take over in the event of a disaster.

Business Continuity Clustering can improve your disaster recovery solution by providing specialized software that manages site-to-site failover of critical workgroup and networking services. BCC works with Open Enterprise Server and Novell Cluster Services to automate cluster configuration, maintenance, and synchronization across two to four geographically separate sites. Services can easily fail over to another cluster in a completely different geographic location. This helps to eliminate downtime, ensure critical services are available, mitigate compliance risk, and minimize the possibility of human error.

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 the benefits of using BCC for disaster recovery of your critical workgroup and networking services.

- ◆ [Section 1.1, “Disaster Recovery Implications,” on page 11](#)
- ◆ [Section 1.2, “Disaster Recovery Implementations,” on page 12](#)
- ◆ [Section 1.3, “Business Continuity Clustering,” on page 19](#)
- ◆ [Section 1.4, “BCC Deployment Scenarios,” on page 20](#)
- ◆ [Section 1.5, “Key Concepts,” on page 23](#)

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, considerations include the following:

- ◆ 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 13](#)
- ♦ [Section 1.2.3, “Stretch Clusters versus 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 products like NetIQ Access Manager and Novell 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 versus 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.

Business Continuity Clustering automates some of the configuration and processes used in a cluster of clusters. For information, see [Section 1.3, “Business Continuity Clustering,” on page 19](#).

- ◆ [“Stretch Clusters” on page 14](#)
- ◆ [“Cluster of Clusters” on page 15](#)

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

Stretch Clusters

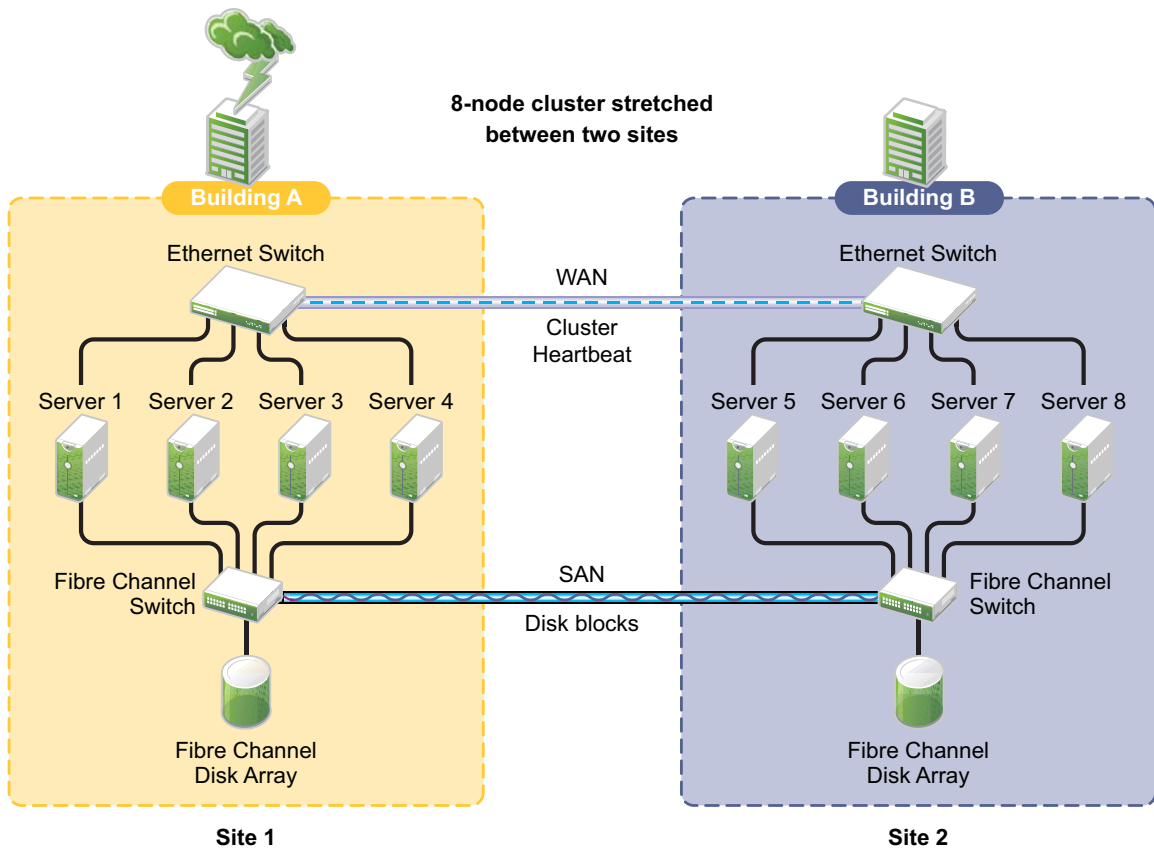
A stretch cluster consists of a single cluster where the nodes are located in two geographically separate data centers. All nodes in the cluster must be in the same eDirectory tree, which requires the eDirectory replica ring to span data centers. The IP addresses for nodes and cluster resources in the 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 13](#). 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

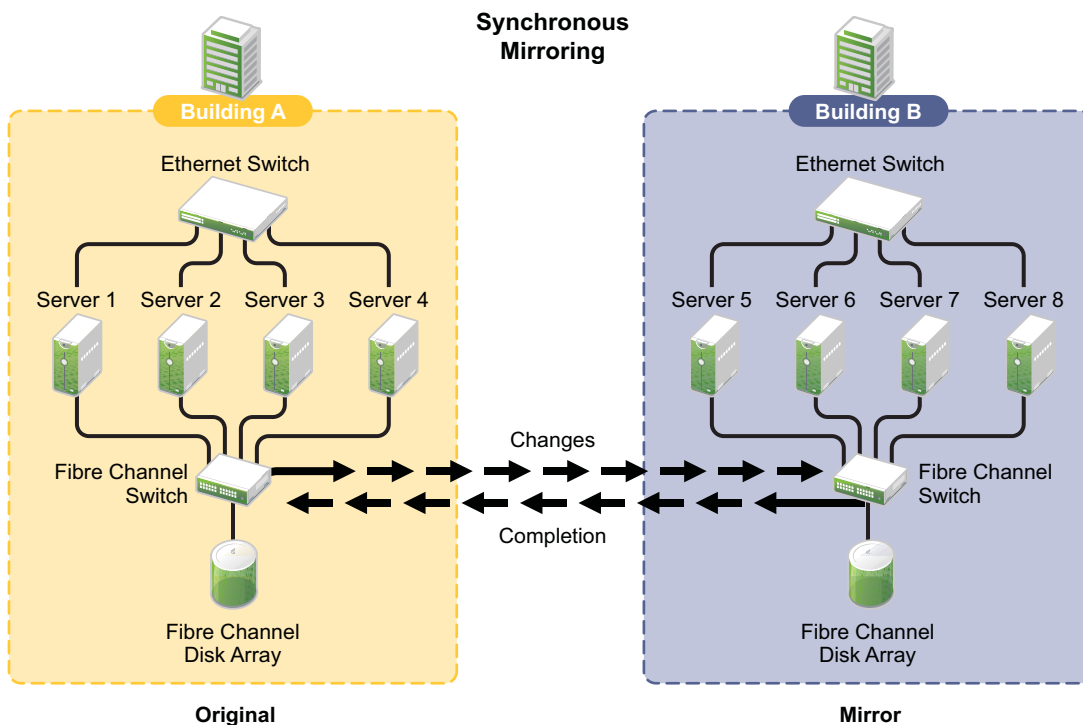
A cluster of clusters consists of multiple clusters in which each cluster is located in a geographically separate data center. Each cluster can be in different Organizational Unit (OU) containers in the same eDirectory tree. Each cluster can be in a different IP subnet.

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 13.

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 with SAN-Based Data Mirroring



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 to four
Number of geographically separated data centers	Two	Two to four
eDirectory trees	Single tree; requires the replica ring to span data centers.	Single tree
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.
Data-block replication between sites	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.
For information about data replication solutions, see	Section 1.2.2, "Host-Based versus Storage-Based Data Mirroring," on page 13.	

Capability	Stretch Cluster	Cluster of Clusters
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 examines why a cluster of cluster solution is less problematic to deploy than a stretch cluster solution. It identifies the advantages, disadvantages, and other considerations for each. Manual configuration is not a problem when using Business Continuity Clustering for your cluster of clusters.

Table 1-3 Evaluation of Stretch Clusters versus Cluster of Clusters

	Stretch Cluster	Cluster of Clusters
Advantages	<ul style="list-style-type: none"> ◆ 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. 	<ul style="list-style-type: none"> ◆ eDirectory partitions don't need to span the cluster. ◆ Each cluster can be in different OUs in the same eDirectory tree. ◆ 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	<ul style="list-style-type: none"> ◆ 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. 	<ul style="list-style-type: none"> ◆ Resource configurations must be manually synchronized. ◆ Storage-based mirroring requires additional configuration steps.

	Stretch Cluster	Cluster of Clusters
Other Considerations	<ul style="list-style-type: none"> ◆ 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. 	<ul style="list-style-type: none"> ◆ Depending on the platform used, storage arrays must be controllable by scripts that run on OES if the SANs are not SMI-S compliant.

1.3 Business Continuity Clustering

A 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 15](#), 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, the BCC solution uses a cluster of clusters. Each geographically separate site hosts an independent cluster that is treated as a “peer cluster” in a larger geographically dispersed cluster of clusters. This allows a site to do fan-out failover of resources to multiple other sites. BCC automates the failover between peer clusters by using eDirectory and policy-based management of the resources and storage systems.

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 shared storage hardware devices to automate the failover process through standards-based mechanisms such as SMI-S.

- ◆ Uses Identity Manager technology to automatically synchronize and transfer cluster-related eDirectory objects from one cluster to another.
- ◆ 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 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 22](#)
- ◆ [Section 1.4.3, “Low-Cost Business Continuity Cluster Solution,” on page 23](#)

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

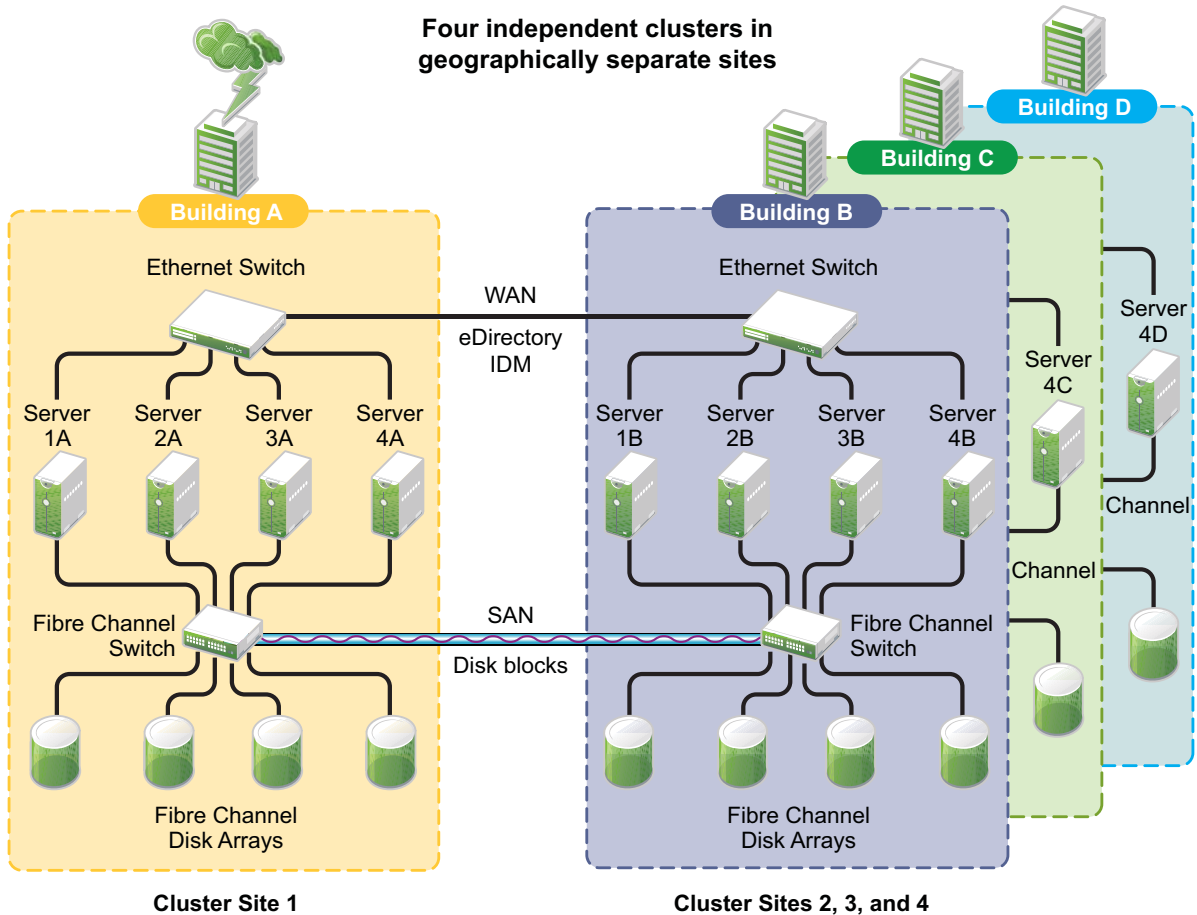


1.4.2 Multiple-Site Business Continuity Cluster Solution

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 Four-Site Business Continuity Cluster



Using additional software, 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 with iSCSI-based block replication or with 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.

- ♦ [Section 1.5.1, “Business Continuity Clusters,” on page 23](#)
- ♦ [Section 1.5.2, “Cluster Resources,” on page 23](#)
- ♦ [Section 1.5.3, “Landing Zone,” on page 23](#)
- ♦ [Section 1.5.4, “BCC Drivers for Identity Manager,” on page 23](#)

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. For information, see [Chapter 6, “Configuring the Identity Manager Drivers for BCC,” on page 67](#).

2 What's New or Changed for BCC

This section describes the changes and enhancements that were made to Business Continuity Clustering (BCC) for Open Enterprise Server (OES) 2018 Support Pack 1 (SP1).

- ♦ [Section 2.1, “BCC Engine,”](#) on page 25
- ♦ [Section 2.2, “What’s Next,”](#) on page 26

2.1 BCC Engine

Beginning with OES 2018 SP1, Novell Business Continuity Cluster (BCC) is available as a pattern in the OES installation pattern list. This eases the distribution and delivery of BCC for OES customers. However, BCC continues to be a separate product with its own license.

Business Continuity Clustering supports up to four peer clusters that are running OES 2018 SP1 and Novell Cluster Services.

The BCC operating environment requires the versions of OES and NetIQ products listed in [Table 2-1](#), with the latest patches applied.

Table 2-1 BCC Operating Environment

Supported Product	For documentation, see:
eDirectory 9.1 This product is included with OES 2018 SP1.	eDirectory 9.1 (https://www.netiq.com/documentation/edirectory-91/) documentation Web site on NetIQ.com
Open Enterprise Server 2018 Support Pack 1	OES 2018 SP1 (http://www.novell.com/documentation/open-enterprise-server-2018) documentation Web site on Novell.com
Novell Cluster Services	<i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i> For information about clustering OES 2018 SP1 services and storage, see the product documentation on the OES 2018 SP1 (https://www.novell.com/documentation/open-enterprise-server-2018/) documentation Web site on Novell.com.
iManager 3.1 This product is included with OES 2018 SP1.	iManager 3.1 (https://www.netiq.com/documentation/imanager-31/) documentation Web site on NetIQ.com
Clusters plug-in for iManager This product is included with OES 2018 SP1.	“Installing or Updating the Clusters Plug-in for iManager” in the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i>
Identity Manager 4.7 The credential to activate IDM is included in the BCC license.	Identity Manager 4.7 (https://www.netiq.com/documentation/identity-manager-47/) documentation Web site on NetIQ.com

2.2 What's Next

For new installations of BCC 2.6 on OES 2018 SP1 clusters, see the following:

- ♦ [Chapter 3, “Planning a Business Continuity Cluster,” on page 27](#)
- ♦ [Chapter 4, “Installation Requirements for BCC,” on page 35](#)
- ♦ [Chapter 5, “Installing Business Continuity Clustering,” on page 49](#)
- ♦ [Chapter 6, “Configuring the Identity Manager Drivers for BCC,” on page 67](#)
- ♦ [Chapter 7, “Configuring BCC for Peer Clusters,” on page 81](#)

For upgrades from BCC 2.x to BCC 2.6 on OES 2018 SP1 clusters, see the following:

- ♦ [Chapter 4, “Installation Requirements for BCC,” on page 35](#)
- ♦ [Appendix G, “Upgrading to Identity Manager 4.7,” on page 177](#)
- ♦ [Appendix H, “Upgrading to BCC 2.6 on OES 2018 SP1,” on page 179](#)

3 Planning a Business Continuity Cluster

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. You can use the guidelines in this section to design your Business Continuity Clustering solution.

- ◆ Section 3.1, “Determining Design Criteria,” on page 27
- ◆ Section 3.2, “Best Practices,” on page 27
- ◆ Section 3.3, “LAN Connectivity Guidelines,” on page 28
- ◆ Section 3.4, “SAN Connectivity Guidelines,” on page 30
- ◆ Section 3.5, “Storage Design Guidelines,” on page 30
- ◆ Section 3.6, “eDirectory Design Guidelines,” on page 31
- ◆ Section 3.7, “Cluster Design Guidelines,” on page 33

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:

- ◆ Ensure that eDirectory and your clusters are stable before implementing BCC.
- ◆ Engage Micro Focus 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
- ♦ 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, “Channel Bonding,” on page 28](#)
- ♦ [Section 3.3.3, “IP Addresses,” on page 29](#)
- ♦ [Section 3.3.4, “Name Resolution and SLP,” on page 30](#)

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 failures.

3.3.2 Channel Bonding

Use channel bonding for adapters for LAN fault tolerance. Channel bonding 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. For information, see `/usr/src/linux/Documentation/bonding.txt`.

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 ensure 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 the guidelines in this section to plan your IP address assignment so that it is consistently applied across all peer clusters.

- ♦ [“Using Unique IP Addresses” on page 29](#)
- ♦ [“Using Dedicated IP Address Ranges” on page 29](#)
- ♦ [“Changing a Resource IP Address” on page 29](#)

Using Unique IP Addresses

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, such as file system resources and service resources that support site-based services like DHCP, DNS, and SLP.
- ♦ Cluster resources that are BCC-enabled, such as file system resources and application resources.

Using Dedicated IP Address Ranges

You can dedicate IP address ranges for BCC-enabled cluster resources. Your IP address plan should provide an IP address range with sufficient addresses for each cluster. With careful planning, the IP address and the name of the virtual server for the cluster resource never need to change.

Changing a Resource IP Address

When a cluster resource is BCC-migrated to a peer cluster, the IP address of the inbound cluster resource is transformed to use an IP address in the same subnet of the target peer cluster. 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. For information about setting up the transformation rules, see [Section 7.3, “Adding Search-and-Replace Values to the Resource Replacement Script,” on page 87](#).

Dynamic DNS (Domain Name Service) is an alternate technique for transforming IP addresses. For information, see [Appendix E, “Using Dynamic DNS with BCC,” on page 155](#).

Unrelated to BCC, if you change the IP address of a BCC-enabled cluster resource, modify the address according to the instructions in [“Changing the IP Address of a Cluster Resource”](#) in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*. Before you bring the resource online, modify the transformation rules as needed to ensure that the resource can still be failed over to a peer cluster.

IMPORTANT: If you change the IP address of a BCC-enabled cluster resource, you might also need to update the transformation rules for the resource.

3.3.4 Name Resolution and SLP

The master IP addresses are stored in the NCS:BCC Peers attribute. Ensure that SLP is properly configured for name resolution.

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 Business Continuity Clustering, ensure that 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.
- ◆ Ensure that 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 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. Using multiple pools per LUN is possible, but is 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 only in a given cluster at any given time. For maximum flexibility, we recommend that you create only one cluster resource per LUN.

- ◆ We recommend that you use only one volume per pool.
- ◆ 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, ensure 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 ensure that all segments of the resource fail over together.

3.6 eDirectory Design Guidelines

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

- ◆ [Section 3.6.1, “Object Location,” on page 31](#)
- ◆ [Section 3.6.2, “Cluster Context,” on page 31](#)
- ◆ [Section 3.6.3, “Partitioning and Replication,” on page 31](#)
- ◆ [Section 3.6.4, “Objects Created by the BCC Drivers for Identity Manager,” on page 32](#)
- ◆ [Section 3.6.5, “Landing Zone,” on page 32](#)
- ◆ [Section 3.6.6, “Naming Conventions for BCC-Enabled Resources,” on page 32](#)


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 you cluster-enable NSS pools, these are the default naming conventions used by NSS:


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. Ensure that your approach is consistent across all peer clusters.

- ◆ **IP addresses:**
 - ◆ Ensure that IP addresses are unique across all BCC peer clusters. For information, see [“Using Unique IP Addresses” on page 29](#).
 - ◆ IP address assignments should be consistently applied within each peer cluster and for all cluster resources. For information, see [“Using Dedicated IP Address Ranges” on page 29](#).
- ◆ **Volume IDs:** Volume IDs of BCC-enabled clustered volumes must be unique across all nodes in every peer cluster. Duplicate volume IDs can prevent resources from going online if the resource is BCC-migrated to a peer cluster.
- ◆ **Configuring Nodes:** As you build nodes for each peer cluster, consider the configuration requirements for each of the services supported across all peer clusters, and for the preferred nodes for each service.
- ◆ **Planning BCC-Enabled Resource Failover:** As you plan your BCC solution, create a failover matrix 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.

4 Installation Requirements for BCC

This section defines the installation requirements for Business Continuity Clustering. Ensure that your clusters meet these requirements before you install BCC in any of the peer clusters.

- ♦ Section 4.1, “Getting a Business Continuity Clustering License,” on page 35
- ♦ Section 4.2, “Downloading Business Continuity Clustering Software,” on page 35
- ♦ Section 4.3, “Minimum System Requirements,” on page 36
- ♦ Section 4.4, “Open Enterprise Server 2018 SP1,” on page 37
- ♦ Section 4.5, “SLP,” on page 38
- ♦ Section 4.6, “Small Footprint CIM Broker and CIMOM,” on page 38
- ♦ Section 4.7, “Novell Cluster Services for Linux,” on page 39
- ♦ Section 4.8, “eDirectory 9.1,” on page 41
- ♦ Section 4.9, “Identity Manager 4.7,” on page 43
- ♦ Section 4.10, “iManager 3.1,” on page 44
- ♦ Section 4.11, “Storage-Related Plug-Ins for iManager 3.1,” on page 45
- ♦ Section 4.12, “Shared Disk Systems,” on page 46
- ♦ Section 4.13, “Mirroring Shared Disk Systems between Peer Clusters,” on page 46
- ♦ Section 4.14, “LUN Masking for Shared Devices,” on page 46
- ♦ Section 4.15, “Link Speeds,” on page 47
- ♦ Section 4.16, “Ports,” on page 47
- ♦ Section 4.17, “Web Browser,” on page 47
- ♦ Section 4.18, “What’s Next,” on page 48

4.1 Getting a Business Continuity Clustering License

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

4.2 Downloading Business Continuity Clustering Software

Beginning with OES 2018 SP1, Novell Business Continuity Cluster (BCC) is available as a pattern in the OES installation pattern list. However, BCC continues to be a separate product with its own license.

To install from the OES pattern list, follow the instructions in [Section 5.4, “Installing and Configuring the Business Continuity Clustering Software,”](#) on page 59.

To install from the network using a YaST auto-configuration file, see [Section 5.7, “Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software,”](#) on page 63.

4.3 Minimum System Requirements

Business Continuity Clustering supports up to four peer clusters that are running OES 2018 SP1 and Novell Cluster Services.

The BCC operating environment requires the versions of OES and NetIQ products listed in [Table 4-1](#), with the latest patches applied.

Table 4-1 BCC Operating Environment

Supported Product	For documentation, see:
eDirectory 9.1 This product is included with OES 2018 SP1.	eDirectory 9.1 (https://www.netiq.com/documentation/edirectory-91/) documentation Web site on NetIQ.com
Open Enterprise Server 2018 Support Pack 1	OES 2018 SP1 (http://www.novell.com/documentation/open-enterprise-server-2018) documentation Web site on Novell.com
Novell Cluster Services	OES 2018 SP1: Novell Cluster Services for Linux Administration Guide For information about clustering OES 2018 SP1 services and storage, see the product documentation on the OES 2018 SP1 (https://www.novell.com/documentation/open-enterprise-server-2018/) documentation Web site on Novell.com.
iManager 3.1 This product is included with OES 2018 SP1.	iManager 3.1 (https://www.netiq.com/documentation/imanager-31/) documentation Web site on NetIQ.com
Clusters plug-in for iManager This product is included with OES 2018 SP1.	“Installing or Updating the Clusters Plug-in for iManager” in the OES 2018 SP1: Novell Cluster Services for Linux Administration Guide
Identity Manager 4.7 The credential to activate IDM is included in the BCC license.	Identity Manager 4.7 (https://www.netiq.com/documentation/identity-manager-47/) documentation Web site on NetIQ.com

Fully configured and tested peer clusters are the most important requirement for BCC. [Figure 4-1](#) shows a business continuity cluster with two peer clusters. It illustrates where the required products are installed. The figure uses the following abbreviations:

BCC: Business Continuity Clustering for OES 2018 SP1

eDir: [eDirectory 9.1](#)


OES Linux: [Open Enterprise Server 2018 SP1](#)

NCS: [Novell Cluster Services for Linux](#)

iManager: [iManager 3.1](#)

IDM: [Identity Manager 4.7](#)

Figure 4-1 Business Continuity Cluster Component Locations



4.4 Open Enterprise Server 2018 SP1

Beginning with OES 2018 SP1, Novell Business Continuity Cluster (BCC) is available as a pattern in the OES installation pattern list. Apply the latest OES 2018 SP1 maintenance patches on all nodes in every peer cluster to automatically upgrade the BCC.

To download OES 2018 SP1, go to [Open Enterprise Server 2018 SP1 \(https://download.novell.com/index.jsp\)](https://download.novell.com/index.jsp) and download the ISO files. Maintenance patches are available through the OES 2018 SP1 patch channel.

See the [OES 2018 SP1: Installation Guide](#) for information on installing, configuring, and updating (patching) OES 2018 SP1. For OES 2018 SP1 system requirements, see “Preparing to Install OES 2018 SP1” in the [OES 2018 SP1: Installation Guide](#).

If you are upgrading an existing business continuity cluster 2.x to BCC 2.6, you must upgrade the Identity Manager servers in the BCC before you upgrade the peer clusters. See the following resources for upgrade requirements:

Table 4-2 *Upgrading Identity Manager and Business Continuity Clustering*

From	To	See the following:
Identity Manager 4.5.6 on OES versions (OES 11 SP2, OES 11 SP3, OES 2015, or OES 2015 SP1)	Identity Manager 4.7 on OES 2018 SP1	Appendix G, “Upgrading to Identity Manager 4.7,” on page 177
Business Continuity Clustering 2.x on OES versions (OES 11 SP2, OES 11 SP3, OES 2015, or OES 2015 SP1)	Business Continuity Clustering 2.6 on OES 2018 SP1	Appendix H, “Upgrading to BCC 2.6 on OES 2018 SP1,” on page 179.

4.5 SLP

You must have SLP (Server Location Protocol) set up and configured properly on each server node in every peer cluster. Typically, SLP is installed as part of the eDirectory installation and set up when you install the server operating system for the server. For information, see [“Configuring OpenSLP for eDirectory”](#) in the *NetIQ eDirectory Administration Guide*.

4.6 Small Footprint CIM Broker and CIMOM

The Small Footprint CIM Broker (SFCB) replaces OpenWBEM for CIMOM activities in BCC. OES 2018 SP1 offer SFCB as the default CIMOM and CIM clients. When you install any OES components that depend on WBEM, SFCB and all of its corresponding packages are installed with the components. For information, see [“Small Footprint CIM Broker \(SFCB\)”](#) in the *OES 2018 SP1: Planning and Implementation Guide*.

IMPORTANT: SFCB must be running and working properly whenever you modify the settings for BCC, the cluster, and the cluster resources.

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.

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 user name and password) between iManager and the `_admin` volume on the OES server you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use HTTPS for transferring data, which ensures that sensitive data is not exposed.

IMPORTANT: SFCB is automatically PAM-enabled for Linux User Management (LUM) as part of the OES installation. Users not enabled for LUM cannot use the CIM providers to manage OES. The user name that you use to log in to iManager when you manage a cluster and the BCC cluster must be a eDirectory user name that has been LUM-enabled.

For more information about the permissions and rights needed by the BCC Administrator user, see [Section 4.8, “eDirectory 9.1,” on page 41](#) and [Section 5.3, “Configuring a BCC Administrator User and Group,” on page 55.](#)

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 SFCB is not running.

You can use the `rscfcb` command to help resolve CIMOM and SFCB issues:

To perform this task	At a command prompt, enter as the root user
To start SFCB service	<code>rscblim-sfcb start</code>
To stop SFCB service	<code>rscblim-sfcb stop</code>
To check SFCB service status	<code>rscblim-sfcb status</code>
To restart SFCB service	<code>rscblim-sfcb restart</code>

For more information, see “Web Based Enterprise Management using SFCB” (https://www.suse.com/documentation/sles-12/book_sle_admin/data/cha_wbem.html) in the *SUSE Linux Enterprise Server 12 Administration Guide* (https://www.suse.com/documentation/sles-12/book_sle_admin/data/book_sle_admin.html).

4.7 Novell Cluster Services for Linux

Business Continuity Clustering supports two to four peer clusters running Novell Cluster Services (NCS) with the latest patches applied. Novell Cluster Services must be installed and running on each node in every peer cluster.

IMPORTANT: Before you install and configure BCC, ensure that you have set up and fully tested the peer clusters.

See the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide* for information on installing, configuring, and managing Novell Cluster Services. NCS maintenance patches are available through the OES patch channel.

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

- ♦ [Section 4.7.1, “Cluster Names,” on page 39](#)
- ♦ [Section 4.7.2, “Cluster Resource Names,” on page 39](#)
- ♦ [Section 4.7.3, “Storage,” on page 40](#)
- ♦ [Section 4.7.4, “Volume IDs,” on page 40](#)
- ♦ [Section 4.7.5, “Cluster Containers,” on page 40](#)
- ♦ [Section 4.7.6, “Peer Cluster Credentials,” on page 40](#)

4.7.1 Cluster Names

Each cluster must have a unique name across all peer clusters.

4.7.2 Cluster Resource Names

Each cluster resource must have a unique name across all peer clusters.

4.7.3 Storage

The storage requirements for Business Continuity Clustering software are the same as for Novell Cluster Services. For information, see the following in the [OES 2018 SP1: Novell Cluster Services for Linux Administration Guide](#):

- ◆ [“Hardware Requirements”](#)
- ◆ [“Shared Disk Configuration Requirements”](#)
- ◆ [“SAN Rules for LUN Masking”](#)
- ◆ [“Multipath I/O Configuration Requirements”](#)

BCC requires some scripted management of your SAN devices to manage the failover between peer clusters. Some storage vendors require you to purchase or license their CLI (Command Line Interface) separately. The CLI for the storage subsystem might not initially be included with your hardware.

Some storage hardware might not be SMI-S compliant and cannot be managed by using SMI-S commands. If the storage subsystems are not SMI-S compliant, the storage subsystems must be controllable by scripts running on the nodes of the cluster. For information, see [Section 8.4, “Adding BCC Load and Unload Scripts,”](#) on page 95.

4.7.4 Volume IDs

In a BCC cluster, the volume IDs assigned to BCC-enabled clustered volumes must be unique across all nodes in every peer cluster. Duplicate volume IDs can prevent resources from coming online.

When you BCC-enable a clustered volume, you must manually edit its load script to ensure that its assigned volume ID is unique across all nodes in every peer cluster. You can use the `ncpcon volumes /v` command on each node in every peer cluster to identify the volume IDs in use by all mounted volumes. Compare the results for each server to identify the clustered volumes that have duplicate volume IDs assigned. Modify the load scripts to manually assign unique volume IDs.

For information about volume IDs and how they are assigned in a cluster, see [“Volume ID Requirements”](#) in the [OES 2018 SP1: Novell Cluster Services for Linux Administration Guide](#).

4.7.5 Cluster Containers

BCC supports clusters in the same eDirectory tree. The recommended configuration is to have each peer cluster in the same eDirectory tree but in different OUs (Organizational Units). For information, see [“eDirectory Containers for Clusters”](#) on page 41.

4.7.6 Peer Cluster Credentials

You can manage Linux BCC peer cluster connections and credentials from the CLI:

```
cluster connections [-a]
cluster credentials [peer_cluster]
```


4.8 eDirectory 9.1

Business Continuity Clustering supports eDirectory 9.1 with the latest patches applied. See the [eDirectory 9.1 documentation \(https://www.netiq.com/documentation/edirectory-91/\)](https://www.netiq.com/documentation/edirectory-91/) for information about using and managing eDirectory.

- ◆ [Section 4.8.1, “eDirectory Containers for Clusters,” on page 41](#)
- ◆ [Section 4.8.2, “eDirectory Read/Write Replica on Each IDM Node,” on page 41](#)
- ◆ [Section 4.8.3, “Rights Needed for Installing BCC,” on page 42](#)
- ◆ [Section 4.8.4, “Rights Needed for BCC Management,” on page 42](#)
- ◆ [Section 4.8.5, “Rights Needed by the BCC Administrator to Manage Peer Clusters,” on page 42](#)
- ◆ [Section 4.8.6, “Rights Needed by BCC Drivers,” on page 43](#)
- ◆ [Section 4.8.7, “eDirectory Requirements for Identity Manager,” on page 43](#)

4.8.1 eDirectory Containers for Clusters

Each cluster that you want to add to a business continuity cluster should reside in its own OU level container.

As a best practice for each peer cluster, put its Server objects, Cluster object, Driver objects, and Landing Zone in the same eDirectory container. See [Table 4-3](#) for an example.

Table 4-3 Sample eDirectory Containers for Peer Clusters

ou=cluster1	ou=cluster2	ou=cluster3
You can use the cluster OU as the landing zone.	You can use the cluster OU as the landing zone.	You can use the cluster OU as the landing zone.
Optionally, you can create an OU in the cluster OU for the landing zone.	Optionally, you can create an OU in the cluster OU for the landing zone.	Optionally, you can create an OU in the cluster OU for the landing zone.
ou=cluster1LandingZone	ou=cluster2LandingZone	ou=cluster3LandingZone
cn=cluster1	cn=cluster2	cn=cluster3
c1_node1 (IDM node with read/write access to ou=cluster1)	c2_node1 (IDM node with read/write access to ou=cluster2)	c3_node1 (IDM node with read/write access to ou=cluster3)
c1_node2	c2_node2	c3_node2
c1_node3	c3_node3	c3_node3
cluster1BCCDriverSet	cluster2BCCDriverSet	cluster3BCCDriverSet
<ul style="list-style-type: none"> ◆ c1toc2BCCDriver ◆ c1toc3BCCDriver 	<ul style="list-style-type: none"> ◆ c2toc1BCCDriver 	<ul style="list-style-type: none"> ◆ c3toc1BCCDriver

4.8.2 eDirectory Read/Write Replica on Each IDM Node

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.

4.8.3 Rights Needed 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 that installs BCC must have the eDirectory credentials necessary to extend the schema.

If the eDirectory administrator user name or password contains special characters (such as \$, #, and so on), you might need to escape each special character by preceding it with a backslash (\) when you enter credentials for some interfaces.

4.8.4 Rights Needed for BCC Management

The BCC Administrator user should have at least Read and Write rights to the All Attribute Rights property on the Cluster object of each peer cluster. Before you install BCC, create the BCC Administrator user and group identities in eDirectory to use when you manage the BCC.

The following trustee settings are recommended for the BCC Administrator user on the Cluster object of each peer cluster:

Property Name	Assigned Rights	Inherit	Description
ACL	None	No	Explicitly removing the rights for the ACL property ensures that no rights flow from eDirectory to the file system.
[All Attributes Rights]	Compare, Read, Write	Yes	Read and Write are required.
[Entry Rights]	Create, Delete	Yes	The Create right allows the trustee to create new objects below the container and also includes the Browse right. The Delete right allows the trustee to delete the target from the directory.

For information, see [Section 5.3, “Configuring a BCC Administrator User and Group,” on page 55.](#)

4.8.5 Rights Needed by the BCC Administrator to Manage Peer Clusters

The BCC Administrator user is not automatically assigned the rights necessary to manage all aspects of each peer cluster. When you manage individual clusters, you must log in as the Cluster Administrator user, or as administrator-equivalent to this 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.

You can assign the BCC Administrator user as an administrator-equivalent account for each peer cluster by configuring the following for the user account:

- ◆ Give the user the Supervisor right to the Server object of each server in the cluster.
- ◆ Linux-enable the user account with Linux User Management (LUM).
- ◆ Make the user a member of a LUM-enabled administrator group that is associated with the servers in the cluster.

For information about configuring permissions for cluster administrator-equivalent users, see “[Configuring Additional Administrators](#)” in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.

4.8.6 Rights Needed by BCC Drivers

Each Identity Manager 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 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.

You can do this by making the Identity Manager Driver object the security equivalent to the BCC Administrator User object after you create the driver.

4.8.7 eDirectory Requirements 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.

IMPORTANT: Full eDirectory replicas are required. Filtered eDirectory replicas are not supported.

4.9 Identity Manager 4.7

The IDM references in this guide refers to the IDM Bundle Edition that is bundled with the Open Enterprise Server. If you already have a subscription or purchased the Identity Manager Standard Edition or Advanced Edition, you can use it with BCC.

Business Continuity Clustering supports Identity Manager 4.7 with the latest patches applied. Identity Manager is installed on one node in each peer cluster. Identity Manager 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 clusters.

IMPORTANT: Before you attempt to modify the BCC configuration or manage the BCC-enabled cluster resources, ensure that the Identity Manager node in each peer cluster is active and Identity Manager is running properly.

- ♦ [Section 4.9.1, “Downloading the Identity Manager and Patches,” on page 44](#)
- ♦ [Section 4.9.2, “Credential for Drivers,” on page 44](#)
- ♦ [Section 4.9.3, “IDM Components,” on page 44](#)
- ♦ [Section 4.9.4, “BCC Cluster Resource Synchronization Template for the eDirectory Driver,” on page 44](#)

4.9.1 Downloading the Identity Manager and Patches

The Identity Manager 4.7 for OES 2018 SP1 allows you to use the Identity Manager software, the eDirectory driver, and the management tools for iManager.

Go to the [Micro Focus Downloads Web site \(http://download.novell.com/\)](http://download.novell.com/) and download the following:

- ♦ [Identity Manager 4.7 for OES 2018 SP1 \(https://download.novell.com/Download?buildid=qWGyyy5jY1g~\)](https://download.novell.com/Download?buildid=qWGyyy5jY1g~)
- ♦ Patches are available on [Micro Focus Patch Finder \(https://download.novell.com/patch/finder/\)](https://download.novell.com/patch/finder/). You must apply all maintenance patches for Identity Manager 4.7 on the IDM node in each of the peer clusters before you install or upgrade BCC on any of the nodes.

4.9.2 Credential for Drivers

The BCC License contains the credential that you need to use Identity Manager drivers beyond the evaluation period. In the Identity Manager interface in iManager, you can enter the credential as you create each driver for BCC, or you can put the credential in a file that you point to.

4.9.3 IDM Components

Before you install Business Continuity Clustering on any cluster node in any peer cluster, you must install the following IDM components on one node in every peer cluster:

- ♦ Identity Manger Metdirectory Server (the IDM engine and Identity Vault)
- ♦ eDirectory Driver
- ♦ iManager Plug-Ins for Identity Manager

For information about installing the IDM components used by BCC, see [Section 5.2.2, “Installing and Configuring the Identity Manager Components Used by BCC,” on page 51](#). For general instructions, see “[Installing Identity Manager](#)” in the *NetIQ Identity Manager Integrated Installation Guide*.

4.9.4 BCC Cluster Resource Synchronization Template for the eDirectory Driver

The BCC Cluster Resource Synchronization template is applied to the eDirectory driver to create BCC-specific drivers that automatically synchronize BCC configuration information between the Identity Manager nodes in peer clusters. The template is added on the iManager/Identity Manager node when you install the BCC IDM module (`novell-business-continuity-cluster-idm.rpm`).

4.10 iManager 3.1

Business Continuity Clustering supports iManager 3.1 with the latest patches applied. iManager must be installed and running on the Identity Manager node in each peer cluster.

The iManager installation also installs the following products:

- ♦ Tomcat
- ♦ Java
- ♦ Novell International Cryptographic Infrastructure (NICI)

The following iManager plug-ins are required for NCS, BCC, and IDM management:

- ♦ iManager Base Content
- ♦ iManager Framework
- ♦ iManager Framework Content
- ♦ Storage Shared (`storagemgmt.npm`, the common code for storage-related plug-ins)
- ♦ Storage Management (`nssmgmt.npm`)
- ♦ Cluster Services (`ncsmgmt.npm`)
- ♦ Identity Manager plug-in

For information about installing and using iManager, see the [iManager documentation Web site \(https://www.netiq.com/documentation/imanager-31/\)](https://www.netiq.com/documentation/imanager-31/).

4.11 Storage-Related Plug-Ins for iManager 3.1

The Clusters plug-in (`ncsmgmt.npm`) provides a BCC management interface when BCC is installed on the cluster you are managing. You must install the Clusters plug-in and the Storage Shared plug-in (`storagemgmt.npm`).

IMPORTANT: The Storage Shared plug-in module (`storagemgmt.npm`) contains common code required by all of the other storage-related plug-ins. Ensure that you install the `storagemgmt.npm` first before you install 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 ensure that the common code works for all plug-ins.

Other storage-related plug-ins are Novell Storage Services (NSS) Storage Management (`nssmgmt.npm`), AFP (`afpmgmt.npm`), CIFS (`cifsmgmt.npm`), and Distributed File Services (`dfsmgmt.npm`). NSS is required in order to use shared NSS pools as cluster resources. The other services are optional.

The plug-ins are delivered in the OES 2018 SP1 ISO images. The plug-ins are also available as the Novell Storage Services Plug-in for iManager download file on the [Micro Focus Downloads Web site \(http://download.novell.com/\)](http://download.novell.com/).

To install or upgrade the Clusters plug-in:

- 1 If you upgraded from OES 2018 and any of the storage-related plug-ins are already installed, uninstall all of them, including `storagemgmt.npm`.
- 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 the Storage Shared plug-in first, then install the other storage-related plug-ins that you need.
- 4 Restart Tomcat by entering the following command at a command prompt:

```
rcnovell-tomcat restart
```

Or

```
systemctl restart novell-tomcat.service
```

4.12 Shared Disk Systems

For Business Continuity Clustering, a shared disk storage system is required for each peer cluster in the business continuity cluster. See [“Shared Disk Configuration Requirements”](#) in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.

Each peer cluster contains shared disks for its own resources plus shared disks for resources in peer clusters that are assigned to fail over to it. You must mirror the data between the sites as described in [Section 4.13, “Mirroring Shared Disk Systems between Peer Clusters,”](#) on page 46.

4.13 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 149.

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

4.14 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 it 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.

4.15 Link Speeds

For real-time mirroring, link latency is the essential consideration. For best performance, use dedicated links and provide bandwidth that can handle the amount of data being transferred between the mirrored devices.

Many factors should be considered for distances greater than 200 kilometers, including the following:

- ♦ The amount of data being transferred
- ♦ The bandwidth of the link
- ♦ Whether or not snapshot technology is being used for data replication

4.16 Ports

If you are using a firewall, the ports must be opened for Small Footprint CIM Broker (SFCB) and the Identity Manager drivers.

Table 4-4 Ports for the BCC Setup

Product	Default Port
SFCB	5989 (secure)
eDirectory driver	8196
Cluster Resources Synchronization driver	2002 (plus the ports for additional instances)

4.17 Web Browser

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

For information about browsers supported by iManager, see the [NetIQ iManager Installation Guide](#).

- ♦ [Section 4.17.1, “Web Browser Language Setting,” on page 47](#)
- ♦ [Section 4.17.2, “Web Browser Character Encoding Setting,” on page 48](#)

4.17.1 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 [iManager documentation \(http://www.netiq.com/documentation/imanager-31/\)](http://www.netiq.com/documentation/imanager-31/) for information about supported languages.

4.17.2 Web Browser Character Encoding Setting

Supported language codes are Unicode (UTF-8) compliant. To avoid display problems, ensure that 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 Firefox 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.

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

4.18 What's Next

For new installations of BCC, see the following:

- ♦ [Chapter 5, "Installing Business Continuity Clustering," on page 49](#)

For upgrades from BCC 2.x on OES versions (OES 11 SP2, OES 11 SP3, OES 2015, or OES 2015 SP1) clusters, see the following:

- ♦ [Appendix G, "Upgrading to Identity Manager 4.7," on page 177](#)
- ♦ [Appendix H, "Upgrading to BCC 2.6 on OES 2018 SP1," on page 179](#)

5 Installing Business Continuity Clustering

This section describes how to install, set up, and configure Business Continuity Clustering for Open Enterprise Server (OES) 2018 Support Pack 1 (SP1) to meet your specific needs.

- ♦ [Section 5.1, “Overview of the BCC Installation and Configuration,” on page 49](#)
- ♦ [Section 5.2, “Installing iManager and Identity Manager on One Node in Each Peer Cluster,” on page 50](#)
- ♦ [Section 5.3, “Configuring a BCC Administrator User and Group,” on page 55](#)
- ♦ [Section 5.4, “Installing and Configuring the Business Continuity Clustering Software,” on page 59](#)
- ♦ [Section 5.5, “Installing the BCC Cluster Resource Template,” on page 62](#)
- ♦ [Section 5.6, “Configuring the Device Rescan for BCC Migration of Resources,” on page 62](#)
- ♦ [Section 5.7, “Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software,” on page 63](#)
- ♦ [Section 5.8, “What’s Next,” on page 66](#)

5.1 Overview of the BCC Installation and Configuration

The BCC installation and configuration involves the following tasks:

1. Set up your Novell Cluster Services clusters and ensure that they are functioning properly.
2. On one node in each peer cluster, install iManager and Identity Manager.
3. Configure a BCC Administrator user and group, and enable them for Linux with Linux User Management (LUM).
4. Install BCC software on every node in each peer cluster. Install the Identity Manager driver template only on the nodes where you installed iManager and Identity Manager.
5. Configure BCC software on every node in each peer cluster.
6. Create Identity Manager driver sets and drivers for BCC. See [Chapter 6, “Configuring the Identity Manager Drivers for BCC,” on page 67](#).
7. Enable BCC for peer clusters. See [Chapter 7, “Configuring BCC for Peer Clusters,” on page 81](#).
8. Enable BCC for cluster resources. See [Chapter 8, “Configuring BCC for Cluster Resources,” on page 91](#).

5.2 Installing iManager and Identity Manager on One Node in Each Peer Cluster

On one node in each peer cluster, install iManager and Identity Manager. This node will be referred to as the IDM node throughout the guide.

Each IDM node must be online in its peer cluster and Identity Manager must be running properly whenever you attempt to modify the BCC configuration or manage the BCC-enabled cluster resources.

- ♦ [Section 5.2.1, “Installing iManager and the Plug-Ins Used by BCC,” on page 50](#)
- ♦ [Section 5.2.2, “Installing and Configuring the Identity Manager Components Used by BCC,” on page 51](#)
- ♦ [Section 5.2.3, “Installing the Identity Manager Plug-in for iManager,” on page 54](#)
- ♦ [Section 5.2.4, “Adding Peer Clusters to Your My Clusters List,” on page 54](#)

5.2.1 Installing iManager and the Plug-Ins Used by BCC

1 Install the iManager Server software on one node in each peer cluster:

1a Log in to the server as the `root` user.

1b In YaST, select **Open Enterprise Server > OES Install and Configuration**.

1c On the Software Management page, scroll down to select the **iManager** pattern.

1d Click **Accept**.

1e On the OES Configuration page, enable configuration of iManager, then click the iManager link.

1f On the iManager Configuration page, select **Don't install plug-ins now**, then click **Next**.

1g Click **Next** and complete the installation.

2 In a Web browser, access iManager, then log in to iManager as an administrator user.

```
https://server_ip_address/nps/iManager.html
```

Replace `server_ip_address` with the IP address or DNS name of the node where you installed iManager in [Step 1](#).

3 On the iManager Management page, install the plug-in modules.

For instructions, see “[Downloading and Installing Plug-in Modules](https://www.netiq.com/documentation/imanager-31/imanager_admin/data/b8qrsq0.html)” (https://www.netiq.com/documentation/imanager-31/imanager_admin/data/b8qrsq0.html) in the *iManager 3.1 Administration Guide* (https://www.netiq.com/documentation/imanager-31/imanager_admin/data/bookinfo.html).

BCC requires the following iManager plug-ins to be installed on the IDM node in each peer cluster. Install the plug-ins in the following order:

- ♦ iManager Base Content
- ♦ iManager Framework
- ♦ iManager Framework Content
- ♦ Storage Shared (`storagemgmt.npm`, the common code for storage-related plug-ins)
- ♦ Storage Management (`nssmgmt.npm`)
- ♦ Cluster Services (`ncsmgmt.npm`)
- ♦ Any other plug-ins

You will install the Identity Manager plug-ins later as part of the Identity Manager software installation.

- 4 View a list of installed plug-ins.
- 5 Exit iManager and close the browser.
- 6 Restart Tomcat. Enter

```
rcnovell-tomcat restart
```

Or

```
systemctl restart novell-tomcat.service
```

- 7 Repeat this procedure on the IDM node in each peer cluster in turn.
- 8 Continue with [Section 5.2.2, “Installing and Configuring the Identity Manager Components Used by BCC,”](#) on page 51.

5.2.2 Installing and Configuring the Identity Manager Components Used by BCC

The procedure in this section installs and configures only the Identity Manager components used by BCC.

Install Identity Manager on the same node in each peer cluster where you installed iManager:

- 1 Log in to the node as the `root` user, then open a terminal console.
- 2 Securely copy the Identity Manager ISO file (`Identity_Manager_4.7_Linux.iso`) to the first node in the peer cluster.
- 3 Mount the `downloaded.iso`.
- 4 From the root directory of the `.iso` file, run the following command.

```
./install.sh
```

- 5 Read through the license agreement and enter `y` to accept the license agreement
- 6 Decide the Identity Manager server edition you want to install. Enter `y` for Advanced Edition and `n` for Standard Edition.
- 7 From the list of components available for installation, select **Identity Manager Engine** component to install engine.
- 8 Navigate to the location where you mounted the `Identity_Manager_4.7_Linux.iso` file.
- 9 Specify the following command at the command line to run the `configure.sh` script:

```
./configure.sh
```

- 9a Decide whether you want to perform a typical configuration or a custom configuration. Enter `'2'` to continue with the custom configuration. Again, press `'Enter'` to confirm.
- 9b Specify the component you want to configure. Enter `'1'` to continue with the Identity Manager Engine configuration. Again, press `'Enter'` to confirm.
- 9c It prompts for `Do you want to set a common password (y/n)?`, enter `'y'` to continue.
- 9d Enter a common password and then re-enter the same password to confirm.
- 9e Specify the component you want to install. Enter `'3'` to continue with the `Add to an Identity Vault existing on local machine configuration mode`. Again, press `'Enter'` to confirm.

- 9f** It prompts for Enter Identity Vault Tree Name, press 'Enter' to continue with the default tree name or enter tree name.
- 9g** It prompts for Enter Identity Vault Administrator name, press 'Enter' to continue with the default administrator name (cn=admin,ou=sa,o=system) or enter the administrator name.
- 9h** It prompts for Enter Identity Vault Administrator password, enter password.
- 9i** It prompts for Enter LDAP SSL port, press 'Enter' to continue with the default port (636) or enter the port number.
- 9j** It prompts for Do you want to install a new driverset (y/n), enter 'y' to continue.
- 9k** It prompts for Enter Identity Vault driver set name, press 'Enter' to continue with the default driver set name (driverset1) or enter the driver set name.
- 9l** It prompts for Enter Identity Vault driver set deploy context, press 'Enter' to continue with the default context (o=system) or enter the driver set deploy context.
- 9m** It prompts for Enter custom driverset ldif file path, press 'Enter' to continue with the default file path (/opt/sample-driverset.ldif) or enter the file path of the sample driver set available at IDM/LDIF/sample-driverset.ldif.

For more information on configuring your component, see [Understanding the Configuration Parameters](#) in the [NetIQ Identity Manager Setup Guide for Linux](#).

The sample Identity Manager configuration is as follows:

```
# ./configure.sh
Verifying installed components...

#####
                    Identity Manager Configuration
                    Wed Jun  6 10:38:00 IST 2018
#####

Select the configuration mode. Typical configuration is for new installation
and demo setup. Custom configuration is for advanced users.

  1 ) Typical Configuration
  2 ) Custom Configuration

Choose ONLY ONE option. To continue, press Enter:2
=====
      The selected component(s) are highlighted below. To confirm, press Enter.
To deselect a component, type the value of the component that you want to
deselect.
=====
  1 ) Typical Configuration
  2+) Custom Configuration

Choose ONLY ONE option. To continue, press Enter:

The following Identity Manager components are available for configuration :

  1 ) Identity Manager Engine

Specify the component(s) that you want to configure. To configure multiple
components, specify the values as a comma-separated list [For example, 1, 2,
3]. To confirm, press Enter: 1

=====
      The selected component(s) are highlighted below. To confirm, press Enter.
```

To deselect a component, type the value of the component that you want to deselect.

=====

1+) Identity Manager Engine

Specify the component(s) that you want to configure. To configure multiple components, specify the values as a comma-separated list [For example, 1, 2, 3]. To confirm, press Enter:

Do you want to set a common password (y/n)? [y]:

Enter a common password:

Confirm Password:

Select the configuration mode :

- 1) Create a new Identity Vault
- 2) Add to an Identity Vault existing on remote machine
- 3) Add to an Identity Vault existing on local machine

Specify the component(s) that you want to install. To install multiple components, specify the values as a comma-separated list [For example, 1, 2, 3]. To confirm, press Enter: 3

=====

The selected component(s) are highlighted below. To confirm, press Enter. To deselect a component, type the value of the component that you want to deselect.

=====

- 1) Create a new Identity Vault
- 2) Add to an Identity Vault existing on remote machine
- 3+) Add to an Identity Vault existing on local machine

Specify the component(s) that you want to install. To install multiple components, specify the values as a comma-separated list [For example, 1, 2, 3]. To confirm, press Enter:

Enter Identity Vault Tree Name [blr*_*_*_*_**_tree]:IDM_PATCH_OES

Enter Identity Vault Administrator name
[cn=admin,ou=sa,o=system]:cn=admin,o=acme

Enter Identity Vault Administrator password:

Enter LDAP SSL port [636]:

Do you want to install a new driverset (y/n) [y]:

Enter Identity Vault driver set name [driverset1]:

Enter Identity Vault driver set deploy context [o=system]:o=acme

Enter custom driverset ldif file path [/opt/sample-driverset.ldif]:/tmp/sample-driverset.ldif

Refer log for more information at /var/opt/netiq/idm/log/idmconfigure.log

- 10 Repeat this procedure on the IDM node in each peer cluster in turn.
- 11 When Identity Manager is running on one node in every peer cluster, continue with [Section 5.2.3, “Installing the Identity Manager Plug-in for iManager,”](#) on page 54.

5.2.3 Installing the Identity Manager Plug-in for iManager

On the IDM node in each peer cluster, ensure that the Identity Manager plug-ins are installed:

- 1 In a Web browser, access iManager running on the IDM node in the peer cluster, then log in to iManager as an administrator user.
- 2 Verify that the Identity Manager plug-in has been added to iManager (an icon appears in the toolbar, or iManager opens by default to the Identity Manager page):
 - ♦ If Identity Manager is present, you are done with this task on this node. Repeat this procedure on the IDM node in each peer cluster in turn, then continue with [Section 5.2.4, “Adding Peer Clusters to Your My Clusters List,”](#) on page 54.
 - ♦ If the Identity Manager plug-in is not present, continue with the next steps to install it.
- 3 In iManager, click the **Configure** icon in the toolbar.
- 4 Click **Plug-In Installation > Available Novell Plug-In Modules**.
- 5 Select **NetIQ Identity Manager Plug-Ins for IDM 4.7**, then click **Install**.
If this option is not in the list, click **Add**, browse to select the Identity Manager .npm file, then click **OK**.
- 6 Read the License Agreement, select **I Agree**, then click **OK**.
- 7 After the installation is complete, click **Close** twice.
- 8 Log out of iManager and close the browser.
- 9 Restart Tomcat. Enter

```
rcnovell-tomcat restart
```


Or

```
systemctl restart novell-tomcat.service
```
- 10 Repeat this procedure on the IDM node in each peer cluster in turn.
- 11 After the iManager plug-in is installed for the iManager instance running on the IDM node in each peer cluster, continue with [Section 5.2.4, “Adding Peer Clusters to Your My Clusters List,”](#) on page 54.


5.2.4 Adding Peer Clusters to Your My Clusters List

Add the peer clusters to the My Clusters page on the iManager instance on the IDM node in each peer cluster:

- 1 Log in to iManager on the IDM node as a cluster administrator.
- 2 In **Roles and Tasks**, select **Clusters > My Clusters**.
The list of clusters is initially empty.



- 3 Click **Add** to open the eDirectory browser pop-up window.
- 4 Browse the tree where you are currently logged in to locate and select the Cluster objects for each of the peer clusters, then click **OK**.
Newly selected clusters are added to your personalized list.



- 5 Repeat the setup for the iManager instance on the IDM node in each peer cluster.
- 6 After a **My Clusters** list is set up in iManager on the IDM node on each peer cluster, continue with [Section 5.3, “Configuring a BCC Administrator User and Group,”](#) on page 55.

5.3 Configuring a BCC Administrator User and Group

You must 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 cluster.

Perform the following tasks to configure the BCC Administrator user and group:

- ♦ [Section 5.3.1, “Accessing iManager,”](#) on page 55
- ♦ [Section 5.3.2, “Creating the BCC Group and Administrator User,”](#) on page 56
- ♦ [Section 5.3.3, “Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects,”](#) on page 57
- ♦ [Section 5.3.4, “Adding the BCC Administrator User to the ncsgroup on Each Cluster Node,”](#) on page 58

5.3.1 Accessing iManager

- 1 Launch a Web browser and enter the URL for iManager:

`https://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 installed (that is, the IDM node).

- 2 Specify the administrator user name and password.

- 3 Specify the IP address of the LDAP server in the tree.
- 4 Click **Login**.

5.3.2 Creating the BCC Group and Administrator User

Before you configure BCC in the cluster, you must create a BCC group (`bccgroup`) and BCC Administrator user (`bccadmin`). Members of the group include the BCC Administrator user and the UNIX workstation objects of each node in every peer cluster. The group must be enabled for Linux User Management (LUM). The group allows the inter-cluster communication to function properly.

IMPORTANT: Linux User Management (LUM) requires case-insensitive names by default. The names you specify must be in all lowercase.

To use mixed case for the BCC group and user names, you must enable the Case Sensitive option in LUM before you attempt to create the BCC group and user.

- 1 In iManager, select the **Roles and Tasks** view.
- 2 Create a BCC group, such as `bccgroup`.
 - 2a Select **Directory Administration > Create Object**.
 - 2b On the Create Object page, select **Group**, then click **OK**.
 - 2c Specify the information for the group, then click **OK**.
- 3 Create a BCC Administrator user, such as `bccadmin`.
 - 3a Select **Directory Administration > Create Object**.
 - 3b On the Create Object page, select **User**, then click **OK**.
 - 3c Specify the information for the user, then click **OK**.
- 4 Add the BCC Administrator user to the BCC group.
 - 4a Select **Directory Administration > Modify Object**.
 - 4b Select the BCC group, then click **OK**.
 - 4c On the group's Properties page, select the **Members** tab.
 - 4d Add the BCC Administrator user as a member of the BCC group.
- 5 Enable the group for Linux.
 - 5a Select **Linux User Management > Enable Groups for Linux**.
 - 5b Browse to select the `bccgroup`, then click **OK**.
 - 5c Enable the group for Linux.

Ensure that you do the following when you LUM-enable `bccgroup`:

- ♦ On the Select Groups page, select the **LUM enable all users in group** option.
- ♦ On the Select Workstations page, add all **UNIXWorkstation** objects for all BCC cluster nodes in all peer clusters for the BCC to the `bccgroup`.

IMPORTANT: If you later add a node or reinstall a node in any of the peer clusters in the BCC, its UNIX workstation object must be added manually to this group.

For information about LUM-enabling groups, see “[Managing User and Group Objects in eDirectory](#)” in the *OES 2018 SP1: Linux User Management Administration Guide*.

- 6 On every node in every peer cluster, refresh the local cache for LUM-enabled users and groups. Log in as the `root` user, open a terminal console, then enter

```
namconfig cache_refresh
```

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

You need to assign trustee rights to the BCC Administrator user for each cluster you plan to add to the business continuity cluster.

- 1 In iManager, select the **Roles and Tasks** view.
- 2 Select **Rights**, then select **Modify Trustees**.
- 3 Browse and select the Cluster object, then click **OK**.
- 4 Click **OK** to view the trustee information for the Cluster object.
- 5 If the BCC Administrator user is not listed as a trustee, click the **Add** (plus) button for **Add Trustee**, browse and select the User object, then click **OK**.
- 6 Click **Assigned Rights** for the BCC Administrator user.
- 7 Click **Add Property**, select **ACL**, then click **OK**.

The [All Attributes Rights] and [Entry Rights] properties should automatically be listed. Add them if they are not present.

- 8 Assign rights and inherit settings for each property:

Property Name	Assigned Rights	Inherit	Description
ACL	None	No	Explicitly removing the rights for the ACL property ensures that no rights flow from eDirectory to the file system.
[All Attributes Rights]	Compare, Read, Write	Yes	Read and Write are required.
[Entry Rights]	Create, Delete	Yes	The Create right allows the trustee to create new objects below the container and also includes the Browse right. The Delete right allows the trustee to delete the target from the directory.

For example:

Modify Trustees ?

Object name: cluster1.clusters.example
Trustee name: bccadmin.example

Remove Selected Add Property

Property Name	Assigned Rights	Inherit
<input type="checkbox"/> ACL	<input type="checkbox"/> Supervisor <input type="checkbox"/> Compare <input type="checkbox"/> Read <input type="checkbox"/> Write <input type="checkbox"/> Self <input type="checkbox"/> Dynamic <input type="checkbox"/> Nested	<input type="checkbox"/>
<input type="checkbox"/> [All Attributes Rights]	<input type="checkbox"/> Supervisor <input checked="" type="checkbox"/> Compare <input checked="" type="checkbox"/> Read <input checked="" type="checkbox"/> Write <input type="checkbox"/> Self <input type="checkbox"/> Dynamic <input type="checkbox"/> Nested	<input checked="" type="checkbox"/>
<input type="checkbox"/> [Entry Rights]	<input type="checkbox"/> Supervisor <input type="checkbox"/> Browse <input checked="" type="checkbox"/> Create <input type="checkbox"/> Rename <input checked="" type="checkbox"/> Delete <input type="checkbox"/> Dynamic <input type="checkbox"/> Nested	<input checked="" type="checkbox"/>

Done Cancel

- 9 Click **Done** to save your changes.
- 10 Repeat [Step 2](#) through [Step 9](#) for the Cluster objects of each peer cluster in your business continuity cluster.

5.3.4 Adding the BCC Administrator User to the ncsgroup on Each Cluster Node

In order for the BCC Administrator user to gain access to the cluster administration files (`/admin/novell/cluster`) on other Linux cluster nodes in your BCC, you must add that user to the Novell Cluster Services administration group (such as `ncsgroup`) on each cluster node.

- 1 Log in as `root` and open the `/etc/group` file.
- 2 Find either of the following lines:


```
ncsgroup::107:
```

or

```
ncsgroup::107:bccd
```

The file should contain one of the above lines, but not both.
- 3 Depending on which line you find, edit the line to read as follows:

```
ncsgroup::107:bccadmin
```

or

```
ncsgroup::107:bccd,bccadmin
```

- 4 Replace `bccadmin` with the BCC Administrator user you created.
 Notice the group ID number of the `ncsgroup`. In this example, the number 107 is used. The actual number is the same on each node in a given cluster; it might be different for each cluster.
- 5 After saving the `/etc/group` file, execute the `id` command from a shell.
 For example, if you named the BCC Administrator user `bccadmin`, enter `id bccadmin`.
 The `ncsgroup` should appear as a secondary group of the BCC Administrator user.

5.4 Installing and Configuring the Business Continuity Clustering Software

Use YaST 2 to install the BCC software on your OES 2018 SP1 cluster nodes.

- ◆ Install the Business Continuity Clustering engine software on each cluster node in each of the peer clusters that will be part of a business continuity cluster. You install the software on the nodes of one cluster at a time.
- ◆ Install the BCC-specific Identity Manager template for iManager on an OES 2018 SP1 Linux server where you have installed iManager.

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

IMPORTANT: Before you begin, ensure that your setup meets the requirements specified in [Chapter 4, “Installation Requirements for BCC,” on page 35](#). The BCC Administrator user and group must already be configured as specified in [Section 5.3, “Configuring a BCC Administrator User and Group,” on page 55](#).

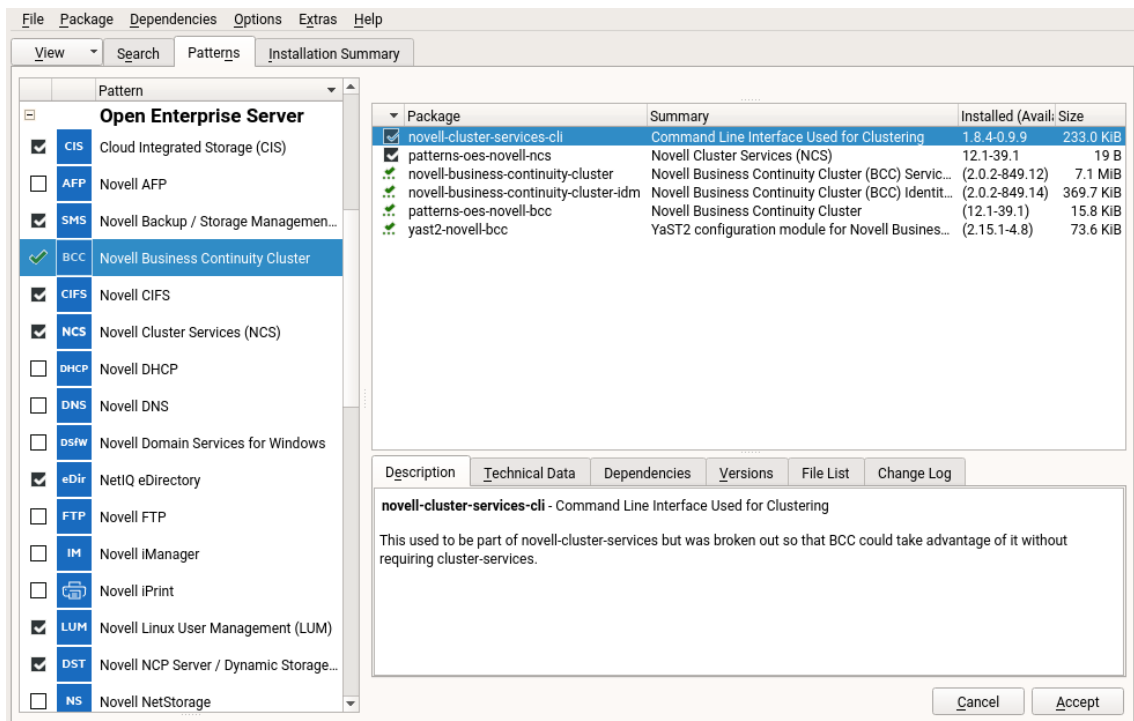
Perform the following tasks in every peer cluster that you want to include in the business continuity cluster:

- ◆ [Section 5.4.1, “Installing the Business Continuity Clustering,” on page 59](#)
- ◆ [Section 5.4.2, “Configuring BCC Software,” on page 61](#)

5.4.1 Installing the Business Continuity Clustering

Perform the following tasks for each of the nodes in every peer cluster:

- 1 Log in to the server as the `root` user.
- 2 In YaST, select **Open Enterprise Server > OES Install and Configuration**.
- 3 On the Software Management page, select the **Novell Business Continuity Cluster** pattern.



This selects the following packages for installation:

```
novell-business-continuity-cluster
novell-business-continuity-cluster-idm
novell-cluster-services-cli
patterns-oes-novell-bcc
patterns-oes-novell-ncs
yast2-novell-bcc
```

- 4 On the non-IDM nodes in the peer cluster, click **Details** to view the list of BCC packages, deselect the `novell-business-continuity-cluster-idm.rpm` package, then click **Finish**.

The package for Identity Manager installs the BCC driver template. This package is needed only on the IDM node in each peer cluster.

- 5 On the Patterns page, click **Accept**, then click **OK** to confirm the Automatic Changes message and resolve dependencies.
- 6 Continue with the installation.
- 7 After the packages are installed, exit YaST.
- 8 Repeat the installation on every node in the peer cluster in turn.
- 9 Repeat the process for each peer cluster in turn.
- 10 Continue with [Section 5.4.2, "Configuring BCC Software,"](#) on page 61.

5.4.2 Configuring BCC Software

Perform the following tasks for each node in every peer cluster:

- 1 Log in to the server as the `root` user.
- 2 Use one of the following methods to open the BCC Configuration page:
 - ♦ In YaST, select **Miscellaneous > Novell-BCC**.
 - ♦ At a command prompt, enter

```
yast2 novell-bcc
```
- 3 The **Novell Business Continuity Cluster (BCC) Configuration** page is displayed and you are prompted to configure LDAP:
 - 3a Click **Continue**.
 - 3b In the **eDirectory Tree** dialog box, specify the administrator user password that you used when you installed the operating system.
 - 3c Click **OK**.
- 4 Specify the **Directory Server Address** by selecting the IP addresses of the master eDirectory server and the local server.
- 5 Select **Start Business Continuity Cluster Services Now** to start the software immediately following the configuration process.

If you deselect the option, you must manually start BCC services later by using the following command:

```
rcnovell-bcc start
```

Or

```
systemctl start novell-bcc.service
```
- 6 Click **Next**.
- 7 Review your setup on the **Novell Business Continuity Cluster (BCC) Configuration Summary** page, then click **Next** to continue the configuration.
- 8 Click **Finish** to save the BCC configuration and exit the tool.
- 9 Verify that the BCC software is running on the server by entering the following at a command prompt:

```
rcnovell-bcc status
```

Or

```
systemctl status novell-bcc.service
```
- 10 Continue with [Section 5.5, “Installing the BCC Cluster Resource Template,”](#) on page 62.

5.5 Installing the BCC Cluster Resource Template

The Cluster Resource template for Business Continuity Clustering is an XML template for the eDirectory driver that is used for synchronizing information about cluster objects between peer clusters. The template is required to configure your business continuity cluster.

You must install the Cluster Resource template on the same server where you installed the Identity Manager Management utilities and iManager. For information, see [Section 4.9, “Identity Manager 4.7,” on page 43](#).

IMPORTANT: If you have not already done so, install the BCC on this server as described in [Section 5.4.1, “Installing the Business Continuity Clustering,” on page 59](#).

The template contains the configuration required to create a BCC driver when applied to the eDirectory driver. It is installed as an Identity Manager template as part of the BCC installation process to the Identity Manager plug-in for iManager on this server. For more information, see [Section 6.4, “Creating a BCC Driver for a BCC Driver Set,” on page 75](#).

Install the BCC cluster resource template on the Identity Manager server in each peer cluster that belongs to the business continuity cluster, then continue with [Section 5.6, “Configuring the Device Rescan for BCC Migration of Resources,” on page 62](#).

5.6 Configuring the Device Rescan for BCC Migration of Resources

When BCC-enabled cluster resources are BCC migrated to a peer cluster, BCC calls a Novell Cluster Services API that automatically runs the `/opt/novell/ncs/bin/device_scan.sh` script on each node that is currently active in the peer cluster. This allows the nodes to recognize the devices and storage objects that are migrated to the cluster.

By default, the script is empty. You must add the Linux shell commands that you need to refresh the nodes in the cluster. Any changes that are made to the script are not overwritten when Novell Cluster Services is upgraded.

WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

On each node in every peer cluster, do the following:

- 1 In a text editor, open the script file `/opt/novell/ncs/bin/device_scan.sh`, add the Linux shell commands that scan your shared devices and storage objects, then save the file.

The following is a sample script:

```
#!/bin/bash

## Logs to /var/log/messages with tag "bccd-scan"
/bin/logger -t bccd-scan "BCC is running script - device_scan.sh -- Scanning
for new devices"

## Rescan devices
rescan-scsi-bus.sh -wcl

## Rescan storage objects, such as for expanded NSS pools
nlvm rescan

## Add multipath command to rebuild maps if applies.
# multipath
```

5.7 Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software

You can install Business Continuity Clustering for Linux core software and Identity Manager management utilities without taking the Business Continuity Clustering software CD or ISO file to different nodes in your cluster. To do this, you must perform the following tasks:

- ♦ [Section 5.7.1, “Creating a YaST Auto-Configuration Profile,” on page 63](#)
- ♦ [Section 5.7.2, “Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media,” on page 64](#)
- ♦ [Section 5.7.3, “Installing and Configuring Business Continuity Clustering on Each Cluster Node,” on page 65](#)
- ♦ [Section 5.7.4, “Removing the NFS Share from Your Server,” on page 65](#)
- ♦ [Section 5.7.5, “Cleaning Up the Business Continuity Clustering Installation Source,” on page 66](#)

5.7.1 Creating a YaST Auto-Configuration Profile

- 1 In a text editor, create a YaST auto-configuration profile XML file named `bccprofile.xml`.

Auto-configuration files are typically stored in the `/var/lib/autoinstall/repository/` directory, but you can use any directory.

The file should appear similar to the example below.

```

<?xml version="1.0"?>
<!DOCTYPE profile SYSTEM "/usr/share/autoinstall/dtd/profile.dtd">
<profile xmlns="http://www.suse.com/1.0/yast2ns" xmlns:config="http://
www.suse.com/1.0/configns">
  <configure>
    <bcc>
      <config-type>New</config-type>
      <start-now>Yes</start-now>
      <cluster-dn>cn=my_cluster.o=novell</cluster-dn>
      <ldap-server>10.1.1.0</ldap-server>
      <ldap-port>389</ldap-port>
      <ldap-secure-port>636</ldap-secure-port>
      <admin-dn>cn=admin.o=novell</admin-dn>
      <admin-password>password</admin-password>
    </bcc>
  </configure>
</profile>

```

Edit the above example to apply to your own specific system settings.

- 2 Copy the XML file you created in [Step 1](#) to each node in the cluster.

Use the same path on each node. You can use the `scp` command to copy the file securely. See the `scp` man page for information on using `scp`.

- 3 Continue with [Section 5.7.2, “Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media,”](#) on page 64.

5.7.2 Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media

- 1 Prepare a directory for an NFS share from within a shell.

To do this, you can either copy the contents of the CD to a local directory, or you can mount the ISO image as a loopback device.

To copy the contents of the CD to a local directory, enter commands similar to the following:

```

mkdir /tmp/bcc_install
cp -r /media/<cd-drive> /tmp/bcc_install

```

To mount the ISO image as a loopback device, enter commands similar to the following:

```

mkdir /mnt/iso
mkdir /tmp/bcc_install
mount path_to_BCC_ISO /tmp/bcc_install -o loop

```

Replace *path_to_BCC_ISO* with the location of the Business Continuity Clustering software ISO image.

- 2 Create an NFS share by opening a shell and running `yast2 nfs_server`. Wait until it is open to continue.
- 3 Select **Start NFS Server**, then click **Next**.
- 4 Click **Add Directory** and enter the following:

```

/tmp/bcc_install

```

- 5 Enter a host wildcard if desired.

- 6 Click **OK**, then click **Finish**.
- 7 Continue with [Section 5.7.3, “Installing and Configuring Business Continuity Clustering on Each Cluster Node,”](#) on page 65.

5.7.3 Installing and Configuring Business Continuity Clustering on Each Cluster Node

You must install BCC software on each cluster node in every cluster that you want to be in the business continuity cluster.

- 1 Create a new YaST software installation source by opening a shell and running `yast2 inst_source`.
- 2 Add **NFS** as the source type.
- 3 Specify the server and directory you entered in [Step 4 on page 64](#), click **OK**, then click **Finish**.
- 4 Install Business Continuity Clustering software by opening a shell and running the following commands in the order indicated:

```
yast2 sw_single -i \  
novell-business-continuity-cluster \  
novell-cluster-services-cli \  
yast2-bcc
```

- 5 Autoconfigure the Business Continuity Clustering software by running the following command from a shell:

```
yast2 bcc_autoconfig path_to_XML_profile
```

Replace `path_to_XML_profile` with the path to the file you created in [Step 1 on page 63](#).

- 6 Remove the installation source you created in [Step 1](#) above by completing the following steps:
 - 6a Open a shell and run `yast2 inst_source`.
 - 6b Select the Business Continuity Clustering installation source, click **Delete**, then click **Finish**.
- 7 Continue with [Section 5.7.4, “Removing the NFS Share from Your Server,”](#) on page 65.

5.7.4 Removing the NFS Share from Your Server

You can remove the Business Continuity Clustering installation directory as an NFS share. After a successful install, it is needed only if you reinstall or uninstall BCC.

- 1 Open a shell and run `yast2 nfs_server`.
- 2 Select **Start Server**, then click **Next**.
- 3 Select the Business Continuity Clustering installation directory, click **Delete**, then click **Finish**.
- 4 Continue with [Section 5.7.5, “Cleaning Up the Business Continuity Clustering Installation Source,”](#) on page 66.

5.7.5 Cleaning Up the Business Continuity Clustering Installation Source

- 1 In a terminal console, run one of the commands below, depending on which method you chose in [Step 1](#) of the procedure in [Section 5.7.2, “Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media,”](#) on page 64.

```
rm -rf /tmp/bcc_install
```

or

```
umount /mnt/iso
```

- 2 Continue with [Section 5.6, “Configuring the Device Rescan for BCC Migration of Resources,”](#) on page 62.

5.8 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 the following steps:

- ♦ [Chapter 6, “Configuring the Identity Manager Drivers for BCC,”](#) on page 67.
- ♦ If you are adding a new cluster to an existing business continuity cluster, follow the instructions in [“Synchronizing Identity Manager Drivers”](#) on page 79 to synchronize the BCC Identity Manager drivers across all of the clusters.
- ♦ [Chapter 7, “Configuring BCC for Peer Clusters,”](#) on page 81.
- ♦ [Chapter 8, “Configuring BCC for Cluster Resources,”](#) on page 91.

6 Configuring the Identity Manager Drivers for BCC

Business Continuity Clustering (BCC) software provides a driver for Identity Manager that is used to synchronize cluster resource information between the clusters in the business continuity cluster. After you install BCC, you must configure the Identity Manager drivers for BCC in order to properly synchronize and manage your business continuity cluster.

IMPORTANT: To assist your planning process, a worksheet is provided in [Appendix B, “Configuration Worksheet for the BCC Drivers for Identity Manager,”](#) on page 141.

- ♦ [Section 6.1, “Understanding the BCC Drivers,”](#) on page 67
- ♦ [Section 6.2, “Prerequisites for Configuring the BCC Drivers for Identity Manager,”](#) on page 72
- ♦ [Section 6.3, “Creating the BCC Driver Set,”](#) on page 73
- ♦ [Section 6.4, “Creating a BCC Driver for a BCC Driver Set,”](#) on page 75
- ♦ [Section 6.5, “Enabling or Disabling the Synchronization of Email Settings,”](#) on page 79
- ♦ [Section 6.6, “Synchronizing Identity Manager Drivers,”](#) on page 79
- ♦ [Section 6.7, “Changing the Identity Manager Synchronization Drivers,”](#) on page 79
- ♦ [Section 6.8, “What’s Next,”](#) on page 79

6.1 Understanding the BCC Drivers

Business Continuity Clustering provides a Cluster Resource Synchronization template that is used with the eDirectory driver in Identity Manager to create the BCC drivers. It is a set of policies, filters, and objects that synchronize cluster resource information between any two of the peer clusters. This template is always used to create drivers for synchronizing information, and must be configured after installing BCC software.

- ♦ [Section 6.1.1, “IDM Driver Set and Driver Guidelines,”](#) on page 68
- ♦ [Section 6.1.2, “Preventing Synchronization Loops for Identity Manager Drivers,”](#) on page 68
- ♦ [Section 6.1.3, “Example 1: Two Peer Clusters,”](#) on page 71
- ♦ [Section 6.1.4, “Example 2: Three Peer Clusters,”](#) on page 71
- ♦ [Section 6.1.5, “Example 3: Four Peer Clusters,”](#) on page 72

6.1.1 IDM Driver Set and Driver Guidelines

Multiple instances of the **Cluster Resource Synchronization** drivers can be added to the same driver set. For example, the driver set has a driver instance for each Identity Manager connection that a given cluster has with another peer cluster.

The BCC drivers are installed and configured on the Identity Manager node in each of the peer clusters in the business continuity cluster. Each of the driver connections has a Publisher channel (sending) and a Subscriber channel (listening) for sharing information between any two peer clusters. The two nodes are not directly connected; they communicate individually with the Identity Manager vault on a port that is assigned for that instance of the driver.

You must assign a unique port for communications between any two peer clusters. The default port in the Cluster Resource Synchronization template is 2002. For each instance of a driver, you can use any ports that are unique and not otherwise allocated. Ensure that the ports are not blocked by the firewall. Examples of port assignments are described elsewhere in this section.

You must specify the same port number for the same driver instance on both cluster nodes. For example, if you specify 2003 as the port number for the Cluster Resource Synchronization driver on one cluster, you must specify 2003 as the port number for the same Cluster Resource Synchronization driver instance on the peer cluster.


The driver needs to have sufficient eDirectory rights in the Cluster context to create, modify, and delete objects related to BCC-enabled cluster resources. This is achieved by making the Driver object security equivalent to the BCC Admin user.

6.1.2 Preventing Synchronization Loops for Identity Manager Drivers

If you have three or more clusters in your business continuity cluster, you should set up synchronization for the 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. Cluster One synchronizes with Cluster Two. 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



In a single-tree scenario with a four-cluster business continuity cluster, Cluster One is an Identity Manager synchronization master in which Cluster One synchronizes data with each of the peer clusters, as illustrated in [Figure 6-4](#).

Figure 6-4 Single-Tree, Four-Cluster Identity Manager Synchronization Scenario



6.1.3 Example 1: Two Peer Clusters

Let's consider a two-cluster business continuity cluster. The Cluster Resource Synchronization driver's Publisher channel in Cluster One communicates with the driver's Subscriber channel in Cluster Two. Conversely, the driver's Publisher channel in Cluster Two communicates with the driver's Subscriber channel in Cluster One. The two clusters send and listen to each other on the same port via the Identity Manager vault, as shown in [Table 6-1](#).

Table 6-1 Two-Cluster Driver Set Example

Cluster Resource	Subscriber Node	
Publisher Node	Cluster One	Cluster Two
Cluster One	Not applicable	CR, port 2002
Cluster Two	CR, port 2002	Not applicable

You install the Cluster Resource Synchronization driver once on Cluster One and once on Cluster Two, as shown in [Table 6-2](#).

Table 6-2 Driver Set Summary for a Two-Cluster Business Continuity Cluster

Driver Instance	Driver Set for Cluster One	Driver Set for Cluster Two
Cluster Resource	C1 to C2, port 2002	C2 to C1, port 2002

6.1.4 Example 2: Three Peer Clusters

If you have more than two clusters in your business continuity cluster, you should set up communications for the drivers in a manner that prevents Identity Manager synchronization loops. Identity Manager synchronization loops can cause excessive network traffic and slow server communication and performance. You can achieve this by picking one of the servers to be the master for the group. Each of the peer clusters' drivers communicates to this node.

For example, let's consider a three-cluster business continuity cluster. You can set up a communications channel for the Cluster Resource Synchronization driver between Cluster One and Cluster Two, and another channel between Cluster One and Cluster Three. Cluster Two does not talk to Cluster Three, and vice versa. You must assign a separate port for each of these communications channels, as shown in [Table 6-3](#) and [Table 6-4](#).

Table 6-3 Three-Cluster Driver Set Example

Cluster Resource	Subscriber Node		
Publisher Node	Cluster One	Cluster Two	Cluster Three
Cluster One (master node)	Not applicable	CR, port 2002	CR, port 2003
Cluster Two	CR, port 2002	Not applicable	No channel
Cluster Three	CR, port 2003	No channel	Not applicable

Table 6-4 Driver Set Summary for a Three-Cluster Business Continuity Cluster

Driver Instance	Driver Set for Cluster One	Driver Set for Cluster Two	Driver Set for Cluster Three
Cluster Resource	C1 to C2, port 2002	C2 to C1, port 2002	C3 to C1, port 2003
Cluster Resource	C1 to C3, port 2003		

6.1.5 Example 3: Four Peer Clusters

When you extend the single-tree example for a four-cluster business continuity cluster, you can set up similar communications channels for the Cluster Resource Synchronization driver between Cluster One and Cluster Two, between Cluster One and Cluster Three, and between Cluster One and Cluster Four. You must assign a separate port for each of these channels, as shown in [Table 6-5](#).

Table 6-5 Four-Cluster Driver Set Example

Cluster Resource	Subscriber Node			
Publisher Node	Cluster One	Cluster Two	Cluster Three	Cluster Four
Cluster One (master node)	Not applicable	CR, port 2002	CR, port 2003	CR, port 2004
Cluster Two	CR, port 2002	Not applicable	No channel	No channel
Cluster Three	CR, port 2003	No channel	Not applicable	No channel
Cluster Four	CR, port 2004	No channel	No channel	Not applicable

You install the drivers on each cluster, with multiple instances in the driver set on Cluster One, but only a single instance in the peer clusters, as shown in [Table 6-6](#).

Table 6-6 Driver Set Summary for a Four-Cluster Business Continuity Cluster

Driver Instance	Driver Set for Cluster One	Driver Set for Cluster Two	Driver Set for Cluster Three	Driver Set for Cluster Four
Cluster Resource	C1 to C2, port 2002	C2 to C1, port 2002	C3 to C1, port 2003	C4 to C1, port 2004
Cluster Resource	C1 to C3, port 2003			
Cluster Resource	C1 to C4, port 2004			

6.2 Prerequisites for Configuring the BCC Drivers for Identity Manager

Before you configure the Identity Manager drivers, ensure that your system meets the requirements in this section.

- ◆ [Section 6.2.1, “Identity Manager,” on page 73](#)
- ◆ [Section 6.2.2, “eDirectory,” on page 73](#)

- [Section 6.2.3, “Landing Zone Container,” on page 73](#)
- [Section 6.2.4, “BCC Admin User and Group,” on page 73](#)

6.2.1 Identity Manager

Before you installed Business Continuity Clustering, you set up and configured the Identity Manager engine and an Identity Manager driver for eDirectory on one node in each cluster. For information, see [Section 5.2, “Installing iManager and Identity Manager on One Node in Each Peer Cluster,” on page 50](#).

The Identity Manager plug-in for iManager requires that eDirectory is running and working properly.

6.2.2 eDirectory

The cluster 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 4.8, “eDirectory 9.1,” on page 41](#).

6.2.3 Landing Zone Container

The landing zone that you specify for drivers must already exist. Typically, this is the OU container you created for the peer cluster’s objects. You can optionally create a separate container in eDirectory specifically for these cluster pool and volume objects.

6.2.4 BCC Admin User and Group

BCC Administrator user and group are used to manage BCC. The driver needs to have sufficient eDirectory rights (see [Section 4.8.6, “Rights Needed by BCC Drivers,” on page 43](#)) in the Cluster context to be able to create, modify, and delete objects related to BCC-enabled cluster resources. This is achieved by making the Driver object security equivalent to the BCC Administrator user.

The procedures in this section assume that you have set up a single BCC Administrator user that is used for all peer clusters. If you have set up a different BCC Administrator user for each peer cluster, then use the corresponding BCC Administrator user when working on drivers for the cluster.

For information about setting up the BCC Administrator user, see [Section 5.3, “Configuring a BCC Administrator User and Group,” on page 55](#).

6.3 Creating the BCC Driver Set

You will create a driver set for each peer cluster, according to the peer connectivity design for your BCC. For information about setting up connectivity to avoid synchronization loops, see [Section 6.1.2, “Preventing Synchronization Loops for Identity Manager Drivers,” on page 68](#).

Repeat the following procedure in each peer cluster to create the cluster’s driver set:


- 1 In iManager, click **Identity Manager > Identity Manager Overview**.
- 2 Browse to select the Identity Manager server in a peer cluster where you want to create drivers.
This is the node in the cluster where you installed the Identity Manager engine and eDirectory driver.
- 3 On the Identity Manager Overview page, click **Driver Sets > New**.



- 4 Type the name of the driver set you want to create for this cluster.

For example, specify *Cluster1 BCC DriverSet*, where *Cluster1* is the name of the cluster where you are configuring a driver instance.

Identity Manager Overview

Specify the container you want to search, then press the search button.

Search in:   



- 5 Browse to select the context that contains the cluster objects for the cluster where you are configuring a driver instance.

For example, *cluster1.clusters.siteA.example*

- 6 Deselect (disable) the **Create a new partition on this driver set** option, click **OK** on the pop-up message, and then click the **OK** button to complete driver set creation and take you to on the Driver Set Overview page.

It is not necessary to create an eDirectory partition for the driver set because the parent cluster container where you are creating the driver set is partitioned.


- 7 Associate the new driver set with the Identity Manager node in this cluster.

7a On the Drive Set Overview page, select the newly created driver set.

7b Select **Servers > Add servers** from the drop-down menu.


Driver Set Overview

Driver Set:



- 7c** Select the Identity Manager node in the cluster.

Driver Set Overview



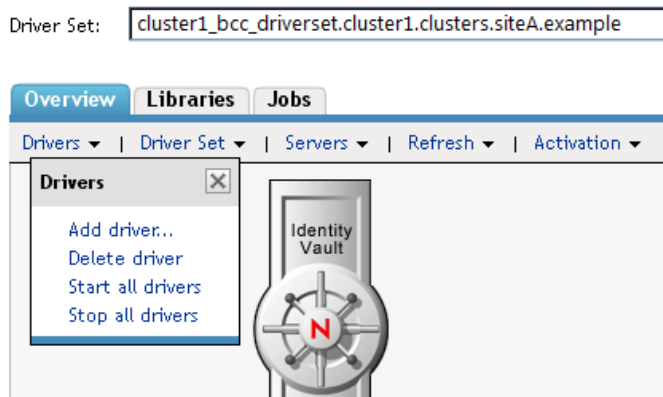
7d Repeat the driver set creation for each peer cluster.

6.4 Creating a BCC Driver for a BCC Driver Set

Repeat the following process for each driver that you want to add to a BCC driver set:



- 1 In iManager, click **Identity Manager > Identity Manager Overview**.
- 2 Browse to select the Identity Manager server in the peer cluster.
This is the node in the cluster where you installed the Identity Manager engine and eDirectory driver.
- 3 If a BCC driver set does not exist for this cluster, create it now.
For information, see [Section 6.3, “Creating the BCC Driver Set,” on page 73](#).
- 4 On the Driver Set Overview page, select the driver set, click **Drivers**, then click **Add driver** from the drop-down menu.

Driver Set Overview



- 5 On the Import Configuration page, specify where you want to place the new driver. Select **In an existing driver set**, verify that the driver set is the one you preselected, then click **Next**.


Where do you want to place the imported configuration?

- In an existing driver set
  
- In a new driver set

6 Import the BCC Cluster Resource Synchronization template.

The template contains the configuration required to create a BCC driver when applied to the eDirectory driver. It is installed as an Identity Manager template as part of the BCC installation process on the Identity Manager node.

- 6a** Browse to select the Identity Manager node in the cluster, then click **Next**.
- 6b** Open the **Show** drop-down menu and select **Configurations not associated with an IDM version**.
- 6c** Open the **Configurations** drop-down menu and select the **BCCClusterResourceSynchronization.xml** file.



Import a configuration into this driver set.

- Import a configuration from the server (.XML file)
Show:
Configurations:

6d Click **Next**.

7 On the BCC Cluster Sync page, specify the driver information.

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:** For this driver instance, specify a unique name to identify its function.
The default name is **BCC Cluster Sync**. We recommend that you indicate the source and destination clusters involved in this driver, such as *Cluster1 to Cluster2 BCC Sync*.
- ◆ **Name of SSL Certificate:** Specify a name for the certificate, such as *Cluster1 to Cluster2 BCC Sync*.
Specify the SSL certificate that was created when you installed OES on the Identity Manager node.
- ◆ **DNS name of the other IDM node:** Specify the DNS name or IP address of the Identity Manager server in the destination cluster for this driver instance. For example, type *10.10.20.21* or type *node1.cluster2.clusters.siteB.example*.
- ◆ **Port number for this connection:** You must specify unique TCP port numbers for each driver instance for a given connection between two clusters. The default port number is 2002 for the cluster resource synchronization.

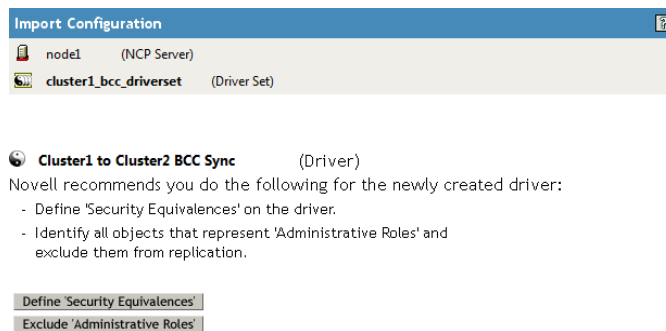
You must specify the same port number for the same template in the destination cluster when you set up the driver instance in that peer cluster. For example, if you specify 2003 as the port number for the resource synchronization driver instance for Cluster1 to Cluster 2, you must specify 2003 as the port number for the Cluster2 to Cluster1 resource synchronization driver instance for the peer driver you create on Cluster2.

- ♦ **Distinguished Name (DN) of this cluster:** Browse to select the Cluster object for this cluster. For example, *cluster1.clusters.siteA.example*. The container name is case-sensitive, so it is best to browse for this name.
- ♦ **Distinguished Name (DN) of the landing zone container for this cluster:** Browse to select the landing zone container where the cluster pool and volume objects in the other cluster are placed when they are synchronized to this cluster. The NCP server objects for the virtual server of a BCC-enabled resource are also placed in the landing zone. The container name is case-sensitive, so it is best to browse for this name.

8 Click **OK**.

The BCC Cluster Resource Synchronization template is imported and merged with the information to create the driver.

This can take a couple of minutes. Be patient and do not interrupt the process. Do not continue until it is complete.



9 Make the Driver object security equivalent to the BCC Administrator User object.

The driver needs to have sufficient eDirectory rights in the Cluster context to create, delete, or modify objects related to BCC-enabled cluster resources. This is achieved by making the driver security equivalent to the BCC Administrator user.

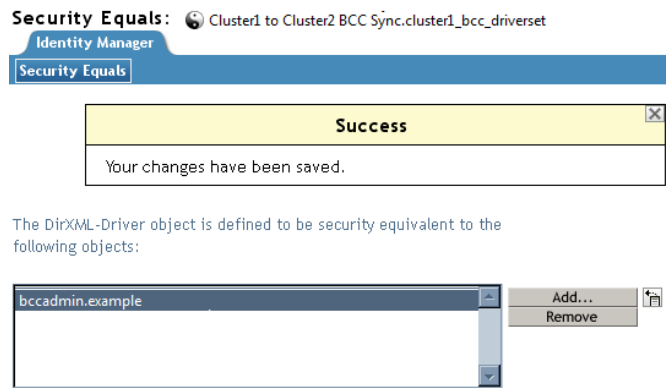
9a On the confirmation page for the driver, click **Define Security Equivalences**.

9b Click **Add**.


9c Browse to select the BCC Administrator User object, then click **OK**.

9d Click **Next**.

The wizard confirms the equivalence of the driver to the selected BCC Administrator user.



- 9e Do not exclude administrator roles. Click **Next**.
In a BCC environment, there is no need to exclude administrator roles.
- 9f Review the summary of the driver configuration, then click **Finish**.
- 10 Verify that the Driver object is set as the security equivalent of the BCC Administrator user.
 - 10a In iManager, select **Directory Administration > Modify Object**.
 - 10b Browse to select the BCC Administrator User object, then click **OK**.
 - 10c Click **OK** to view the properties of the selected user, then select the **Security** tab.
 - 10d On the Security Equal to Me page, verify that the Driver object appears in the **Equivalent to Me** list.



- 10e Click **Cancel** to abandon the Modify Object task.
- 11 Repeat **Step 1** through **Step 9** above on the peer clusters in your business continuity cluster.
This includes creating a new driver and driver set for each cluster.
- 12 After you have configured the BCC drivers on the IDM node in each peer cluster, you must upgrade the drivers to the Identity Manager 4.7.
 - 12a In iManager, click **Identity Manager**, then click **Identity Manager Overview**.
 - 12b Search for the driver sets that you have added, then click the driver set link to bring up the **Driver Set Overview**.
 - 12c Click the red **Cluster Sync** icon, and you should be prompted to upgrade the driver.

6.5 Enabling or Disabling the Synchronization of Email Settings

You can modify the Identity Manager driver filter to enable or disable the synchronization of email settings. For example, you might need to prevent email settings from being synchronized between two peer clusters when you are debugging your BCC solution to isolate problems in a given peer cluster.

6.6 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 Log in to iManager as the BCC Administrator user.
- 2 Go to the Identity Manager Administration page.
- 3 In **Roles and Tasks**, select **Identity Manager > Identity Manager Overview**.
- 4 Search to locate the BCC driver set, then click the driver set link.
- 5 Click the red **Cluster Sync** icon for the driver you want to synchronize, click **Migrate**, then click **Migrate from Identity Vault** from the drop-down menu.
- 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.

6.7 Changing the Identity Manager Synchronization Drivers

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

6.8 What's Next

After the Identity Manager drivers for BCC are configured, you are ready to set up BCC for the clusters and cluster resources. See [Chapter 7, "Configuring BCC for Peer Clusters," on page 81](#).

7 Configuring BCC for Peer Clusters

After you have installed and configured Identity Manager, the Business Continuity Clustering software, and the Identity Manager drivers for BCC, you are ready to set up the Novell Cluster Services clusters to form a business continuity cluster.

IMPORTANT: Identity Manager must be configured and running on one node in each peer cluster before configuring clusters for business continuity. Ensure that the Identity Manager server is part of the cluster and that it is working properly whenever you make BCC configuration changes to the cluster. For information, see [Chapter 6, “Configuring the Identity Manager Drivers for BCC,”](#) on page 67.

Perform the following tasks on each peer Novell Cluster Services cluster that you want to be part of the business continuity cluster:

- ♦ [Section 7.1, “Enabling Clusters for Business Continuity,”](#) on page 81
- ♦ [Section 7.2, “Adding Peer Cluster Credentials,”](#) on page 83
- ♦ [Section 7.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 87
- ♦ [Section 7.4, “Enabling Linux POSIX File Systems to Run on Secondary Clusters,”](#) on page 88
- ♦ [Section 7.5, “Verifying BCC Administrator User Trustee Rights and Credentials,”](#) on page 89
- ♦ [Section 7.6, “Disabling BCC for a Peer Cluster,”](#) on page 89
- ♦ [Section 7.7, “What’s Next,”](#) on page 90

7.1 Enabling Clusters for Business Continuity

You can enable two to four clusters to form a business continuity cluster. Enable BCC for each Novell Cluster Services cluster that you want to add to the business continuity cluster. When you enable a cluster for BCC, BCC adds an Auxiliary class to its Cluster object that provides the attributes **NCC:BCC Peer** and **NCS:BCC Revision**.

- 1 Log in to iManager as the BCC Administrator user.
- 2 Ensure that the BCC-specific Identity Manager drivers are running:
 - 2a Go to the Identity Manager page.
 - 2b Select **Identity Manager Administration > Identity Manager Overview**.
 - 2c Search the eDirectory Container or tree for the BCC-specific Identity Manager drivers.
 - 2d For each driver, click the upper-right corner of the driver icon to see if a driver is started or stopped.
 - 2e If the driver is stopped, start it by selecting **Start**.
- 3 Go to the Roles and Tasks page.
- 4 In **Roles and Tasks**, select **Clusters. > My Clusters**.
- 5 In your personalized My Clusters list, select the check box next to the name of the cluster you want to manage.

If the cluster does not appear in your list, add the cluster to your list. Click **Add**, browse to select the Cluster object, then click **OK**.

- 6 Select **Actions > Edit Properties** to view the cluster's properties.

You can also access the cluster properties with the **Properties** button on the **Cluster Options** tab when you manage the cluster.

- 7 In the **Properties** dialog box, select the **Business Continuity** tab.

- 8 Select the **Enable Business Continuity Features** check box, then click **OK**.

cluster1.clusters.siteA.example ?

Policies **Priorities** **Protocols** **RME Groups** **Business Continuity**

Main | Auto-Failover

Enable or disable business continuity for this cluster. Create, edit, or delete cluster resource script replacement values. View cluster peers or edit peer credentials.

Enable Business Continuity Features

Resource Replacement Scripts 0 Item(s)

New | Edit | Delete

Search Value Replacement Value Description Use Regular Expressions

No Items

Connections 0 Item(s)

Edit

Name	State	User	Address
<i>No Items</i>			

OK Cancel Apply

- 9 Click **OK** when you are prompted to confirm.

A confirmation message is briefly displayed, and a BCC peer connection is added to the **Connections** table. However, because no credentials have been supplied so far, the connection cannot be established.

Connections 1 Item(s)

Edit

Name	State	User	Address
<input type="checkbox"/> cluster1	Invalid Credentials		10.10.10.60

- Repeat [Step 5](#) through [Step 9](#) for each cluster that you want to add to the business continuity cluster.

As additional peer clusters are enabled, the cluster is added in the **Connections** table.

Connections			
Edit			2 Item(s)
Name	State	User	Address
<input type="checkbox"/> cluster1	Invalid Credentials		10.10.10.60
<input type="checkbox"/> cluster2	Invalid Credentials		10.10.20.70

- Wait for the BCC Identity Manager drivers to synchronize.
You can use the `cluster connections` command to list the clusters. The drivers are synchronized when all of the peer clusters are present in the list.
- Continue with [Adding Peer Cluster Credentials](#).

7.2 Adding Peer Cluster Credentials

Clusters must be able to authenticate to themselves and to peer clusters. In order for one cluster to connect to a second cluster, the first cluster must be able to authenticate to the second cluster. For each node, add the authentication credentials (user name and password) of the user that the selected cluster will use to authenticate to a selected peer cluster.

Configure peer cluster credentials on one node in each peer cluster in the business continuity cluster:

- Log in as the `root` user on the cluster node where you want to add peer credentials in this cluster, then open a terminal console.

You can use any node in the cluster to specify credentials.

- At the command prompt, enter

```
cluster connections <cluster_name>
```

For example, this lists the BCC peer clusters and credential information. At this point, the credentials have not been provided and are reported as invalid.

```
cluster connections cluster1
```

```
Connection status for cluster: cluster1
```

Cluster Name	Username	Connection Status
cluster1	<unknown>	Invalid Credentials
cluster2	<unknown>	Invalid Credentials

- Verify that all clusters are present in the list.

If the clusters are not present, the Identity Manager drivers are not synchronized.

If synchronization is in progress, wait for it to complete, then try `cluster connections` again.

If you need to synchronize, see [“Synchronizing Identity Manager Drivers” on page 79](#).

- For each cluster in the list, enter the following command at the command prompt, then enter the `bccadmin` user name and password when you are prompted:

```
cluster credentials <cluster_name>
```

For example, issue the command once for each peer cluster:

```
cluster credentials cluster1
Enter the credentials for the specified cluster. Press CTRL+C to cancel.
Username: bccadmin
Password:
Please wait...
The credentials for cluster cluster1 have been saved.
```

```
cluster credentials cluster2
Enter the credentials for the specified cluster. Press CTRL+C to cancel.
Username: bccadmin
Password:
Please wait...
The credentials for cluster cluster2 have been saved.
```

If you created different BCC Administrator users for each peer cluster, you will provide the respective credentials for each cluster.

- 5 Verify that the peer clusters are recognized as being members of the same BCC. Enter the `cluster connections` command and `cluster view` command to view the status.

For example, on cluster1:

```
cluster connections cluster1

Connection status for cluster: cluster1
  Cluster Name      Username      Connection Status
  cluster1         bccadmin     OK
  cluster2         bccadmin     OK

cluster view

  Cluster cluster1
  This node c1_node1 [epoch 1 master node c1_node2]
  Cluster nodes [c1_node1, c1_node2]
  BCC peer clusters [cluster1, cluster2]
```

If you created different BCC Administrator users for each peer cluster, each peer cluster would report a different user name in the `cluster connections` command.

NOTE: If the cluster view report shows only one peer cluster instead of showing all the peer clusters, execute the command `cluster refresh -p` on that node. For more information, [Section 10.2.1, “Incorrect Status in the Cluster View Report After Enabling BCC on Peer Clusters,” on page 118.](#)

- 6 Add the BCC Administrator user to the `ncsgroup` for the cluster. Repeat the following steps for every node in the peer cluster:

- 6a As the `root` user, open the `/etc/group` file in a text editor.

- 6b Locate the line that reads `ncsgroup`, then modify it to include the `bccadmin` user.

For example, change

```
ncsgroup:!:107:
```

to

```
ncsgroup:!:107:bccadmin
```

For example, change

```
ncsgroup:!:107:bccd
```

to

```
ncsgroup:!:107:bccd,bccadmin
```

The file should contain one of the above lines, but not both.

Notice the group ID number of the `ncsgroup`. In this example, the number 107 is used. This number can be different for each cluster node.

6c Save the `/etc/group` file.

6d At the server console prompt, enter the following to verify that the `bccadmin` user is a member of the `ncsgroup`.

```
id bccadmin
```

7 Repeat the previous steps on a node in each peer cluster in turn.

The nodes in the other peer clusters know about the peer clusters, but they do not yet have the credentials needed to connect.

8 (Optional) Use iManager to verify that the peer clusters are communicating their status.

8a Log in to iManager as the BCC administrator.

8b Select **Clusters > My Clusters**.

BCC-enabled clusters are identified by a check in the BCC column.

The screenshot shows the 'My Clusters' page in iManager. On the left is a navigation pane with categories like 'Archive Versioning', 'Clusters', 'My Resources', 'DHCP (OES Linux)', 'Directory Administration', and 'Distributed File Services'. The main area displays a table of clusters. The table has columns for Status, Name, Master Node, IP Address, Nodes, Epoch, and BCC. Two clusters are listed: 'cluster1' with IP 10.10.10.60 and 'cluster2' with IP 10.10.20.70. Both have a green status icon and a checkmark in the BCC column.

Status	Name	Master Node	IP Address	Nodes	Epoch	BCC
	cluster1	c1_node1	10.10.10.60	2/2	3	✓
	cluster2	c2_node1	10.10.20.70	2/2	1	✓

8c Click the cluster name of a peer cluster.

8d Select the **BCC Manager** tab.

The connection status is good for this peer cluster. There are no BCC-enabled resources at this time.



cluster1.clusters.siteA.example

View cluster connection and cluster node status. View or change cluster resource status.

Cluster Manager | **BCC Manager** | Cluster Event Log | Cluster Options

Connections

Details

<input type="checkbox"/>	<u>Name</u>	<u>State</u>
<input type="checkbox"/>	cluster1	 Normal
<input type="checkbox"/>	cluster2	 Normal

BCC Enabled Resources

[BCC Migrate](#) | [Details](#) | [Refresh](#)▼

<input type="checkbox"/>	<u>Type</u>	<u>Name</u>	<u>State</u>	<u>Location</u>	<u>Available Peers</u>
<i>No items</i>					

- 8e Return to the My Clusters page, select another peer cluster, then go to its **BCC Manager** tab.

The connection status is good for the other peer cluster. There are no BCC-enabled resources at this time.



cluster2.clusters.siteB.example

View cluster connection and cluster node status. View or change cluster resource status.

Cluster Manager | **BCC Manager** | Cluster Event Log | Cluster Options

Connections

Details

<input type="checkbox"/>	<u>Name</u>	<u>State</u>
<input type="checkbox"/>	cluster1	 Normal
<input type="checkbox"/>	cluster2	 Normal

BCC Enabled Resources

[BCC Migrate](#) | [Details](#) | [Refresh](#)▼

<input type="checkbox"/>	<u>Type</u>	<u>Name</u>	<u>State</u>	<u>Location</u>	<u>Available Peers</u>
<i>No items</i>					

7.3 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 peer clusters. You need to add the search-and-replace strings that are used to transform cluster resource load and unload scripts from another cluster to the one where you create the replacement script. Replacement scripts are for inbound changes to scripts for objects being synchronized from other clusters, not outbound.

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

For example, consider two clusters where ClusterA uses subnet 10.10.10.x and ClusterB uses subnet 10.10.20.x. For ClusterA, you create a replacement script to replace the IP addresses of inbound resources. IP addresses starting with “10.10.20.” are replaced with IP addresses starting with “10.10.10.” so that they work in ClusterA’s subnet. For ClusterB, you create a replacement script to replace the IP addresses of inbound resources. IP addresses starting with “10.10.10.” are replaced with IP addresses starting with “10.10.20.” so that they work in ClusterB’s subnet.

The scripts are not changed for a cluster until a synchronization event comes from the other cluster. To continue the example, you can force an immediate update of the scripts for ClusterB by opening the script for ClusterA, add a blank line, then click **Apply**. To force an immediate update of the scripts for ClusterA, open the script for ClusterB, add a blank line, then click **Apply**.

You can see the IP addresses that are currently assigned to resources on a given node by entering the `ip addr show` command at the Linux terminal console on that node. It shows only the IP addresses of resources that are online when the command is issued. You must be logged in as `root` to use this command. Repeat the command on each node in the cluster to gather information about all IP addresses for resources in that cluster.

To add search-and-replace values to the cluster replacement script:

- 1 In iManager, click **Clusters > My Clusters**.
- 2 Select the check box next to the cluster, then select **Actions > Edit Properties** to view the cluster’s properties.
You can also access the cluster properties with the **Properties** button on the **Cluster Options** tab when you manage the cluster.
- 3 In the **Properties** dialog box, select the **Business Continuity** tab.
- 4 In the **Resource Replacement Script** section of the Business Continuity Cluster Properties page, click **New**.
- 5 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 specified 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 done 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 IP addresses, and replaces them with the 192.168.1.1, 192.168.1.2, 192.168.1.3, and 192.168.1.4 IP addresses, respectively.

IMPORTANT: 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”.

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

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

- 7 Click **Apply** to save your changes.

Clicking **OK** does not apply the changes to the directory.

- 8 Verify that the change has been synchronized with the peer clusters by the Identity Vault.
- 9 Continue with [Section 8.4, “Adding BCC Load and Unload Scripts,” on page 95.](#)

7.4 Enabling Linux POSIX File Systems to Run on Secondary Clusters

If you are using Linux POSIX file systems in cluster resources on the clusters in your BCC and you want to migrate or fail over those file systems to peer clusters, you must add a script to convert the container for the file system. Without the script, the file system cannot be mounted and the cluster resource cannot be brought online on another cluster.

NOTE: The script is necessary only for Linux POSIX file systems that were created on OES 2 systems and used EVMS Cluster Segment Manager containers. It is not used for clustered NSS pools or LVM volume groups.

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **Cluster Options** tab.
- 4 Under **Cluster Objects**, select the BCC-enabled cluster resource that contains the Linux file system, then click **Details**.
Cluster resources that are enabled for business continuity have the BCC label on the resource type icon.
- 5 Click the **Business Continuity** tab, then click **Storage Management**.
- 6 Under **BCC Load Scripts**, click **New** to bring up a wizard that lets you create a new script.
- 7 In the wizard, specify the following information:
 - ◆ **Name:** Convert CSM container.

- ♦ **Description:** Converts the CSM container so the specified container is available on the secondary cluster.
- ♦ **CIM Enabled:** Ensure that this option is deselected (not checked).
- ♦ **Synchronous:** Ensure that this option is selected (checked)

See [Step 6b](#) in “[Adding BCC Load and Unload Scripts](#)” on [page 95](#) for more information and descriptions of the information in the value fields.

8 Under **Script Parameters**, click **New**, then specify the following:

- ♦ **Name:** Specify the variable name as `CONTAINER_NAME`. This value is case-sensitive and should be entered as

`CONTAINER_NAME`
- ♦ **Value:** Specify the name of the container. You assigned this name to the container when you created it. This value is case-sensitive and should exactly match the container name.

9 Using a text editor, copy and paste the `bcc_csm_util.pl` script into the **Script Parameters** text box.

The script is located in the `/nsmi_scripts/linux/` directory on the Business Continuity Clustering CD or ISO image.

10 Click **Apply** to save your changes.

11 Repeat [Step 1](#) through [Step 10](#) for each peer cluster in your BCC.

The information you provide in the steps should be unique for each cluster.

7.5 Verifying BCC Administrator User Trustee Rights and Credentials

You must ensure that the BCC Administrator user is a LUM-enabled user. For instructions, see “[Creating the BCC Group and Administrator User](#)” on [page 56](#).

You must ensure that the user who manages your BCC (that is, 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. For instructions, see “[Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects](#)” on [page 57](#).

In order for the BCC Administrator user to gain access to the cluster administration files (`/admin/novell/cluster`) on other Linux cluster nodes in your BCC, you must add that user to the `ncsgroup` on each cluster node. For instructions, see “[Adding the BCC Administrator User to the ncsgroup on Each Cluster Node](#)” on [page 58](#).

7.6 Disabling BCC for a Peer Cluster

Before you disable BCC for a given peer cluster, you must first disable BCC for each of the cluster resources running on that cluster. Ensure that you remove the secondary peer clusters from the cluster resource’s **Assigned** list before you disable BCC for the resource on the primary peer cluster. For information, see [Section 8.7, “Disabling BCC for a Cluster Resource,”](#) on [page 99](#).

After you have disabled BCC for all resources running on that cluster, remove the secondary peer clusters from the Assigned list of preferred nodes for that cluster, then disable BCC for the cluster.

7.7 What's Next

Enable the cluster resources in each peer cluster that you want to be able to fail over between them. For information, see [Chapter 7, “Configuring BCC for Peer Clusters,”](#) on page 81.

8

Configuring BCC for Cluster Resources

After you have set up the Novell Cluster Services clusters for a business continuity cluster by using Business Continuity Clustering software, you are ready to configure the cluster resources for BCC. You can enable BCC for one or multiple cluster resources in each peer cluster. For each resource, you can specify the preferred peer clusters for failover.

- ◆ [Section 8.1, “Requirements for Cluster Resources,” on page 91](#)
- ◆ [Section 8.2, “BCC-Enabling Cluster Resources,” on page 93](#)
- ◆ [Section 8.3, “Configuring Search-and-Replace Values for an Individual Cluster Resource,” on page 94](#)
- ◆ [Section 8.4, “Adding BCC Load and Unload Scripts,” on page 95](#)
- ◆ [Section 8.5, “Assigning Preferred Peer Clusters for the Resource,” on page 97](#)
- ◆ [Section 8.6, “Assigning Preferred Nodes in Peer Clusters,” on page 98](#)
- ◆ [Section 8.7, “Disabling BCC for a Cluster Resource,” on page 99](#)
- ◆ [Section 8.8, “Deleting or Unsharing a BCC-Enabled Shared NSS Pool Resource,” on page 100](#)
- ◆ [Section 8.9, “Permanently Relocating a Cluster Resource to a Peer Cluster,” on page 102](#)
- ◆ [Section 8.10, “Renaming a BCC-Enabled Pool or Volume,” on page 103](#)

8.1 Requirements for Cluster Resources

- ◆ [Section 8.1.1, “LUNs for Cluster Pool Resources,” on page 91](#)
- ◆ [Section 8.1.2, “Volumes for Cluster Pool Resources,” on page 91](#)
- ◆ [Section 8.1.3, “Shared Disk Cluster Resources,” on page 92](#)
- ◆ [Section 8.1.4, “Cluster Resources for OES 2018 SP1 Services,” on page 92](#)

8.1.1 LUNs for Cluster Pool Resources

In a business continuity cluster, you should have only one NSS pool for each LUN that can be failed over to another cluster. This is necessary because in a business continuity cluster, entire LUNs fail over to other peer clusters. A pool is the entity that fails over to other nodes in a given cluster.

Multiple LUNs can be used as segments in a pool if the storage systems used in the clusters can fail over groups of LUNs, sometimes called consistency groups. In this case, a given LUN can contribute space to only one pool.

8.1.2 Volumes for Cluster Pool Resources

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. The clusters in the other eDirectory tree cannot decrypt the NSS volumes.

8.1.3 Shared Disk Cluster Resources

See [Table 8-1](#) for information about how to create shared disk cluster resources for OES 2018 SP1 clusters.

Table 8-1 Clustering Shared Disk Cluster Resources with Novell Cluster Services

Shared Disk Cluster Resources	Procedure
Dynamic Storage Technology shadow volume pairs	“Configuring DST Shadow Volume Pairs with Novell Cluster Services” in the <i>OES 2018 SP1: Dynamic Storage Technology Administration Guide</i>
Linux LVM volume groups and logical volumes	“Configuring and Managing Cluster Resources for Shared LVM Volume Groups” in the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i>
NCP (NetWare Core Protocol) volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 2018: SP1 NCP Server for Linux Administration Guide</i>
Novell Storage Services (NSS) pools and volumes	“Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes” in the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i>
Legacy CSM cluster resources from OES 2 SP3 clusters (managing only)	“Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers” in the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i>
Xen virtual machines	“Virtual Machines as Cluster Resources” in the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i>

8.1.4 Cluster Resources for OES 2018 SP1 Services

See [Table 8-2](#) for information about how to create cluster resources for OES 2018 SP1 services.

Table 8-2 Clustering OES 2018 SP1 Services with Novell Cluster Services

OES 2018 SP1 Service	Procedure
Apache Web Server	“Apache HTTP Server” in the <i>OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide</i>
Certificate Server	<p>The eDirectory Certificate Server is not cluster-enabled. The Certificate Server service issues Server Certificate objects that might need to reside on each node in a cluster, depending on the service that is clustered.</p> “eDirectory Server Certificates” in the <i>OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide</i>

OES 2018 SP1 Service	Procedure
DFS VLDB (Distributed File Services volume location database)	“Clustering Distributed File Services” in the <i>OES 2018 SP1: Distributed File Services Administration Guide for Linux</i>
DHCP Server	See the following sections in the <i>OES 2018 SP1: DNS/DHCP Services for Linux Administration Guide</i> : <ul style="list-style-type: none"> ◆ “Configuring DHCP with Novell Cluster Services for the NSS File System” ◆ “Configuring DHCP with Novell Cluster Services for the Linux File System”
DNS Server	“Configuring DNS with Novell Cluster Services” in the <i>OES 2018 SP1: DNS/DHCP Services for Linux Administration Guide</i>
eDirectory	eDirectory is not clustered because it has its own replica system.
File, AFP (Apple Filing Protocol)	“Configuring AFP with Novell Cluster Services for an NSS File System” in the <i>OES 2018 SP1: Novell AFP for Linux Administration Guide</i>
File, CIFS (Windows File Services)	“Configuring CIFS with Cluster Services for an NSS File System” in the <i>OES 2018 SP1: Novell CIFS for Linux Administration Guide</i>
File, FTP	“Cluster Enabling Pure-FTPD in an OES Environment” in the <i>OES 2018 SP1: Planning and Implementation Guide</i>
File, NetStorage	“Configuring NetStorage with Novell Cluster Services” in the <i>OES 2018 SP1: NetStorage Administration Guide for Linux</i>
iPrint	“Configuring iPrint with Novell Cluster Services” in the <i>OES 2018 SP1: iPrint Administration Guide</i>
MySQL	A MySQL template is available that uses a shared LVM volume group that you have already created. For information, see “Configuring MySQL with Novell Cluster Services” in the <i>OES 2015 SP1: Web Services and Applications Guide</i> .

8.2 BCC-Enabling Cluster Resources

Cluster resources must be enabled for business continuity on the primary cluster before they can be synchronized and appear as resources in the peer 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.

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **Cluster Options** tab.
- 4 Select the desired cluster resource from the list of Cluster objects, then click the **Details** link to view the resource Properties page.

- 5 Select the **Business Continuity** tab.
- 6 Enable a cluster resource or cluster pool resource for business continuity by selecting the **Enable Business Continuity Features** check box, then click **Apply**.
- 7 Select the **Scripts** tab, then view the load script and ensure that the volume ID used by the resource is unique across all nodes of every peer cluster.

If the volume ID is not unique, modify the volume ID value to assign a unique ID, then click **Apply**. The change is not applied until you take the resource offline and then bring it online.

You can use the `ncpcon volumes` command on each node in every peer cluster to identify the volume IDs in use by all mounted volumes. Compare the results for each server to identify the duplicates.

- 8 If you have not created a cluster replacement script to specify search-and-replace values for the entire cluster, do one of the following:
 - ♦ **Cluster Replacement Script:** (Recommended) Create a replacement script that contains search-and-replace values to use for the entire cluster.
 1. Follow the procedure in [Section 7.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 87.
 2. Continue with [Section 8.5, “Assigning Preferred Peer Clusters for the Resource,”](#) on page 97.
 - ♦ **Individual Resource Replacement Script:** (Not recommended) Create a replacement script that contains search-and-replace values only for an individual resource. Continue with [Section 8.3, “Configuring Search-and-Replace Values for an Individual Cluster Resource,”](#) on page 94.

8.3 Configuring Search-and-Replace Values for an Individual Cluster Resource

Some values for a cluster resource’s properties (such as IP addresses) are specified in the resource load and unload scripts when you create the resource. Whenever a resource fails over to a different peer cluster, these properties must be changed to values that work in the destination peer cluster. BCC provides a Search-and-Replace option that dynamically modifies the load and unload scripts for the resource so that it has the values it needs when it is migrated to or fails over between peer clusters.

You can create a replacement script that contains search-and-replace values for the entire cluster, or for an individual resource. Replacement scripts are for inbound changes to scripts for objects being synchronized from other clusters, not outbound.

IMPORTANT: We recommend that you create a replacement script for the entire cluster instead of for an individual resource. For information, see [Section 7.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 87.

Before you create a replacement script for an individual resource, you should contact [Micro Focus Support \(http://support.novell.com\)](http://support.novell.com).

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

To create a replacement for an individual cluster resource:

- 1 In the **Resource Script Replacements** section of the Business Continuity 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 you can 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 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.

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

See [“Adding Search-and-Replace Values to the Resource Replacement Script” on page 87](#) for more information on resource script search-and-replace values and adding those values to the entire cluster.

- 3 If you are creating a new cluster resource, click **Next**.
- 4 Continue with [Section 8.5, “Assigning Preferred Peer Clusters for the Resource,” on page 97](#).

8.4 Adding BCC Load and Unload Scripts

You can create BCC load and unload scripts for each BCC-enabled resource in each peer cluster. They are used to handle the transfer of SAN devices between clusters during a failover.

Scripts written for a SAN that mirrors data between two clusters should demote/mask a LUN (or group of LUNs) for a running resource on its current cluster, swap the synchronization direction, then promote/unmask the LUN(s) for the resource on the other cluster.

You can add commands that are specific to your storage 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.

The scripts or commands you add are stored in eDirectory. If you add commands to call outside scripts, those scripts must exist in the file system in the same location on every server in the cluster.

IMPORTANT: Scripts are not synchronized by Identity Manager.

Consider the following guidelines when creating and using scripts:

- ♦ Scripts must be written in Perl or have a Perl wrapper around them.
- ♦ Log files can be written to any location, but the BCC cluster resource information is logged to SYSLOG (`/var/log/messages`).
- ♦ Error codes can be used and written to a control file so that you know why your script failed. BCC checks only whether the script was successful. If an error is returned from the script, the resource does not load and remains in the offline state.
- ♦ The BCC scripts are run from the Master_IP_Address_Resource node in the cluster.
- ♦ Perl script code that you customize for your SAN can be added to a BCC-enabled cluster resource load script and unload script through the BCC management interface.
 - ♦ You can include parameters that are passed to each Perl script. BCC passes the parameters in the format of `%parm1%`, `%parm2%`, and so on.

- ◆ There can be multiple scripts per resource, but you need to use a common file to pass information from one script to another.
- ◆ The BCC load script and unload script for a BCC-enabled cluster resource must be unique on each cluster node.

To add storage management configuration information:

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **Cluster Options** tab.
- 4 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.

- 5 Click the **Business Continuity** tab, then click **BCC Scripts**.
- 6 Create BCC storage management load and unload scripts:

- 6a 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 in which load scripts execute by clicking the **Move Up** and **Move Down** links.

- 6b Specify values for the following parameters on the Storage Management Script Details page:

Parameter	Description
Name	Specify a name for the script you are creating.
Description	If desired, specify a description of the script you are creating.
CIM enabled	Select this check box if your SAN storage system supports SMI-S. This causes the CIM-specific fields to become active on this page. It is selected by default.
CIMOM IP or DNS	If you selected the CIM Enabled check box, specify the IP address or DNS name that is used for SAN management.
Namespace	If you selected the CIM Enabled check box, accept the default namespace, or specify a different namespace for your storage system. Namespace determines which models and classes are used with your storage system. Consult the vendor documentation to determine which namespace is required for your storage system.
User name and password	If you selected the CIM Enabled check box, specify the user name and password that is used to connect to and manage your SAN.
Port	If you selected the CIM Enabled check box, accept the default port number or specify a different port number. This is the port number that CIMOM (your storage system manager) uses. Consult your storage system documentation to determine which port number you should use.
Secure	If you selected the CIM Enabled check box, select the Secure check box if you want storage management communication to be secure (HTTPS). Deselect the Secure check box to use non-secure communications (HTTP) for storage management communications.

Parameter	Description
Synchronous	<p>Select this check box to run scripts sequentially (that is, one at a time). This is the default.</p> <p>Deselect this check box to allow multiple scripts to run concurrently. Most SAN storage system vendors do not support running multiple scripts concurrently.</p>
Script parameters	<p>If desired, specify variables and values for the variables that are used in the storage management script.</p> <p>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.</p> <p>These script commands are specific to your SAN hardware. You can add Perl script commands that can be run on Linux. If you add commands to call outside scripts, those scripts must exist on every server in the cluster.</p>
Script parameters text box	<p>Use this text box to add script commands to the script you are creating.</p> <p>These script commands are specific to your SAN hardware. You can add Perl script commands that can be run on Linux.</p> <p>IMPORTANT: If you add commands to call outside scripts, those scripts must exist with the same name and path on every server in the cluster.</p>

- 6c** 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.

8.5 Assigning Preferred Peer Clusters for the Resource

The cluster resource is configured on its primary cluster. Preferred peer clusters are the other clusters that a cluster resource can be migrated to in a business continuity cluster. You specify to the resource's **Assigned** clusters list where the resource can be migrated. You can migrate a resource only to one of the clusters that has been selected.

When you add a new peer cluster to the business continuity cluster (or remove a cluster and add it back in), the new cluster is automatically listed in the **Not Assigned** list for each of the existing cluster resources. You must manually assign the peer cluster to the **Assigned** list of preferred clusters for each of the cluster resources that you want to be able to fail over to it.

- 1 In the **Assigned** list, specify the preferred peer clusters for fail over and their preferred failover order.

Select a peer cluster, then click the left-arrow or right-arrow button to move it to the **Assigned** list or **Unassigned** list. Select a peer cluster, then use the up-arrow and down-arrow buttons to change the preferred failover order of the peer clusters in the **Assigned** list.

IMPORTANT: By default, the cluster resource can fail over to any node in the preferred peer cluster. That is, all nodes in a peer cluster appear on the **Assigned Nodes** list for that resource on that cluster.

- 2 Do one of the following:

Cluster Resource	New or Existing	Action
Linux POSIX file system	New	Click Finish .
Linux POSIX file system	Existing	Click Apply .
NSS pool	New	Click Next , then add the storage management configuration information. For information, see “Adding BCC Load and Unload Scripts” on page 95 .
NSS pool	Existing	Add the storage management configuration information. For information, see “Adding BCC Load and Unload Scripts” on page 95 .

- 3 After the cluster resource information is synchronized to all of the peer clusters in the resource’s **Assigned** clusters list, you must specify your node preferences separately for each of the peer clusters. For information, see [Section 8.6, “Assigning Preferred Nodes in Peer Clusters,” on page 98](#).

8.6 Assigning Preferred Nodes in Peer Clusters

After the resource information has been synchronized to the peer clusters, its preferred nodes settings in the **Assigned Nodes** list in each peer are set to all nodes in the cluster by default. You can change the resource’s preferred nodes settings on peer clusters by specifying the preferred nodes within each of the peer clusters on the resource **Assigned Nodes** list.

Before you begin, ensure that you have assigned the preferred peer clusters for the resource. For information, see [Section 8.5, “Assigning Preferred Peer Clusters for the Resource,” on page 97](#).

For each BCC-enabled cluster resource, do the following for each peer cluster where you plan to fail over the cluster resource:

- 1 In iManager, click **Clusters > My Clusters**, then click the cluster name.
- 2 Click the **Cluster Options** tab.
- 3 Select the box next to the resource whose preferred node list you want to view or edit, then click the **Details** link.
If this is a peer cluster where the resource is currently running, the cluster resource has a status of primary. On other peer clusters, the status is secondary.
- 4 Click the **Preferred Nodes** tab.
- 5 From the **Unassigned Nodes** list, select the server you want the resource assigned to, then click the right-arrow button to move the selected server to the **Assigned Nodes** list.
Repeat this step for all servers you want assigned to the resource.
- 6 From the **Assigned Nodes** list, select the servers 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.

- 7 Click the up-arrow and down-arrow buttons to change the preferred failover order of the servers assigned to the resource or volume.
- 8 Click **Apply** to save the node assignment changes.
The preferences apply immediately.

8.7 Disabling BCC for a Cluster Resource

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 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.

Before you disable BCC for a BCC-enabled resource, 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 attempt to disable BCC on the same resource from multiple peer clusters.

To disable BCC for a cluster resource:

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **Cluster Options** tab.
- 4 Select the check box next to the desired cluster resource from the list of Cluster objects.
- 5 Click the **Details** link.
- 6 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.
 - 6a Click the **Preferred Nodes** tab.
 - 6b From the **Assigned Nodes** list, select the nodes that reside in the secondary peer clusters, then click the 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.
 - 6c Click **Apply** to save the node assignment changes.
- 7 On the **Details** page, click the **Business Continuity** tab, deselect the **Enable Business Continuity Features** check box, then click **Apply**.

8 Wait for Identity Manager to synchronize the changes.

This could take from 30 seconds to one minute, depending on your configuration.

9 Delete the Cluster Resource object on the clusters where you no longer want the resource to run.

In iManager, use the Clusters plug-in to delete the cluster resource.

9a In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.

9b If the cluster resource is online, take it offline:

On the Cluster Manager page, select the check box next to the cluster resource, then click **Offline**.

Refresh the Cluster Manager page to verify that the resource has a status of **Offline** before you continue.

9c Select the **Cluster Options** tab.

9d Select the check box next to the resource, then click **Delete**.

This deletes the Cluster Resource object. It does not delete the storage or service represented by the object.

For a pool cluster resource, this also deletes the related NCP virtual server object, Pool object, and Volume object.

NOTE: Repeat [Step 9 on page 100](#) on all the clusters where you no longer want the resource to run.

8.8 Deleting or Unsharing a BCC-Enabled Shared NSS Pool Resource

You might need to unshare a pool resource if you decide not to cluster the pool. You might decide to delete an NSS pool if the data stored on it is no longer needed.

Consider the following guidelines when planning to delete or unshare a BCC-enabled pool resource:

- ♦ All resource configuration must happen from the master node. On the Cluster Options page for iManager, connect to the Cluster object, not to Cluster Node objects. On the **Storage > Pools** page for iManager, connect to the master node. Run NSSMU only on the master node.
- ♦ You must disable BCC for a shared NSS pool before you delete it or unshare it. This allows the pool to be removed from the list of resources kept by the peer clusters.
- ♦ You must offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering; then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

To delete or unshare a BCC-enabled shared NSS pool resource:

- 1 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 2 Click the **Cluster Options** tab.
- 3 If the resource is on a non-master node in the cluster, migrate it to the master node. Select the resource, click **Migrate**, then specify the master node.

- 4 Select the shared NSS pool resource, click **Properties**, then select **Business Continuity**.
- 5 Disable BCC for the pool cluster resource, then click **Apply**.

For information, see [Section 8.7, “Disabling BCC for a Cluster Resource,” on page 99](#).

Wait for Identity Manager to synchronize the change to all peer clusters before continuing. You can check the other peer clusters to ensure that the cluster resource you BCC-disabled no longer appears in their list of resources.

- 6 Go to the **Cluster Manager** page, select the check box next to the NSS pool cluster resource, then click **Offline**.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

- 7 Do one of the following:

Task	Procedure
Keep the shared NSS pool to use it only in the current cluster.	<ol style="list-style-type: none"> 1. Click the Cluster Options tab. 2. Select NSS pool resource to view Properties, then select Scripts. 3. Modify the load script by removing any commands related to the BCC, then click Apply. 4. Modify the unload script by removing any commands related to the BCC, then click Apply. 5. Modify the monitor script by removing any commands related to the BCC, then click Apply. 6. Click OK. 7. Go to the Cluster Manager page, select the check box next to the NSS pool cluster resource, then click Online.
Keep the NSS pool, but unshare it to use it as a non-clustered pool.	<ol style="list-style-type: none"> 1. In iManager, select Storage > Pools, then select the master node of the cluster. 2. Select the pool to view its Pool Details. (It should be deactive since the resource was offline when you modified the script.) 3. Mark the pool as Not Shareable for Clustering, then click Apply. 4. In Roles and Tasks, click Clusters > My Clusters, then click the cluster name. 5. Click the Cluster Options tab. 6. Select the check box next the Cluster Resource object you want to delete, then click Delete.

Task	Procedure
Unshare the pool, then delete it.	<ol style="list-style-type: none"> 1. In iManager, select Storage > Pools, then select the master node of the cluster. 2. Select the pool to view its Pool Details. (It should be deactive since the resource was offline when you modified the script.) 3. Mark the pool as Not Shareable for Clustering, then click Apply. 4. Select Clusters > My Clusters, select the cluster, then select the Cluster Options tab. 5. Select the check box next the Cluster Resource object you want to delete, then click Delete. 6. In iManager, select Storage > Pools, then select the node of the cluster. 7. Select the pool to view the Pool Details. 8. Click Delete, then confirm the deletion.

8.9 Permanently Relocating a Cluster Resource to a Peer Cluster

In some disaster situations, it is possible that the pool cluster resource has been permanently relocated to the current peer cluster. You must create new Pool and Volume objects in its new permanent location in order in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.

Because these changes need to send information to eDirectory, you should perform them from the master node of its new cluster location.

To cluster-enable the pool in its current cluster (that is, in a different peer cluster than where it was originally created):

- 1 Save any unique or custom information for this resource, such as custom load, unload, or monitoring scripts, BCC scripts for SAN management, and resource replacement values.
- 2 Disable BCC on the pool cluster resource.
- 3 Take the resource offline on the peer cluster where you want it to permanently reside.
- 4 Delete the Cluster Resource objects from eDirectory for the original cluster and the peer cluster.

You should be able to use the Clusters plug-in in iManager to delete the cluster resource. This should clean up the related resource objects, particularly the Virtual NCP Server object for the resource. You can alternatively delete the objects manually from eDirectory.

Check eDirectory manually to ensure that the Virtual NCP Server object, the Pool object, and the Volume object are also deleted for the resource you deleted. These objects reside in the landing zone of the peer cluster and in the normal context on the original cluster.

- 5 Use the **Update eDirectory** option in NSSMU or in the Storage plug-in in iManager to create new Storage objects for the existing pool and its volumes.

The new names will be based on the server where you create the objects.

- 6 (Optional) After the new objects have been synchronized in eDirectory, you can perform tasks on the pool and volume, such as renaming the pool or volume, or adding new volumes.

- 7 Use the Clusters plug-in to cluster-enable the pool and its volumes in its current cluster.
- 8 BCC enable the resource, then reconfigure any of the custom settings for the resource based on its current cluster.

8.10 Renaming a BCC-Enabled Pool or Volume

You can rename a BCC-enabled pool or volume in its original cluster, that is, where it was created rather than from a peer cluster.

IMPORTANT: If you cannot BCC-migrate the resource back to its original cluster, such as in a disaster situation, you must permanently relocate the pool to its current cluster as described in [Section 8.9, “Permanently Relocating a Cluster Resource to a Peer Cluster,” on page 102](#). You can rename the pool or volume as part of this re-assignment process.

- 1 Migrate the resource to the master node of its original cluster.
- 2 Save any unique or custom information for this resource, such as custom load, unload, or monitoring scripts, BCC scripts for SAN management, and resource replacement values.
- 3 Disable BCC for the resource as described in [Section 8.7, “Disabling BCC for a Cluster Resource,” on page 99](#), following the process to keep it cluster-enabled in its original pool.
- 4 (Optional) Rename the clustered pool as described in “Renaming a Clustered NSS Pool ” in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.
- 5 (Optional) Rename a clustered volume as described in “Renaming a Clustered NSS Volume ” in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.
- 6 Re-enable BCC for the resource as described in [Section 8.2, “BCC-Enabling Cluster Resources,” on page 93](#).
- 7 Reconfigure any of the custom settings for the resource based on its current cluster.

9 Managing a Business Continuity Cluster

This section can help you effectively manage a business continuity cluster with the 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.

IMPORTANT: Identity Manager must be configured and running on one node in each peer cluster before any BCC-enabled cluster resource changes are made. Ensure that the Identity Manager server is part of the cluster and that it is working properly whenever you make BCC configuration changes to the BCC-enabled cluster resources. For information, see [Chapter 6, “Configuring the Identity Manager Drivers for BCC,” on page 67.](#)

For information about using console commands to manage your business continuity cluster, see [Appendix A, “Console Commands for BCC,” on page 137.](#)

- ♦ [Section 9.1, “Migrating a Cluster Resource to a Peer Cluster,” on page 105](#)
- ♦ [Section 9.2, “Bringing a Downed Cluster Back in Service,” on page 107](#)
- ♦ [Section 9.3, “Changing Peer Cluster Credentials,” on page 107](#)
- ♦ [Section 9.4, “Viewing the Current Status of a Business Continuity Cluster,” on page 109](#)
- ♦ [Section 9.5, “Generating a Cluster Report,” on page 109](#)
- ♦ [Section 9.6, “Resolving Business Continuity Cluster Failures,” on page 110](#)

9.1 Migrating a Cluster Resource to a Peer Cluster

Although 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, see [Appendix C, “Setting Up Auto-Failover,” on page 145.](#)

- ♦ [Section 9.1.1, “Understanding BCC Resource Migration,” on page 105](#)
- ♦ [Section 9.1.2, “Migrating Cluster Resources between Clusters,” on page 106](#)

9.1.1 Understanding BCC Resource Migration

A cluster resource can be migrated or failed over to nodes in the same cluster or to nodes in a peer cluster. Typically, you migrate or fail over locally to another node in the same cluster whenever it makes sense to do so. If one site fails (all nodes in a given cluster are not functional), you can use iManager to manually BCC migrate resources to any of the peer clusters. Each resource starts on its preferred node on the peer cluster where you have BCC migrated the resources.

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 storage management unload script that is associated with the pool resource is run.

4. The `cluster scan for new devices` command is executed on the peer cluster so that the cluster is aware of LUNs that are no longer available.
5. On the destination peer cluster, the resource changes from secondary to primary so that it can be brought online.
6. Any storage management load script that is associated with the pool resource is run.
If an error is returned from the BCC load script, the resource is not brought online and remains in the offline, not comatose, state.
7. The `cluster scan for new devices` command is executed on the destination peer cluster so that the cluster is aware of LUNs that are now available.
8. Resources are brought online and load on the most preferred node in the cluster (that is, on the first node in the preferred node list).

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

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

9.1.2 Migrating Cluster Resources between Clusters

WARNING: Do not migrate resources for a test failover if the storage connection between the source and destination cluster is down. Possible disk problems and data corruption can occur if the down connection comes up and causes a divergence in data. 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 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **BCC Manager** tab.
- 4 Select one or more cluster resources, then click **BCC Migrate**.

The cluster you chose is shown as the **Current Cluster**.

Current Cluster is associated with the first table on the BCC Migrate page. It tells you what cluster you selected to manage the BCC. Information is provided from the point of view of that cluster.

For example, if the resource is assigned to a node in the cluster you are managing from, the cluster resource status is the same as if you were looking at the cluster itself, such as **Running** or **Offline**. If the cluster resource is not assigned to the cluster you are managing from (that is, not in the current cluster), then the status is shown as **Secondary**.

- 5 In the list of clusters, 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.

- 6 View the state of the migrated resources by selecting **Clusters > BCC Manager**, then select the Cluster object of the cluster where you have migrated the resources.

The migration or failover of NSS pools or other resources between peer clusters can take additional time compared to migrating between nodes in the same cluster. Until the resource achieves a **Running** state, iManager shows multiple states as the resource is unloaded from the node in the source cluster and loaded successfully on its preferred node in the destination cluster.

9.2 Bringing a Downed Cluster Back in Service

If a cluster has been totally downed (all nodes are down concurrently), the peer clusters do not automatically recognize the cluster if you bring the nodes back online.

To bring the downed cluster back into service in the business continuity cluster:

- 1 Bring up only a single node in the cluster.
- 2 At the terminal console of this node, enter

```
cluster resetresources
```

- 3 Bring up the remainder of the nodes.

9.3 Changing Peer Cluster Credentials

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

Configure the new peer cluster credentials on one node in each peer cluster in the business continuity cluster:

- 1 Log in as the `root` user on the cluster node where you want to add peer credentials in this cluster, then open a terminal console.

You can use any node in the cluster to specify credentials.

- 2 At the command prompt, enter

```
cluster connections <cluster_name>
```

For example, this lists the BCC peer clusters and credential information. At this point, the credentials have been changed but not entered for the peer clusters. They are reported as invalid.

```
cluster connections cluster1
```

```
Connection status for cluster: cluster1
```

Cluster Name	Username	Connection Status
cluster1	<unknown>	Invalid Credentials
cluster2	<unknown>	Invalid Credentials

- 3 For each cluster in the list, enter the following command at the command prompt, then enter the `bccadmin` user name and password when you are prompted.

```
cluster credentials <cluster_name>
```

For example, issue the command once for each peer cluster:

```
cluster credentials cluster1
Enter the credentials for the specified cluster. Press CTRL+C to cancel.
Username: bccadmin
Password:
Please wait...
The credentials for cluster cluster1 have been saved.
```

```
cluster credentials cluster2
Enter the credentials for the specified cluster. Press CTRL+C to cancel.
Username: bccadmin
Password:
Please wait...
The credentials for cluster cluster2 have been saved.
```

If you created different BCC Administrator users for each peer cluster, you will provide the respective credentials for each cluster.

- 4 Verify that the peer clusters are recognized as being members of the same BCC. Enter the `cluster connections` command and `cluster view` command to view the status.

For example, on cluster1:

```
cluster connections cluster1

Connection status for cluster: cluster1
  Cluster Name      Username      Connection Status
  cluster1         bccadmin     OK
  cluster2         bccadmin     OK
```

```
cluster view

  Cluster cluster1
  This node c1_node1 [epoch 1 master node c1_node2]
  Cluster nodes [c1_node1, c1_node2]
  BCC peer clusters [cluster1, cluster2]
```

If you created different BCC Administrator users for each peer cluster, each peer cluster would report a different user name in the `cluster connections` command.

- 5 Add the BCC Administrator user to the `ncsgroup` for the cluster. Repeat the following steps for every node in the peer cluster:

5a As the `root` user, open the `/etc/group` file in a text editor.

5b Locate the line that reads `ncsgroup`, then modify it to include the `bccadmin` user.

For example, change

```
ncsgroup:!:107:
```

to

```
ncsgroup:!:107:bccadmin
```

For example, change

```
ncsgroup:!:107:bccd
```

to

```
ncsgroup:!:107:bccd,bccadmin
```

The file should contain one of the above lines, but not both.

Notice the group ID number of the `ncsgroup`. In this example, the number 107 is used. This number can be different for each cluster node.

5c Save the `/etc/group` file.

5d At the server console prompt, enter the following to verify that the `bccadmin` user is a member of the `ncsgroup`.

```
id bccadmin
```

6 Repeat the previous steps on a node in each peer cluster in turn.

The nodes in the other peer clusters know about the peer clusters, but they do not yet have the credentials needed to connect.

9.4 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 9.4.1, “Using iManager to View the Cluster Status,” on page 109](#)
- ♦ [Section 9.4.2, “Using Console Commands to View the Cluster Status,” on page 109](#)

9.4.1 Using iManager to View the Cluster Status

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**, then click the cluster name.
- 3 Click the **BCC Manager** tab.
- 4 View the status of the BCC resources in the business continuity cluster.
- 5 Repeat the process for each peer cluster to view the status from each cluster’s point of view.

9.4.2 Using Console Commands to View the Cluster Status

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

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

9.5 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 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**.
- 3 Select the check box next to the cluster.
- 4 Select **Actions > Run Report**.

9.6 Resolving Business Continuity Cluster Failures

There are several failure types associated with a business continuity cluster. 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 storage-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. However, your storage vendor documentation might use different terms, such as *mask* and *unmask*.

- ♦ [Section 9.6.1, “Storage-Based Mirroring Failure Types and Responses,” on page 110](#)
- ♦ [Section 9.6.2, “Host-based Mirroring Failure Types and Responses,” on page 112](#)

9.6.1 Storage-Based Mirroring Failure Types and Responses

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

- ♦ [“Primary Cluster Fails but Primary Storage System Does Not” on page 110](#)
- ♦ [“Primary Cluster and Primary Storage System Both Fail” on page 111](#)
- ♦ [“Secondary Cluster Fails but Secondary Storage System Does Not” on page 111](#)
- ♦ [“Secondary Cluster and Secondary Storage System Both Fail” on page 111](#)
- ♦ [“Primary Storage System Fails and Causes the Primary Cluster to Fail” on page 111](#)
- ♦ [“Secondary Storage System Fails and Causes the Secondary Cluster to Fail” on page 111](#)
- ♦ [“Inter-Site Storage System Connectivity Is Lost” on page 111](#)
- ♦ [“Inter-Site LAN Connectivity Is Lost” on page 111](#)

Primary Cluster Fails but Primary Storage System 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 BCC migrate the resources to a peer cluster. Next, work to restore the failed 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 peer cluster.

Prior to bringing up the original cluster servers, ensure that the storage system and SAN interconnect are in a state in which the cluster resources cannot come online and cause a divergence in data. Divergence in data occurs when connectivity between storage systems has been lost and both clusters assert that they have ownership of their respective disks. Ensure that the former primary storage system is demoted to secondary before bringing cluster servers back up. If the former primary storage system has not been demoted to secondary, you might need to demote it manually. Consult your storage hardware documentation for instructions on demoting and promoting LUNs. You can use the `cluster resetresources` console command to change resource states to offline and secondary.

Primary Cluster and Primary Storage System Both Fail

Bring the primary storage system back up. Follow your storage vendor's instructions to remirror it. Promote the former primary storage system back to primary. Then bring up the former primary cluster servers, and fail back the cluster resources.

Secondary Cluster Fails but Secondary Storage System Does Not

Secondary clusters are not currently running the resource. 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 Storage System Both Fail

Secondary clusters are not currently running the resource. Bring the secondary storage system back up. Follow your storage 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 Storage System Fails and Causes the Primary Cluster to Fail

When the primary storage system fails, the primary cluster also fails. BCC migrate the resources to a peer cluster. Bring the primary storage system back up. Follow your storage vendor's instructions to remirror. Promote the former primary storage system back to primary. You might need to demote the LUNs and resources to secondary on the primary storage 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 the resources.

Secondary Storage System Fails and Causes the Secondary Cluster to Fail

Secondary clusters are not currently running the resource. When the secondary storage system fails, the secondary cluster also fails. Bring the secondary storage back up. Follow your storage vendor's instructions to remirror. Then bring the secondary cluster back up. When you bring the secondary storage system and cluster back up, resources are still in a secondary state.

Inter-Site Storage System Connectivity Is Lost

Recover the connection. If divergence of the storage systems occurred, remirror from the good side to the bad side.

Inter-Site LAN Connectivity Is Lost

User connectivity might be lost to a given service or data, depending on where the resources are running and whether multiple clusters run the same service. Users might not be able to access servers in the cluster that they usually connect to, but can possibly access servers in another peer cluster. If users are co-located with the cluster that runs the service or stores the data, nothing additional is required. An error is displayed. Wait for connectivity to resume.

If you have configured the auto-failover feature, see [Appendix C, "Setting Up Auto-Failover," on page 145](#).

9.6.2 Host-based Mirroring Failure Types and Responses

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

- ◆ “Primary Cluster Fails but Primary Storage System Does Not” on page 112
- ◆ “Primary Cluster and Primary Storage System Both Fail” on page 112
- ◆ “Secondary Cluster Fails but Secondary Storage System Does Not” on page 112
- ◆ “Secondary Cluster and Secondary Storage System Both Fail” on page 112
- ◆ “Primary Storage System Fails and Causes the Primary Cluster to Fail” on page 112
- ◆ “Secondary Storage System Fails and Causes the Secondary Cluster to Fail” on page 113
- ◆ “Inter-Site Storage System Connectivity Is Lost” on page 113
- ◆ “Inter-Site LAN Connectivity is Lost” on page 113

Primary Cluster Fails but Primary Storage System Does Not

The initial response is to BCC migrate the resources to a peer cluster. Next, work to restore the failed cluster to normal operations. The long-term response is total recovery from the failure. Do not disable MSAP (Multiple Server Activation Prevention), which is enabled by default. When the former primary cluster is recovered, bring up the former primary cluster servers, and fail back the cluster resources.

Primary Cluster and Primary Storage System Both Fail

Bring up your primary storage system 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 the downed cluster servers back up. Then bring up the former primary cluster servers, and fail back the cluster resources.

Secondary Cluster Fails but Secondary Storage System Does Not

Secondary clusters are not currently running the resource. No additional response is necessary for this failure other than recovering the secondary cluster. When you bring the secondary cluster back up, the storage system is still secondary to the primary cluster.

Secondary Cluster and Secondary Storage System Both Fail

Secondary clusters are not currently running the resource. Bring up your secondary storage system 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 storage system is still secondary to the primary cluster.

Primary Storage System Fails and Causes the Primary Cluster to Fail

If your primary storage system fails, all nodes in your primary cluster also fail. BCC migrate the resources to a peer cluster. Bring the primary storage system back up. Bring up your primary cluster servers. Ensure that remirroring completes before failing back resources to the former primary cluster.

Secondary Storage System Fails and Causes the Secondary Cluster to Fail

Secondary clusters are not currently running the resource. When the secondary storage system fails, the secondary cluster also fails. Bring the secondary storage back up. Bring up your secondary cluster servers. Ensure that remirroring completes on the secondary storage system. When you bring the secondary storage system and cluster back up, resources are still in a secondary state.

Inter-Site Storage System Connectivity Is Lost

Recover the connection. If divergence of the storage systems occurred, remirror from the good side to the bad side.

Inter-Site LAN Connectivity is Lost

User connectivity might be lost to a given service or data, depending on where the resources are running and whether multiple clusters run the same service. Users might not be able to access servers in the cluster that they usually connect to, but can possibly access servers in another peer cluster. If users are co-located with the cluster that runs the service or stores the data, nothing additional is required. An error is displayed. Wait for connectivity to resume.

If you have configured the auto-failover feature, see [Appendix C, "Setting Up Auto-Failover," on page 145](#).

10 Troubleshooting Business Continuity Clustering

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

- ◆ [Section 10.1, “Identity Manager,” on page 115](#)
- ◆ [Section 10.2, “Peer Clusters,” on page 117](#)
- ◆ [Section 10.3, “BCC-Enabled Cluster Resources,” on page 121](#)
- ◆ [Section 10.4, “BCC Startup Flags,” on page 125](#)
- ◆ [Section 10.5, “BCC Error Codes,” on page 126](#)

10.1 Identity Manager

- ◆ [Section 10.1.1, “Identity Manager Plug-Ins Do Not Appear in iManager,” on page 115](#)
- ◆ [Section 10.1.2, “Identity Manager Drivers for Cluster Synchronization Do Not Start,” on page 115](#)
- ◆ [Section 10.1.3, “Identity Manager Drivers Do Not Synchronize Objects from One Cluster to Another,” on page 116](#)
- ◆ [Section 10.1.4, “Tracing Identity Manager Communications,” on page 116](#)
- ◆ [Section 10.1.5, “SSL Certificates Are Missing,” on page 117](#)

10.1.1 Identity Manager Plug-Ins Do Not Appear in iManager

After you properly install Identity Manager plug-ins in iManager, the plug-ins sometimes disappear from iManager for the tree you want to manage.

Identity Manager plug-ins for iManager require that eDirectory is running and working properly in the tree you are trying to manage. If the plug-in does not appear in iManager, ensure that the eDirectory daemon (`ndsd`) is running on the server that contains the eDirectory master replica for that tree.

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

```
rcndsd restart
```

10.1.2 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:

- ◆ 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 [Chapter 6, “Configuring the Identity Manager Drivers for BCC,” on page 67](#) 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 10.1.1.1:2002:2002.

- ◆ 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 10.1.4, “Tracing Identity Manager Communications,” on page 116](#).

10.1.3 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 eDirectory daemon (`ndsd`) is not running on the server that contains the eDirectory master replica in the tree. To restart `ndsd` on the master replica server, enter the following command at its command prompt as the `root` user:

```
rcndsd restart
```

- ◆ 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 iPrint Server and you try to synchronize a resource named iPrint_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 10.1.1.1:2002:2002.

See [Section 10.1.4, “Tracing Identity Manager Communications,” on page 116](#).

10.1.4 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).

For information about setting trace levels for driver sets, see [“Configuring Trace Levels” in the *NetIQ Identity Manager Driver Administration Guide*](#).

For information about using `ndstrace`, see “Using DSTrace” in the *NetIQ eDirectory Administration Guide*.

The trace messages are written to the `ndstrace.log` file located in the directory where eDirectory is installed. By default, it is `/var/nds`. You might want to delete this file before starting a trace so that the events logged in the file are specific to the actions you are tracing.

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

- 1 Modify two attributes on both DriverSet objects:
 - 1a Log in to iManager as the BCC Administrator user.
 - 1b Click the **View Objects** button at the top of the iManager page.
 - 1c In the tree view, browse to and right-click the desired DriverSet object, then select **Modify Object**.
 - 1d Click the **General** tab. In the list of valued attributes, click **DirXML-DriverTraceLevel**, then click **Edit**.
 - 1e Ensure that the Trace Level is set to 4, then click **OK**.
 - 1f Repeat [Step 1d](#) and [Step 1e](#) for the **DirXML-XSLTraceLevel** attribute, also setting the trace level to 4.
 - 1g Repeat [Step 1c](#) through [Step 1f](#) for the other driver sets you want to trace.
- 2 At the Linux terminal console, log in as the `root` user, then start the DSTrace utility by entering

```
 /opt/novell/eDirectory/bin/ndstrace
```
- 3 Configure DSTrace by entering

```
 ndstrace inline -all +dvrs +dxml +time
```
- 4 Exit the DSTrace utility by entering

```
 exit
```
- 5 Enable DSTrace by again entering `ndstrace on` at the Linux terminal console.
- 6 Run the desired actions for the information you want traced.
- 7 Disable DSTrace by entering `ndstrace off` at the Linux terminal console.

10.1.5 SSL Certificates Are Missing

If SSL certificates are not present or have not been created, Identity Manager drivers might not start or function properly.

10.2 Peer Clusters

- ◆ [Section 10.2.1, “Incorrect Status in the Cluster View Report After Enabling BCC on Peer Clusters,” on page 118](#)
- ◆ [Section 10.2.2, “Administration of Peer Clusters Is Not Functional,” on page 118](#)
- ◆ [Section 10.2.3, “Peer Cluster Communication Is Not Working,” on page 118](#)
- ◆ [Section 10.2.4, “Cluster Connection States,” on page 118](#)
- ◆ [Section 10.2.5, “Driver Port Number Conflicts,” on page 119](#)

- ◆ [Section 10.2.6, “Security Equivalent User,” on page 120](#)
- ◆ [Section 10.2.7, “Clusters Cannot Communicate,” on page 121](#)

10.2.1 Incorrect Status in the Cluster View Report After Enabling BCC on Peer Clusters

After the set up of BCC and enabling it on peer clusters, the cluster view report on one of the node in the first cluster shows only one peer cluster as configured instead of showing all the peer clusters. However, the cluster view report on all the other nodes correctly shows all the peer clusters configured. To view the correct cluster view report on the node in first cluster where the issue is seen, execute the command `cluster refresh -p`.

10.2.2 Administration of Peer Clusters Is Not Functional

This problem is normally caused by the BCC Administrator user not having file system rights to the cluster administration files. For information, see the following:

- ◆ [“eDirectory 9.1” on page 41](#)
- ◆ [“Configuring a BCC Administrator User and Group” on page 55](#)
- ◆ [“Verifying BCC Administrator User Trustee Rights and Credentials” on page 89](#)

10.2.3 Peer Cluster Communication Is 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.

- ◆ A firewall is blocking port 5988 or 5989 (CIM communications using OpenWBEM).

10.2.4 Cluster Connection States

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.

There are several different cluster connection states:

Table 10-1 BCC Connection States

BCC Connection State	Number	Description	Possible Actions
Normal	0	The connections between clusters are functioning normally.	None required.

BCC Connection State	Number	Description	Possible Actions
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 user name or password for the selected peer cluster.	Enter the correct user name 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.	<p>Ping the peer cluster to see if it is up and reachable.</p> <p>Ensure that Novell Cluster Services is running on the servers in the peer cluster, then ensure that BCC is running on the peer clusters.</p> <p>Ensure that OpenWBEM is running on the peer cluster.</p> <p>Ensure that a firewall is not preventing access on OpenWBEM ports 5988 and 5989.</p> <p>Ensure that the Admin file system is running. To do this, enter <code>systemctl status adminfs</code>.</p>
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 57.
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.

10.2.5 Driver Port Number Conflicts

If your Identity Manager driver or drivers will not 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 Log in to iManager as the BCC Administrator user.
- 2 Go to the Identity Manager page.
- 3 Click **Identity Manager Administration > Identity Manager Overview**.
- 4 Select **Search Entire Tree**, then click **Search**.
- 5 Select the driver you want to check by clicking the red **Cluster Sync** icon.
- 6 Click the 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 that 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** icon, then click **Restart driver**.
- 9 If you changed the port number in [Step 7](#), change the port numbers to be the same for the corresponding driver in the other cluster.

You can do this by repeating the process for the Identity Manager driver on the other cluster.

10.2.6 Security Equivalent User

If resources or peers do not appear in peer clusters in your BCC, it is possible that either a cluster resource 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 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 sometimes 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.

To make the Cluster Resource Synchronization Driver object the security equivalent to a User object with administrative rights:

- 1 Log in as the BCC Administrator user.
- 2 Go to the Identity Manager page.
- 3 Click **Identity Manager Administration > Identity Manager Overview**.
- 4 Choose **Search Entire Tree**, then click **Search**.
- 5 Select the driver you want to check by clicking the red **Cluster Sync** icon.
- 6 Click the icon again, then click the **Identity Manager** tab if it is not already selected.
- 7 Click **Security Equals**, then view or add a user as needed to be its security equivalent.
- 8 Repeat [Step 5](#) through [Step 7](#) for the other drivers in your BCC.

Ensure that the BCC Administrator user is a LUM-enabled user. To LUM-enable a user, see [“Managing User and Group Objects in eDirectory”](#) in the *OES 2018 SP1: Linux User Management Administration Guide*.

10.2.7 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 Log in as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Rights**, then click **Modify Trustees**.
- 3 Browse to select the Cluster object name, then click **OK**.
- 4 Click **OK** to view the trustee information for the selected object.
- 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 2](#) through [Step 7](#) for the other Cluster objects in your BCC.

10.3 BCC-Enabled Cluster Resources

- ♦ [Section 10.3.1, “Resource Cannot Be Brought Online,”](#) on page 121
- ♦ [Section 10.3.2, “Resource Does Not Migrate to a Peer Cluster,”](#) on page 122
- ♦ [Section 10.3.3, “Blank Error String iManager Error Appears While Bringing a Resource Online,”](#) on page 122
- ♦ [Section 10.3.4, “Clustered Pool Is Stuck in an eDirectory Synchronization State,”](#) on page 122
- ♦ [Section 10.3.5, “Mapping Drives in Login Scripts,”](#) on page 122
- ♦ [Section 10.3.6, “Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable,”](#) on page 123
- ♦ [Section 10.3.7, “Resource Script Search-and-Replace Functions Do Not Work,”](#) on page 124
- ♦ [Section 10.3.8, “Virtual NCP Server IP Addresses Won’t Change,”](#) on page 124
- ♦ [Section 10.3.9, “The IP Address, Virtual Server DN, or Pool Name Does Not Appear on the iManager Cluster Configuration Page,”](#) on page 125

10.3.1 Resource Cannot Be Brought Online

If you cannot bring a BCC-enabled resource online, 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. This sets 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 117](#) for an example of how to modify object attributes.

10.3.2 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 10.2.3, “Peer Cluster Communication Is Not Working,”](#) on page 118.
- ♦ Perl script errors.
Special characters must be escaped in scripts. For example, the script might have a single % character in it that needs to be an escape (%% instead of %). Also, hashes in Perl need to have escape characters.

10.3.3 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.

10.3.4 Clustered Pool Is Stuck in an eDirectory Synchronization State

Pools might get stuck in an eDirectory synchronization state if all of the volumes in that pool are not active and mounted.

To resolve this problem:

- 1 In NSSMU or iManager, activate and mount the deactive volumes.
- 2 In iManager, click **Clusters**, then select the cluster resource for the clustered pool.
- 3 Click **Offline** to dismount the pool's volumes and deactivate the pool.
Wait until the resources are successfully offline before continuing.
- 4 Click **Online** to bring the resource back online.

10.3.5 Mapping Drives in Login Scripts

Consider the following when mapping drives in login scripts in a BCC:

- ♦ Using a fully distinguished name for cluster-enabled volume (such as `map s:=BCCP_Cluster_HOMES.servers.:shared`) 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) (<http://support.novell.com/docs/Tids/Solutions/10057730.html>) for information on modifying the server cache Time To Live (TTL) value on the Client for Open Enterprise Server.

10.3.6 Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable

Consider the following when mapping drives to home directories 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 the resource fails 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) (<http://support.novell.com/docs/Tids/Solutions/10057730.html>) for information on modifying the Server Cache Timeout value on the Client for Open Enterprise Server.

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.

10.3.7 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 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 then verify that the change has synchronized with the other cluster before deleting it.

10.3.8 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, unload, and monitor script pages.
For information, see [“Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster”](#) in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.
- ♦ The virtual server has an extra IP address.
IP address changes should always be made on the Protocols page of the iManager cluster plugin, and not in load and unload scripts.
This is the only way to change the IP address on the virtual NCP server object in eDirectory.

10.3.9 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 **Business Continuity** tab.

The eDirectory pointers for the cluster resource are missing or are invalid. The following shows the required attributes and the format:

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

10.4 BCC Startup Flags

There are three optional flags that can be used by BCC during startup that can help with troubleshooting:

Table 10-2 Optional Startup Flags

Startup Flag	Description
d	Keeps the bccd from forking (keeps the process in foreground). Log messages are printed to the running terminal screen (<code>stdout</code>) along with the normal <code>syslog</code> .
v	Turns on more verbose logging.
t	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.

There are two usage options:

- ♦ To use the `v` and `t` flags, edit the `/etc/init.d/novell-bcc` file and change the `NOVELL_BCCD_ARGS=` line to `NOVELL_BCCD_ARGS=-flags`. Replace `flags` with any combination of `v` and/or `t`. This could include `vt`, `t`, or `tv`.

Do not use the `d` flag with this option.

- ◆ Stop BCC by entering `rcnovell-bccd stop` or `systemctl stop novell-bccd.service` at the server console. Then restart it by entering `/opt/novell/bcc/sbin/bccd -flags`. Replace *flags* with any combination of `v`, `t`, or `d`.

10.5 BCC Error Codes

The following table lists the different BCC error codes by number and gives a brief description for each error code.

Table 10-3 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.
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.

Error Code Number	Message
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.

The path to the log file on Linux is `/var/log/messages`. You can search for BCCD to view BCC related messages and entries in the log file.

11

Security Considerations

Security administrators can use information in this section to understand how to configure and maintain a business continuity cluster in the most secure way possible.

- ◆ [Section 11.1, “Security Features,” on page 129](#)
- ◆ [Section 11.2, “Security Configuration,” on page 130](#)
- ◆ [Section 11.3, “General Security Guidelines,” on page 134](#)
- ◆ [Section 11.4, “Security Information for Dependent Products,” on page 134](#)

11.1 Security Features

The following table contains a summary of the security features of Novell Business Continuity Clustering.

Table 11-1 Business Continuity Clustering Security Features

Feature	Yes/No	Details
Users are authenticated	Yes	Administrative users are authenticated via eDirectory. For information about configuring rights needed by BCC administrators, see Section 4.8, “eDirectory 9.1,” on page 41 .
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 or security auditing is done	Yes	Syslog on Linux.
Data on the wire is encrypted by default	Yes	The following data is encrypted on the wire: <ul style="list-style-type: none">◆ 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 user names and passwords are encrypted. Cluster credentials are stored encrypted in eDirectory.
Security is on by default	Yes	

11.2 Security Configuration

This section provides a summary of security-related configuration settings for Business Continuity Clustering.

- ◆ [Section 11.2.1, “BCC Configuration Settings,” on page 130](#)
- ◆ [Section 11.2.2, “Changing the NCS: BCC Settings Attributes in the BCC XML Configuration,” on page 131](#)
- ◆ [Section 11.2.3, “Disabling SSL for Inter-Cluster Communication,” on page 132](#)
- ◆ [Section 11.2.4, “Restricting the Network Address for Administration,” on page 134](#)

11.2.1 BCC Configuration Settings

[Table 11-2](#) lists the BCC configuration settings that are security-related or that impact the security of BCC.

Table 11-2 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
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)

Configuration Setting	Possible Values	Default Value	Recommended Value for Best Security
BCC Administrator group	Any LUM-enabled eDirectory group	bccgroup	Unique group used for BCC administration. See Section 5.3, "Configuring a BCC Administrator User and Group," on page 55.
Peer cluster CIMOM URL (same as the Inter-cluster communication scheme)	<code>http://cluster_ip_address</code> <code>cluster_ip_address</code> where <code>https://</code> is assumed	<code>cluster_ip_address</code> <code>https://</code> is assumed	Default value

11.2.2 Changing the NCS: BCC Settings Attributes 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 Micro Focus Support. Doing so can have adverse effects 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 in order for these settings to take effect. These are advanced settings that are intentionally not exposed in the BCC plug-in for iManager.

```
<!-- The contents of this file are generated automatically. Changes made to this
file will not be saved. -->
<bccSettings>
  <adminGroupName>bccgroup</adminGroupName>
  <authorizationCacheTTL>300</authorizationCacheTTL>
  <cimConnectTimeout>15</cimConnectTimeout>
  <cimReceiveTimeout>15</cimReceiveTimeout>
  <cimSendTimeout>15</cimSendTimeout>
  <idlePriorityThreshold>3</idlePriorityThreshold>
  <initialNormalThreads>3</initialNormalThreads>
  <initialPriorityThreads>2</initialPriorityThreads>
  <ipcResponseTimeout>45</ipcResponseTimeout>
  <maximumPriorityThreads>20</maximumPriorityThreads>
  <minimumPriorityThreads>2</minimumPriorityThreads>
  <resourceOfflineTimeout>300</resourceOfflineTimeout>
  <resourceOnlineTimeout>300</resourceOnlineTimeout>
  <scanForNewDevicesDelay>5</scanForNewDevicesDelay>
</bccSettings>
```

On Linux, the above XML is written to the `/etc/opt/novell/bcc/bccsettings.xml` file.

IMPORTANT: This file might be overwritten by Business Continuity Clustering at any time. Therefore, any changes to this file on Linux are ignored and lost. All changes should be made in eDirectory.

Table 11-3 provides additional information on each setting.

Table 11-3 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 that the authorization rights are cached in the BCC OpenWBEM provider.	300 seconds
<cimConnectTimeout>	BCC CIM client connect timeout in seconds.	15 seconds
<cimReceiveTimeout>	BCC CIM client receive timeout in seconds.	15 seconds
<cimSendTimeout>	BCC CIM client send timeout in seconds.	15 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.	45
<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
<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

11.2.3 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. 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. For more information on using iManager, see [NetIQ iManager Administration Guide](#).

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.

[Table 11-4](#) provides an example of possible combinations of scheme and port specifier for the <address> tag for values of the NCS:BCC Peers attribute.

Table 11-4 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

11.2.4 Restricting the Network Address for Administration

You can restrict the network address to the loopback address (127.0.0.1) to increase the security for the BCC Administrator user (bccadmin).

BCC makes a secure connection to OpenWBEM over port 5989 on both the remote and local boxes. This port can be changed.

The `cluster connection` command reports the status of the OpenWBEM connection from the last time a status update was performed. Typically, this occurs every 30 seconds on the cluster's master node, and every hour on its slaves. Running the following command forces a status update:

```
cluster refresh -p
```

OpenWBEM then makes an NCP connection to check the rights of the user who authenticated to OpenWBEM. The NCP connection itself goes to the loopback address.

11.3 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.

11.4 Security Information for Dependent Products

Table 11-5 provides links to security-related information for other products that can impact the security of BCC:

Table 11-5 Security Information for Other Products

Product Name	Links to Security Information
eDirectory	Security for eDirectory is provided by NCI (Novell International Cryptographic Infrastructure). See the <i>NICI 2.7x Administration Guide</i> (https://www.netiq.com/documentation/nici27x/nici_admin_guide/data/a20gkue.html).

Product Name	Links to Security Information
Identity Manager	“Security Best Practices” in the <i>NetIQ Identity Manager Security Guide</i>
iSCSI	For information about securing iSCSI targets and initiators for Linux iSCSI solutions on SLES, see “Mass Storage over IP Networks—iSCSI” (https://www.suse.com/documentation/sles-12/stor_admin/data/cha_iscsi.html) in the <i>SUSE Linux Enterprise Server 12 Administration Guide</i> (https://www.suse.com/documentation/sles-12/book_sle_admin/data/book_sle_admin.html).
Linux User Management (LUM)	See the <i>OES 2018 SP1: Linux User Management Administration Guide</i> .
Novell Cluster Services for Linux	In the <i>OES 2018 SP1: Novell Cluster Services for Linux Administration Guide</i> , see <ul style="list-style-type: none"> ◆ “IP Address Requirements” ◆ “Assigning Install Rights for Container Administrators or Non-Administrator Users”
Novell Storage Services for Linux	In the <i>OES 2018 SP1: NSS File System Administration Guide for Linux</i> , see: <ul style="list-style-type: none"> ◆ “Access Control for NSS” ◆ “File Access for Users” ◆ “Securing Access to NSS Volumes, Directories, and Files” ◆ “Security Considerations”
Small Footprint CIM Broker (SFCB)	“Web Based Enterprise Management using SFCB” (https://www.suse.com/documentation/sles-12/book_sle_admin/data/cha_wbem.html) in the <i>SUSE Linux Enterprise Server 12 Administration Guide</i> (https://www.suse.com/documentation/sles-12/book_sle_admin/data/book_sle_admin.html)

A

Console Commands for BCC

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.

IMPORTANT: For Novell Cluster Services console commands, see “[Console Commands for Novell Cluster Services](#)” in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.

- ◆ [Section A.1, “Using Console Commands,” on page 137](#)
- ◆ [Section A.2, “Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh,” on page 140](#)

A.1 Using Console Commands

[Table A-1](#) lists the server console commands for managing a business continuity cluster and gives a brief description of each command.

To execute a cluster console command, enter `cluster` followed by the command. For example, if this cluster belongs to 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.

IMPORTANT: You must be logged in as the `root` user, or any other user in `admin` or `ncsgroup`.

Some commands prompt the user for input. When using the command in a script, the prompt is not needed. You can use the `--noprompt` option for the command to bypass the prompted interaction.

For example, the `cluster migrate` command prompts the user to specify whether to continue migration if the cluster is experiencing connection issues. When you use the `cluster migrate` command in a script, add the `--noprompt` option to the command line to skip the prompt:

```
cluster migrate --noprompt [source/resource] [destination/nodename]
```

Table A-1 Console Commands for Business Continuity Clustering

Console Command	Description
<code>cluster connections [-a]</code>	Displays the connection status of the cluster. Specifying <code>-a</code> attempts to show the connection status of all clusters in the business continuity cluster.
<code>cluster credentials [peer_cluster]</code>	Lets you change the administrator user name and password that this cluster uses to connect to the specified peer cluster. The cluster you specify must belong to a cluster that has already been enabled for BCC.

Console Command	Description
cluster disable [resource]	<p>Disables BCC for the specified cluster resource. The resource you specify must belong to a cluster that has already been enabled for BCC. If no resource is specified, the entire cluster is disabled for BCC.</p> <p>IMPORTANT: Before you disable BCC for a given peer cluster, you must first disable BCC for each of the cluster resources running on that cluster.</p> <p>Before you disable BCC for a given cluster resource, remove the secondary peer clusters from its Assigned list. For information, see Section 8.7, “Disabling BCC for a Cluster Resource,” on page 99.</p> <p>After you have disabled BCC for all resources running on the cluster, remove the secondary peer clusters from the Assigned list of preferred nodes, then disable BCC for the cluster.</p> <p>If you disable BCC for a cluster by using the <code>cluster disable</code> console command, BCC is automatically disabled for those cluster resources that have been enabled for BCC. If you re-enable BCC for the cluster, you must re-enable BCC for each individual cluster resource that you want to be enabled for business continuity.</p> <p>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 BCC for an entire cluster.</p>
cluster enable [resource]	<p>Enables BCC for the specified resource. The resource you specify must belong to a cluster that has already been enabled for BCC. If no resource is specified, the entire cluster is enabled for BCC.</p> <p>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 business continuity cluster on a resource. Assigning a resource to specific clusters still must be done through iManager.</p>
cluster migrate [source/resource] [destination/nodename]	<p>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.</p>

Console Command	Description
<code>cluster nsmi</code>	<p>Execute a test run of an NSMI script for a single business continuity resource load or unload script. This command is to be used only for testing Business Continuity Clustering scripts.</p> <p>IMPORTANT: Ensure that the appropriate resource is offline in all peer clusters in the BCC before executing this command.</p> <p>The command prompts you for the resource name (such as <code>pool6_server</code>) and the type of BCC script (load or unload). The available scripts are listed, and you are prompted to enter the ID next to the script you want to use. Output is sent to the system log.</p>
<code>cluster refresh</code>	<p>This command should not be used except under the direction of Micro Focus Support.</p>
<code>cluster resetresources</code>	<p>Changes the state of all resources on this cluster to offline and secondary. This is a recovery procedure that should be run when a peer cluster in a business continuity cluster is brought back into service.</p> <p>The server where the command is issued must be a member of the cluster, but it does not need to be the one node that is currently in the cluster.</p> <p>You should run this command when only one node is a member of the cluster.</p> <ol style="list-style-type: none"> 1. After a failure, bring up one node in the cluster. All other nodes should remain powered off. 2. Run the <code>cluster resetresources</code> command. 3. Bring up the remaining nodes in the cluster.
<code>cluster resources [resource]</code>	<p>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.</p>
<code>cluster status</code>	<p>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.</p>
<code>cluster view</code>	<p>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 belongs to a business continuity cluster.</p>

A.2 Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh

Novell Cluster Services provides a framework for customized device scanning to BCC for Linux through an API that reads and executes commands in the `/opt/novell/ncs/bin/device_scan.sh` script. BCC automatically initiates a cluster scan for new devices for BCC-enabled cluster resources that have BCC load/unload scripts when cluster resources are migrated between peer clusters. When BCC calls the API, the script is executed on each node that is currently active in a cluster.

By default, the script is empty. You must add the Linux shell commands that you need to refresh the nodes in the cluster. Any changes that are made to the script are not overwritten when Novell Cluster Services is upgraded.

WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

- 1 In a text editor, open the script file `/opt/novell/ncs/bin/device_scan.sh`, add the Linux shell commands that scan your shared devices, then save the file.

The following is a sample script:

```
#!/bin/bash

## Logs to /var/log/messages with tag "bccd-scan"
/bin/logger -t bccd-scan "BCC is running script - device_scan.sh -- Scanning
for new devices"

## Rescan devices
rescan-scsi-bus.sh -wcl

## Rescan storage objects (OES 11 and later), such as for expanded NSS pools
nlvm rescan

## Add multipath command to rebuild maps if it applies.
# multipath
```

B

Configuration Worksheet for the BCC Drivers for Identity Manager

Use this worksheet to gather the information you need to configure the Business Continuity Clustering drivers for Identity Manager. Repeat this process for each connection between clusters. For information about how to set up the drivers, see [Chapter 6, “Configuring the Identity Manager Drivers for BCC,”](#) on page 67.

- ♦ [Section B.1, “Cluster1 to Cluster2,”](#) on page 141
- ♦ [Section B.2, “Cluster2 to Cluster1,”](#) on page 142

B.1 Cluster1 to Cluster2

First, create a driver instance on the Identity Manager node for the first cluster in order to synchronize information from the first cluster to the second cluster.

Clusters for This Driver Instance

Clusters Synchronized by This Driver Instance	Value for Your Network
Name of the source cluster	
Example: cluster1	
Name of the destination cluster	
Example: cluster2	

New Driver Set for the Source Cluster

Driver Set Information	Value for Your Network
Name of the BCC driver set for the source cluster	
Example: Cluster1 BCC Driver Set	
Context that contains the source cluster	
Example: cluster1.clusters.siteA.example	
TIP: Browse to select the context.	
Distinguished name of the Identity Manager node in the source cluster	
Example: cl1svr1.cluster1.clusters.siteA.example	
TIP: Browse to select the server.	
Create a new partition on this driver set	Deselect this option.

Driver Instance for the Connection from Source to Destination

Driver Set Information	Value for Your Network
Pre-configured driver template to import:	
	<code>BCCClusterResourceSynchronization.xml</code>
Driver name for this driver instance	
Example: cluster1_to_cluster2 BCCCR Sync	
Default: BCC Cluster Sync	
SSL certificate name to use for this driver instance	
Example: cluster1_to_cluster2 BCCCR Sync	
Default: BCC Cluster Sync (kmo)	
IP address or DNS name of the Identity Manager node in the destination cluster for this driver instance	
Example: 10.10.20.21	
Example: cl2svr1.cluster2.clusters.siteB.example	
Unique IP port to use for this driver instance	
Example: 2002 (default for resource sync)	
TIP: Both clusters must use the same port for this connection.	
Distinguished name of the source cluster	
Example: cluster1.clusters.siteA.example	
TIP: Browse to select the cluster.	
Landing zone for this driver instance	
Specify the distinguished name of the container where the cluster-enabled pool, NCP server, and volume objects for the destination cluster will be placed when they are synchronized to the source cluster for this driver instance.	
The container must already exist and must be specified using dot format without the tree name.	
Example: clusters.siteA.example	
TIP: Browse to select the context.	

B.2 Cluster2 to Cluster1

Next, create a second driver instance on the Identity Manager node for the second cluster in order to synchronize information from the second cluster to the first cluster.

Clusters for This Driver Instance

Clusters Synchronized by This Driver Instance	Value for Your Network
Name of the source cluster	
Example: cluster2	
Name of the destination cluster	
Example: cluster1	

New Driver Set for the Source Cluster

Driver Set Information	Value for Your Network
Name of the BCC driver set for the source cluster	
Example: Cluster2 BCC Driver Set	
Context that contains the source cluster	
Example: cluster2.clusters.siteB.example	
TIP: Browse to select the context.	
Distinguished name of the Identity Manager node in the source cluster	
Example: cl2svr1.cluster2.clusters.siteB.example	
TIP: Browse to select the server.	
Create a new partition on this driver set	Deselect this option.

Driver Instance for the Connection from Source to Destination

Driver Set Information	Value for Your Network
Preconfigured driver template to import:	
<code>BCCClusterResourceSynchronization.xml</code>	
Driver name for this driver instance	
Example: cluster2_to_cluster1 BCCCR Sync	
Default: BCC Cluster Sync	
SSL Certificate name for this driver instance	
Example: cluster2_to_cluster1 BCCCR Sync	
Default: BCC Cluster Sync (kmo)	
IP address or DNS name of the Identity Manager node in the destination cluster for this driver instance	
Example: 10.10.10.21	
Example: cl1svr1.cluster1.clusters.siteA.example	

Driver Set Information	Value for Your Network
<p>Unique IP port to use for this driver instance (same as the value you used for the Cluster1-to-Cluster2 driver instance)</p>	<p>Example: 2002 (default for resource sync)</p>
<p>TIP: Both clusters must use the same port for this connection.</p>	
<p>Distinguished name of the source cluster for this driver instance</p>	<p>Example: cluster2.clusters.siteB.example</p>
<p>TIP: Browse to select the cluster.</p>	
<p>Landing zone for this driver instance</p>	<p>Specify the distinguished name of the container where the cluster-enabled pool, NCP server, and volume objects for the destination cluster will be placed when they are synchronized to the source cluster for this driver instance.</p>
<p>The container must already exist and must be specified using dot format without the tree name.</p>	<p>Example: clusters.siteB.example</p>
<p>TIP: Browse to select the context.</p>	

C Setting Up Auto-Failover

Auto-failover is available for Business Continuity Clustering. To set up the auto-failover feature, you must enable it, and 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 might include but are not limited to the following:

- ◆ 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 clusters 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.

-
- ◆ [Section C.1, “Enabling Auto-Failover,” on page 145](#)
 - ◆ [Section C.2, “Creating an Auto-Failover Policy,” on page 146](#)
 - ◆ [Section C.3, “Refining the Auto-Failover Policy,” on page 146](#)
 - ◆ [Section C.4, “Adding or Editing Monitor Configurations,” on page 147](#)

C.1 Enabling Auto-Failover

To enable auto-failover for all Business Continuity Cluster resources in a cluster:

- 1 Log in to iManager as the BCC Administrator user.
- 2 In **Roles and Tasks**, click **Clusters > My Clusters**.
- 3 Select the check box next to the cluster, then select **Actions > Properties**.
You can also click the **Properties** button on the **Cluster Options** page for the cluster.
- 4 In the **Cluster Properties** dialog box, click the **Business Continuity** tab.
- 5 On the Business Continuity page, click the **Main** link.
- 6 Click the **AutoFailover** button.
- 7 Click the **Auto-Failover** link just under the tabs.
- 8 Select the **Enable Automatic Failover of Business Continuity Cluster Resources** check box, then click **Apply**.

- 9 Continue with [Section C.2, “Creating an Auto-Failover Policy,”](#) on page 146 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, and 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 **Apply** 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 that are automatically created by using the Cluster Membership Monitoring Settings table, you can add a condition that tests whether 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 whether 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. BCC ships with two monitors (nodecnt and nodepnt) that report an indication of health that represents either the percentage or number of nodes that do not belong to 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 BCC. 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 Linux 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.

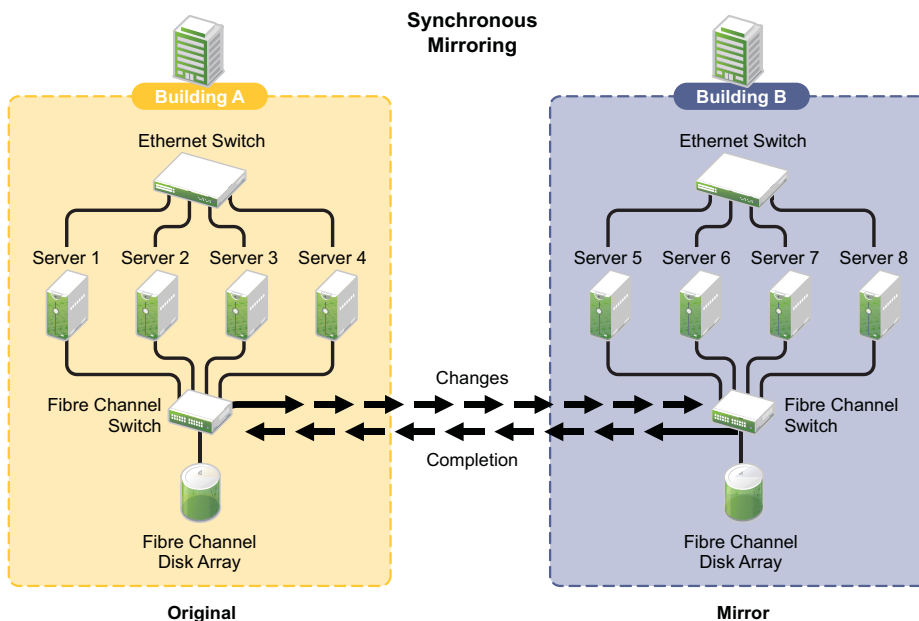
- 6 Click **Apply** to save your settings.

NOTE: See the [BCC NDK documentation \(http://developer.novell.com/documentation/cluster/ncss_enu/data/bktitle.html\)](http://developer.novell.com/documentation/cluster/ncss_enu/data/bktitle.html) for more information on creating custom failover policies.

D Configuring Host-Based File System Mirroring for NSS Pools

Several methods and scenarios exist for mirroring data between geographically separate sites. Each method has its own strengths and weaknesses. For a Business Continuity Clustering system, you need to choose either host-based mirroring or storage-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 storage-based mirroring or host-based file system 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 OES. NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. It is an alternative to storage-based synchronous replication options.

IMPORTANT: NSS pool snapshot technology does not work in a business continuity cluster.

- ♦ [Section D.1, "Creating and Mirroring NSS Pools on Shared Storage," on page 150](#)
- ♦ [Section D.2, "Creating NSS Volumes," on page 152](#)

- ♦ [Section D.3, “Novell Cluster Services Configuration and Setup,” on page 153](#)
- ♦ [Section D.4, “Checking NSS Volume Mirror Status,” on page 153](#)

D.1 Creating and Mirroring NSS Pools 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.

Prior to creating and mirroring NSS pools on shared storage, ensure that you have the following:

- ♦ All servers in the cluster are connected to a shared storage system.
- ♦ One or more drive arrays are configured on the shared storage system.
- ♦ The drives on the shared storage system have been initialized.
- ♦ NSS is installed and running. For information, see [“Installing and Configuring Novell Storage Services”](#) in the *OES 2018 SP1: NSS File System Administration Guide for Linux*.
- ♦ CIFS for Linux and AFP for Linux are available as advertising protocols for NSS pool cluster resources. If you plan to mark CIFS or AFP as an advertising protocol for the NSS pool resource, ensure that these protocols are installed and running when you create the pool resource. If you install the protocols after you create the pool, you can use the Clusters plug-in for iManager to add CIFS or AFP as advertising protocols.
- ♦ You need a static IP address for the pool resource. It must be in the same subnet as the cluster master IP address.

IMPORTANT: NSS pool snapshot technology is not supported for pool resources in a Novell Cluster Services cluster, and also does not work in a business continuity cluster.

To create and mirror NSS pools:

- 1 Start NSSMU by entering `nssmu` at the server console of a cluster server.
- 2 Select **Devices** from the NSSMU main menu and mark all shared devices as sharable for clustering.

On Linux, shared disks are not by default marked sharable for clustering. With a device marked as sharable for clustering, all partitions on that device are automatically sharable.

You can press F6 to individually mark devices as sharable.

- 3 From the NSSMU main menu, select **Pools**, press the Insert key, then type a name for the new pool you want to create.
- 4 Select the device on your shared storage where you want the pool created.
Device names might be labelled something like `/dev/sdc`.
- 5 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 6](#). If you want to cluster-enable the pool at a later date, change the default entry from **Yes** to **No**, select **Create**, and then go to [“Creating NSS Volumes” on page 152](#).

- 6 On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	<p>(Optional) The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is <code>cluster1</code> and the pool cluster resource name is <code>POOL1_SERVER</code>, then the default virtual server name is <code>CLUSTER1-POOL1-SERVER</code>.</p> <p>You can use the suggested name, or specify a different Virtual Server name for the cluster resource.</p> <p>You can modify the cluster resource name after the resource has been created by using the <code>cluster rename</code> command. Changing the resource name does not modify the pool name or the virtual server name.</p>
CIFS Server Name	<p>(Optional) If CIFS is installed and running, you can use this field to specify the name of the CIFS virtual server that CIFS clients see when they browse the network.</p> <p>If CIFS is installed and running, but CIFS is disabled as an advertising protocol, this field is not available (dimmed).</p> <p>If CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code>.</p> <p>CIFS is disabled by default as an advertising protocol. You can select the CIFS check box to enable it. By default, the NCP virtual server name is suggested as the CIFS virtual server name. You can use the suggested name or specify a customized name for the CIFS virtual server name.</p> <p>If desired, specify a new name for the CIFS virtual server. The name can be up to 15 characters, which is a restriction of the CIFS protocol.</p>
IP Address	<p>Specify an IP address for the pool cluster resource. Tab between the address fields. The address is IPv4 format, such as <code>10.10.10.243</code>.</p> <p>Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster hosts the pool.</p>
Advertising Protocols	<p>Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the OES Trustee model.</p> <p>Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.</p> <p>If the CIFS or AFP protocols are not installed and running, selecting the corresponding CIFS or AFP check box has no effect.</p>

Parameter	Action
Online Resource after Create	<p>The check box is deselected by default and dimmed so that you cannot change the setting.</p> <p>The pool is currently active on the server. You must deactivate the pool from the server before attempting to bring the resource online. You should also configure the resource load, unload, and monitor scripts before you bring the resource online.</p>
Define Additional Properties	<p>Select the Define Additional Properties check box.</p> <p>This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.</p>

- 7 Select **Create** to create and cluster-enable the pool.
- 8 Repeat [Step 3](#) to [Step 7](#) for each additional pool you want to create on shared storage.
- 9 Select **Partitions** from the NSSMU main menu.
- 10 Select the partition you want to mirror (this is the partition that was created when you created the pool), then press F3.
- 11 Enter a name for the software RAID 1 device that will be created.

For more information on configuring software RAID for NSS, see “[Managing NSS Software RAID Devices](#)” in the *OES 2018 SP1: NSS File System Administration Guide for Linux*.
- 12 Select the shared device with free space that you want to use as the mirror, 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. You can select up to three shared devices.
- 13 Repeat [Step 9](#) through [Step 12](#) for each additional shared pool that you want to mirror.
- 14 Continue with “[Creating NSS Volumes](#)” on page 152.

D.2 Creating NSS Volumes

After you create the pool and its mirror, you must create an NSS volume on each pool resource so that it works properly as a cluster resource. To create an NSS volume on a shared pool, follow the instructions in “[Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes](#)” in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.

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 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 volume on the mirrored NSS partition and pool, if you did not cluster-enable the NSS pool on the mirrored partition when you created it, do so by following the instructions in “[Cluster-Enabling an Existing NSS Pool and Its Volumes](#)” in the *OES 2018 SP1: Novell Cluster Services for Linux 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 an assigned server by entering `nsscon volumes` at the server console.
- 2 Enter `exit` to exit the NSSCON utility.
- 3 Check the mirror status of the mirrored partition by starting NSSMU and selecting the RAID device you created in [Step 11 on page 152](#).

NSSMU displays the mirror status of the RAID device. If the partitions are still synchronizing, it could take time for NSSMU to display the mirror status.

- 4 Migrate the pool to another server in the cluster. Check again to ensure that the volumes on the pool are mounted by entering `nsscon volumes` at the server console.

Enter `exit` to exit the NSSCON utility.

- 5 Using NSSMU, check the mirror status of the partition again.

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 `nlvm rescan` command on a server in each of the peer clusters.

E

Using Dynamic DNS with BCC

One deployment consideration is how the client computers reconnect to services when cluster resources fail over to a peer cluster in the disaster recovery location. Typically, the disaster recovery site is a geographically dispersed site in a different network subnet, which forces the cluster resources to use a different IP address range. This section describes how to configure Business Continuity Clustering (BCC) to act as a dynamic DNS (Domain Name System) client.

- ♦ [Section E.1, “Requirements and Assumptions,” on page 155](#)
- ♦ [Section E.2, “Configuring the DNS Server for Dynamic DNS,” on page 156](#)
- ♦ [Section E.3, “Configuring the Cluster Resources for Dynamic DNS,” on page 162](#)
- ♦ [Section E.4, “Testing the Dynamic DNS Solution,” on page 166](#)

E.1 Requirements and Assumptions

- ♦ [Section E.1.1, “Software Requirements,” on page 155](#)
- ♦ [Section E.1.2, “DNS Server,” on page 155](#)
- ♦ [Section E.1.3, “TSIG Keys,” on page 156](#)
- ♦ [Section E.1.4, “DNS Record Time-to-Live Values,” on page 156](#)

E.1.1 Software Requirements

Install and configure your business continuity cluster using the following components:

- ♦ Business Continuity Clustering
- ♦ Novell Cluster Services
- ♦ Open Enterprise Server 2018 SP1

To integrate dynamic DNS in your BCC solution, use the `bcc_dyn_dns.pl` NSMI script that is included in the BCC software CD in the `<media>/nsmi_scripts/linux` directory.

E.1.2 DNS Server

The examples used in this BCC solution use the ISC BIND 9 DNS server on Linux. It is not a requirement that you use the ISC BIND DNS server, but your DNS server must meet the following requirements:

- ♦ Your DNS server must support the Dynamic DNS standard in RFC 2136.
- ♦ If you use the ISC BIND DNS server, you should install the `bind` and `bind-utils` RPM packages on each node in your business continuity cluster.

We recommend that your DNS strategy involve redundant DNS servers that make use of secure zone transfers. Furthermore, redundant DNS servers should be implemented at each individual disaster recovery site.

Another option for your DNS servers is to put them in your Novell Cluster Services cluster. This creates a DNS service that is extremely resilient to failure. For information, see “[Configuring DNS with Novell Cluster Services](#)” in the *OES 2018 SP1: DNS/DHCP Services for Linux Administration Guide*.

E.1.3 TSIG Keys

TSIG (Transaction Signature) keys are used in the examples to authenticate dynamic updates of the DNS server. It is not a requirement that you use TSIG. Other methods of authorizing updates can be used instead, such as DNSSEC (DNS Security Extensions). In addition, good security requires more than authorization keys. Logging, monitoring, firewalls, intrusion detection systems, and so on should all be employed to keep your systems and network safe from unwanted access. However, it is beyond the scope of this document to cover these alternatives.

The TSIG `dnssec-keygen` utility should be automatically installed as part of the ISC BIND 9 DNS server package.

E.1.4 DNS Record Time-to-Live Values

Selecting the proper Time-to-live (TTL) value for the DNS records can be challenging. If the values are too short, the DNS traffic on your network can increase dramatically. If the values are too long, the end users are unable to reconnect to the cluster resources after a BCC migration until the DNS records expire. There is no perfect TTL value. Each customer and environment is unique and has different needs and goals. You should experiment with the TTL values while monitoring the DNS traffic on your network to find the ideal value for your network.

E.2 Configuring the DNS Server for Dynamic DNS

Begin by configuring the DNS server so that it accepts dynamic updates to a particular zone and authenticates these updates using TSIG (Transaction Signature) keys.

- [Section E.2.1, “Creating the TSIG Keys for DNS Server Authentication,” on page 156](#)
- [Section E.2.2, “Configuring the DNS Server with the Public Key,” on page 157](#)
- [Section E.2.3, “Configuring the DNS Server Zones,” on page 158](#)
- [Section E.2.4, “Testing the DNS Server,” on page 158](#)

E.2.1 Creating the TSIG Keys for DNS Server Authentication

TSIG keys are used to authenticate dynamic updates of the DNS server. Use the `dnssec-keygen` utility to create the public and private TSIG key files in the following format:

```
K<cluster_dns_name>.+157+<random number>.key  
K<cluster_dns_name>.+157+<random number>.private
```

- 1 On a node in one of the peer clusters, log in as the `root` user, then open a terminal console.
- 2 Use the `dnssec-keygen` utility to create the public and private TSIG keys by entering

```
dnssec-keygen -a HMAC-MD5 -b 512 -n HOST cluster_dns_name
```

The `-a` option specifies the cryptographic algorithm. For dynamic DNS, this must be HMAC-MD5.

The `-b` options specifies the number of bits in the key. You should use the strongest encryption possible, which for HMAC-MD5 is 512.

The `-n` option is the name type. Because a computer is updating the DNS server, use the `HOST` name type.

Replace `cluster_dns_name` with the name of the host. For BCC, the cluster node that hosts the Novell Cluster Services Master IP Address resource updates the DNS server. Because this can be any node in the cluster, use the fully qualified name of the cluster as the host name.

For example, enter

```
dnssec-keygen -a HMAC-MD5 -b 512 -n HOST cluster1.clusters.site1.company.com
```

This generates the public and private key files:

```
Kcluster1.clusters.site1.company.com.+157+60303.key
```

```
Kcluster1.clusters.site1.company.com.+157+60303.private
```

`60303` represents a randomly generated number created by the utility.

- 3 Store these files in a secure location.

The DNS administrator uses these keys to configure your master DNS server.

- 4 Continue with [Section E.2.2, "Configuring the DNS Server with the Public Key,"](#) on page 157.

E.2.2 Configuring the DNS Server with the Public Key

Modify the DNS Server configuration to use the public TSIG key you generated in [Section E.2.1, "Creating the TSIG Keys for DNS Server Authentication,"](#) on page 156. You can place the public key information directly in the `/etc/named.conf` file, but it is more secure to place it in a separate location where the key file can be protected.

- 1 On the DNS Server, open a terminal console, then log in as the `root` user.
- 2 Open the `/etc/named.conf` file in a text editor, add the following line before the zone configuration, then save the changes:

```
include "keys.conf";
```

- 3 Go to the `/var/lib/named` directory, then use a text editor to create a `keys.conf` file.
- 4 In the `keys.conf` file, create a section for each public key you need to add.

The format of the key section is:

```
key <cluster_dns_name>. {  
    algorithm <cryptographic algorithm>;  
    secret "<the public key secret>";  
};
```

The `cluster_dns_name` is the same name you used when creating the key with the `dnssec-keygen` utility. This name is also found in the public key file that `dnssec-keygen` created.

The cryptographic algorithm must be `HMAC-MD5`.

The public key secret is the Base64-encoded secret found in the public key file that the `dnssec-keygen` utility created. You can copy and paste the secret from the public key file to the `/var/lib/named/keys.conf` file. To continue our example, the key section for the `/var/lib/named/keys.conf` file might look like this:

```
key cluster1.clusters.site1.company.com. {  
    algorithm HMAC-MD5;  
    secret "SCUT8rIUoGBYvcI1Iok7tY7YvcEaHaM3zusCxXmboBxVcJvUxr335HCg  
lXcDQRPrJrzIKQhH4dJ4cY10ebOJFw==" ;  
};
```

- 5 Save the file.
- 6 Continue with [Section E.2.3, “Configuring the DNS Server Zones,”](#) on page 158.

E.2.3 Configuring the DNS Server Zones

To configure the DNS zones to accept authorized DNS updates:

- 1 On the DNS server, open a terminal console, then log in as the `root` user.
- 2 Open the `/etc/named.conf` file in a text editor, then add the `allow-update` keyword and key in the zone configuration sections for regular lookups and reverse lookups.

For example, this is a sample zone section:

```
zone "clusters.site1.company.com" in {
    file "dyn/clusters.site1.company.com";
    type master;
    allow-update {
        key cluster1.clusters.site1.company.com.;
    };
};

zone "1.1.10.in-addr.arpa" in {
    file "dyn/10.1.1.zone";
    type master;
    allow-update {
        key cluster1.clusters.site1.company.com.;
    };
};
```

- 3 Save the changes.
- 4 Restart the DNS Server to ensure that the new configuration is imported.
Your DNS Server is now configured to accept secure dynamic updates.
- 5 Continue with [Section E.2.4, “Testing the DNS Server,”](#) on page 158.

E.2.4 Testing the DNS Server

Before you continue to set up the BCC for dynamic DNS updating, verify that your DNS server accepts secure dynamic updates.

- 1 Install the `bind-utils` RPM on a Linux client computer.
- 2 Ensure that the Linux client computer is configured to use the DNS Server you want to test.
Alternately, you can force the `dig` utility to query a specific DNS Server.
- 3 Use the `dig` utility to perform a baseline test on any server to view its current IP address.

Test Equipment	Sample Value
DNS Server IP address	10.1.1.172
Linux client computer	wkstn1.clusters.site1.company.com
FTP server (testing to see its IP address)	10.1.1.215 ftp.clusters.site1.company.com

For example, enter:

```
dig #10.1.1.172 ftp.clusters.site1.company.com
```

The output shows that the IP address for ftp.clusters.site1.company.com is 10.1.1.215:

```
; <<>> DiG 9.3.2 <<>> @10.1.1.172 ftp.clusters.site1.company.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 47449
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 1

;; QUESTION SECTION:
ftp.clusters.site1.company.com.    IN      A

;; ANSWER SECTION:
ftp.clusters.site1.company.com. 10 IN      A      10.1.1.215

;; AUTHORITY SECTION:
clusters.site1.company.com. 120 IN      NS
wkstnl.clusters.site1.company.com

;; ADDITIONAL SECTION:
wkstnl.clusters.site1.company.com. 120 IN  A      10.1.1.172

;; Query time: 0 msec
;; SERVER: 10.1.1.172#53(10.1.1.172)
;; WHEN: Tue Aug 14 17:19:55 2008
;; MSG SIZE rcvd: 98
```

- 4 Use the `-x` option for the dig utility to perform a baseline test to check the reverse lookup records in the DNS server by the IP address.

For example, enter

```
dig @10.1.1.172 -x 10.1.1.215
```

The output shows the DNS name for the FTP server is ftp.clusters.site1.company.com:

```
; <<>> DiG 9.3.2 <<>> @10.1.1.172 -x 10.1.1.215
; (1 server found)
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 34957
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
215.1.1.10.in-addr.arpa.    IN      PTR

;; ANSWER SECTION:
215.1.1.10.in-addr.arpa. 10 IN      PTR      ftp.clusters.site1.company.com.

;; AUTHORITY SECTION:
1.1.10.in-addr.arpa.    120 IN      NS
wkstnl.clusters.site1.company.com.1.1.10.in-addr.arpa.

;; Query time: 0 msec
;; SERVER: 10.1.1.172#53(10.1.1.172)
;; WHEN: Tue Aug 14 17:32:11 2008
;; MSG SIZE rcvd: 127
```

- 5 Securely copy the public and private keys created in [Section E.2.1, "Creating the TSIG Keys for DNS Server Authentication,"](#) on page 156 to your home directory on the Linux client computer (such as `/home/yourhomedir/`).

These keys have filenames of the form `K<name>.+157+<random number>.key` (the public key) and `K<name>.+157+<random number>.private` (the private key).

- 6 Use the `nsupdate` utility to update an A record on the DNS server to change its IP address to 10.1.1.216.

```
nsupdate -v -k path_to_private_key_file
```

For example, enter

```
nsupdate -v -k /home/yourhomedir/  
Kcluster1.clusters.site1.company.com.+157+60303.private  
  
> server 10.1.1.172 53  
> update delete ftp.clusters.site1.company.com. A  
> update add ftp.clusters.site1.company.com. 300 A 10.1.1.216  
> show  
Outgoing update query:  
;; ->>HEADER<<- opcode: UPDATE, status: NOERROR, id:      0  
;; flags: ; ZONE: 0, PREREQ: 0, UPDATE: 0, ADDITIONAL: 0  
;; UPDATE SECTION:  
ftp.clusters.site1.company.com. 0 ANY      A  
ftp.clusters.site1.company.com. 300 IN     A          10.1.1.216  
  
> send  
> quit
```

- 7 Use the `nsupdate` utility to update the PTR record used for reverse lookups.

For example, enter

```
nsupdate -v -k /home/yourhomedir/  
Kcluster1.clusters.site1.company.com.+157+60303.private  
  
> server 10.1.1.172 53  
> update delete 215.1.1.10.in-addr.arpa PTR  
> update add 216.1.1.10.in-addr.arpa 300 PTR ftp.clusters.site1.company.com  
> show  
Outgoing update query:  
;; ->>HEADER<<- opcode: UPDATE, status: NOERROR, id:      0  
;; flags: ; ZONE: 0, PREREQ: 0, UPDATE: 0, ADDITIONAL: 0  
;; UPDATE SECTION:  
215.1.1.10.in-addr.arpa. 0      ANY      PTR  
216.1.1.10.in-addr.arpa. 300   IN       PTR      ftp.clusters.site1.company.com.  
  
> send  
> quit
```

- 8 Use the `dig` utility to verify that the changes made in [Step 6](#) and [Step 7](#) occurred on the DNS Server:

8a Use the `dig` utility with the DNS name to verify the IP address.


```

dig @10.1.1.172 ftp.clusters.sitel.company.com

; <<> DiG 9.3.2 <<> @10.1.1.172 ftp.clusters.sitel.company.com
; (1 server found)
;; global options: printcmd
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 35080
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 1

;; QUESTION SECTION:
;ftp.clusters.sitel.company.com.    IN      A

;; ANSWER SECTION:
ftp.clusters.sitel.company.com. 300 IN    A      10.1.1.216

;; AUTHORITY SECTION:
clusters.sitel.company.com. 120    IN      NS
wkstn1.clusters.sitel.company.com.

;; ADDITIONAL SECTION:
wkstn1.clusters.sitel.company.com. 120 IN    A      10.1.1.172

;; Query time: 0 msec
;; SERVER: 10.1.1.172#53(10.1.1.172)
;; WHEN: Tue Aug 14 17:50:13 2008
;; MSG SIZE rcvd: 98

```

8b Use the dig utility with the IP address to find the DNS name.

```

dig @10.1.1.172 -x 10.1.1.216

; <<> DiG 9.3.2 <<> @10.1.1.172 -x 10.1.1.216
; (1 server found)
;; global options: printcmd
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 14497
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;216.1.1.10.in-addr.arpa.          IN      PTR

;; ANSWER SECTION:
216.1.1.10.in-addr.arpa. 300    IN      PTR
ftp.clusters.sitel.company.com.

;; AUTHORITY SECTION:
1.1.10.in-addr.arpa.      120    IN      NS
wkstn1.clusters.sitel.company.com.1.1.10.in-addr.arpa.

;; Query time: 6 msec
;; SERVER: 10.1.1.172#53(10.1.1.172)
;; WHEN: Tue Aug 14 17:55:01 2008
;; MSG SIZE rcvd: 127

```

9 Use the dig utility in a reverse lookup for the old IP address to ensure that it does not return an answer.

```

dig @10.1.1.172 -x 10.1.1.215

; <<>> DiG 9.3.2 <<>> @10.1.1.172 -x 10.1.1.215
; (1 server found)
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 49360
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;215.1.1.10.in-addr.arpa.      IN      PTR

;; AUTHORITY SECTION:
1.1.10.in-addr.arpa.      120     IN      SOA  wkstn1.clusters.sitel.company.com.
root.wkstn1.clusters.sitel.company.com. 2007032716 10800 3600 604800 86400

;; Query time: 0 msec
;; SERVER: 10.1.1.172#53(10.1.1.172)
;; WHEN: Tue Aug 14 17:55:07 2008
;; MSG SIZE rcvd: 105

```

- 10** If your setup passes the tests in this section, continue with [Section E.3, “Configuring the Cluster Resources for Dynamic DNS,”](#) on page 162. Otherwise, go back to try again.

E.3 Configuring the Cluster Resources for Dynamic DNS

After the DNS server is configured for dynamic DNS, you must configure each of the BCC-enabled cluster resources to take advantage of it. You modify the BCC load scripts for each of the cluster resources so that the script automatically updates the DNS server with the correct IP address of the given resource during a BCC migration to a peer cluster.

IMPORTANT: In each of the peer clusters, the keys and the BCC load script for each resource will differ. You assign the resource an IP address that is consistent with the subnet into which the BCC-enabled cluster resource is being migrated or failed over. This allows the DNS server to be updated when the resource fails over from the primary cluster to the secondary cluster. The same thing needs to be done on the primary cluster so the DNS server is automatically updated on failback from the secondary cluster to the primary cluster. You must perform these steps for every BCC-enabled resource in the business continuity cluster.

The BCC load and unload scripts for BCC-enabled cluster resources are used to automate any tasks that need to occur when the resource is failed over to a peer cluster. For example, during a BCC migration, the BCC scripts manage the storage and update eDirectory. The BCC scripts are based on Perl, which means that you need to create Perl-based wrappers for the `nsupdate` utility commands before placing them in the BCC load script. Typically, the line for the `nsupdate` utility would be the last command to run in the script, but this might not be true for all environments.

- ◆ [Section E.3.1, “Modifying the BCC Load Script,”](#) on page 163
- ◆ [Section E.3.2, “Public and Private Keys,”](#) on page 165
- ◆ [Section E.3.3, “Testing the Perl Wrapper Script,”](#) on page 166

E.3.1 Modifying the BCC Load Script

- 1 Open iManager, then log in as the BCC Administrator user.
- 2 In **Roles and Tasks**, select **Clusters > My Clusters**, then click the cluster name of the secondary peer cluster for the resource. (That is, choose a peer cluster where the resource is not currently assigned and running.)
- 3 Click the **Cluster Options** tab.
- 4 Click the name link of the BCC-enabled resource where you want to add dynamic DNS support. This opens the Properties page for the selected resource.
- 5 Select the **Business Continuity** tab, click **SAN Management**, then click **New** in the **BCC Load Script** table.

SAN Management Script Details ?

View or modify the name, CIM details, parameters, or logic in the SAN management script.

Name:

Description:

CIMOM IP/DNS:

Namespace: Port: Secure

Username:

Password:

Script Parameters

New | Edit | Delete 0 Item(s)

Name	Value
No Items	

CIM Enabled Synchronous Edit Flags Flags:

- 6 Specify the following settings for the script:

Parameter	Description
CIM-Enabled	Used for CIM (Common Information Model) and SMI-S (Storage Management Initiative Specification) enabled scripts that manage a physical SAN. These are not used for the dynamic DNS script. Deselect the check box. After this check box is deselected, the CIMON IP/DNS, Namespace, Port, Secure, User name, and Password controls are all disabled.
Name	Specify the name of the script that is shown in the BCC log files (such as in /var/log/messages). The name should give enough information that the entries in the log file are meaningful. For example, a name of <code>Dynamic DNS Update - Resource Test1</code> identifies what the script does and which resource it is acting upon.
Description	Specify information that helps you understand the nature of the script. This information is not displayed in any of the BCC log files.

Parameter	Description
CIMOM IP/DNS	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
Namespace	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
Port	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
Secure	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
User name	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
Password	This field is used by CIM or SMI-S enabled scripts. It is not needed for the dynamic DNS script.
Script Parameters	The script parameters are used to customize the dynamic DNS script to work with a particular DNS server. Go to Step 7 to enter the parameters.
Script	Copy and paste the dynamic DNS script into this edit box.
Synchronous	When enabled (selected), this option synchronizes the execution of multiple BCC load and unload scripts. This is not necessary for the dynamic DNS script, so leave it disabled (deselected).
Edit Flags	This is an advanced option that should only be enabled when instructed to do so by Micro Focus Support. Leave it disabled (deselected).

7 Specify the Script Parameters.

7a Click **New** to insert a new editable row into the **Script Parameters** table.

7b Add the parameter name and value.

The left field is the parameter name and the right field is the parameter value.

7c When you are done adding a name/value pair, click **OK** to save the parameter.

7d Repeat the process for each name/value pair.

The following are the name/value pairs that must be entered into the **Script Parameters** table for the dynamic DNS script:

Name	Sample Value	Description
DNS_SERVER_ADDR	10.1.1.172	Specify the IP address of the DNS master server.
HOST_NAME	cluspool1.clusters.site1.company.com	Specify the hostname of the cluster resource whose IP address needs to be updated when it is migrated to this peer cluster.
HOST_RECORD_TTL	60	Specify the time-to-live value of the DNS record in seconds.

Name	Sample Value	Description
HOST_IP	10.1.20.216	Specify the new IP address for the cluster resource in the subnet of this peer cluster.
KEY_FILE	<code>/mnt/bcc-master/dyndns/keys/ Kcluster1.clusters.sitel.company.c om.+157+60303.private</code>	Specify the location of the private key file. Remember that the public key must be in this same directory.

- Click **OK** to return to the Cluster Resource Properties page for the resource you are modifying, then click **OK** again on the Cluster Resource Properties page to save the new dynamic DNS BCC load script.

E.3.2 Public and Private Keys

The BCC load and unload scripts always run on the node that is hosting the Novell Cluster Services master resource (that is, the `Master_IP_Address_Resource`). This resource can be hosted on any node in the Novell Cluster Services cluster, which means the BCC load and unload scripts can also be executed on any node in the cluster.

The Perl wrapper script for the `nsupdate` utility needs access to both the private and public keys created in [Section E.2.1, “Creating the TSIG Keys for DNS Server Authentication,” on page 156](#). The `nsupdate` utility needs access only to the private key. However, for historical reasons, the public key must be in the same location as the private key. The files that contain the keys must be available on all nodes in the cluster. This can be accomplished in either of the following ways:

- ♦ **Copy to the same location on each node in the cluster (not recommended).** Although this is simple and relatively quick to do initially, it can create significant maintenance issues. If the keys ever change, they must be copied to all nodes in the cluster. Failure to copy the keys to all nodes in the cluster creates the potential for failure in the dynamic DNS update process. You must also remember to copy the keys to any nodes you add to the cluster at a later date.
- ♦ **Create a cluster resource that contains the keys (recommended).** Create a shared volume using any standard journaled Linux POSIX file system, such as Ext3. Configure the cluster resource with the **Resource Follows Master** setting enabled on the **Policies > Resource Behavior** area on the Cluster Resource Properties page. This setting forces the given resource to always be hosted by the same node that is hosting the Novell Cluster Services master resource. The keys can then be copied to the file system hosted by this resource, which makes them available to the same node that is hosting the Novell Cluster Services master resource, and to the BCC load and unload scripts. This option takes a bit more configuration time, but results in easier maintenance. If the keys change, they only need to be copied to the file system hosted by this resource. In addition, if a node is added to the cluster, the new node automatically has access to the keys if it ever becomes the Novell Cluster Services master.

- Create a 10 MB shared volume with the Ext3 file system mounted at `/mnt/bcc-master`.

For example, name the resource `bcc-master`. Ensure that it is configured with the **Resource Follows Master** enabled.

For information about creating a shared Linux POSIX file system, see [“Configuring and Managing Cluster Resources for Shared LVM Volume Groups”](#) in the *OES 2018 SP1: Novell Cluster Services for Linux Administration Guide*.

- 2 Create the `/mnt/bcc-master/dyndns/keys` directory on the shared volume.
- 3 Copy the public and private key files to the directory.
- 4 Ensure that the Linux POSIX permissions are set so that the `root` user is the only user who has access.

E.3.3 Testing the Perl Wrapper Script

The dynamic DNS script cannot be compiled directly by a Perl interpreter because it does not recognize the BCC variables on lines 62–66 of the script. It is a simple matter to test the dynamic DNS script by manually and temporarily replacing the variables with values.

- 1 Copy the script to the local file system on the Novell Cluster services master node (such as `/tmp/bcc_dyn_dns.pl`).
- 2 Open the copy of the script in a text editor, then modify the BCC variables by modifying lines 62 through 66 of the script.

The uppercase string surrounded by the percent character (`%`) is the BCC variable and should be replaced with the real value (such as `%DNS_SERVER_ADDR%`).

For example, lines 62–66 of the script could be modified like the following:

```
my $dns_server_addr = "10.1.1.172";

my $host_name      = "ftp.clusters.site1.company.com";

my $host_record_ttl = "60";

my $host_ip        = "10.1.1.216";

my $key_file       = "/mnt/bcc/dyndns/keys/
Kcluster1.clusters.site1.company.com.+157+60303.private";
```

Remember to replace the values with the actual values for your environment.

- 3 Test the solution by invoking the Perl interpreter on the script via the following command

```
perl -w /tmp/bcc_dyn_dns.pl
```

This executes the dynamic DNS script and modifies your DNS server. You can test the results as outlined in [Section E.2.4, “Testing the DNS Server,” on page 158](#).

E.4 Testing the Dynamic DNS Solution

After you have created the BCC load scripts for each resource in each of the peer clusters, you are ready to test the overall dynamic DNS solution.

- 1 Perform a BCC migrate (failover) of the test resource from the primary cluster to a secondary peer cluster.

The dynamic DNS script runs before the cluster resource is brought online on the secondary cluster. If everything works correctly, the DNS server should be updated by the time the cluster resource comes online on the secondary cluster.

- 2 Test the results of the dynamic DNS update on the secondary cluster by following the steps in [Section E.2.4, “Testing the DNS Server,” on page 158](#).

- 3 Perform a BCC migrate (failback) of the test resource from the secondary peer cluster to the primary cluster.

The dynamic DNS script runs before the cluster resource is brought online on the primary cluster. If everything works correctly, the DNS server should be updated by the time the cluster resource comes online on the primary cluster.

- 4 Test the results of the dynamic DNS update on the primary cluster by following the steps in [Section E.2.4, “Testing the DNS Server,”](#) on page 158.

F Using Virtual IP Addresses with BCC

One deployment consideration is how the client computers reconnect to services when cluster resources fail over to a peer cluster in the disaster recovery location. Typically, the disaster recovery site is a geographically dispersed site in a different network subnet, which forces the cluster resources to use a different IP address range. In Business Continuity Clustering (BCC), you can configure BCC to use virtual IP addresses for BCC-enabled resources as an alternate approach to assigning secondary IP addresses to resources.

- ◆ [Section F.1, “Understanding Internal Virtual IP Networks,” on page 169](#)
- ◆ [Section F.2, “Virtual IP Address Benefits,” on page 170](#)
- ◆ [Section F.3, “Planning a Virtual IP Network Implementation,” on page 172](#)
- ◆ [Section F.4, “Configuring a Virtual Router with OSPF,” on page 173](#)
- ◆ [Section F.5, “Configuring Virtual IP Addresses,” on page 174](#)

F.1 Understanding Internal Virtual IP Networks

To use virtual IP addresses, you set up an internal virtual IP network on each cluster node. The network consists of the following:

- ◆ [Section F.1.1, “Virtual Adapter,” on page 169](#)
- ◆ [Section F.1.2, “Host Mask,” on page 169](#)
- ◆ [Section F.1.3, “Internal Router,” on page 170](#)

F.1.1 Virtual Adapter

A virtual adapter is a software-based adapter. A virtual adapter behaves like a conventional loopback interface with external visibility. LAN routers can maintain accurate information on available routes to the virtual adapter destination. The last hop on a router path to cluster resources occurs inside the cluster node itself. A cluster resource’s IP address is bound to the virtual adapter instead of to a physical adapter.

F.1.2 Host Mask

Virtual adapters support configuring virtual IP addresses with a host mask. The mask treats each resource as a separate network segment. This allows each cluster resource to have its own entry in the routing table of the internal router on the cluster node where it resides.

F.1.3 Internal Router

An internal router is a software-based router that runs on the operating system in a cluster node. On Linux, the internal router is set up with the open source Quagga Routing Software Suite.

Quagga Routing Software Suite is an open source implementation of major routing protocols for IPv4 and IPv6, including OSPFv2, OSPFv3, RIPv1, RIPv2, RIPv3, and BGPv4. It is available under the [GNU General Public License version 2.0](http://www.gnu.org/licenses/old-licenses/gpl-2.0.html) (<http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>). For information, see www.quagga.net. For information about Quagga, see the [Quagga Routing Suite Web site](http://www.nongnu.org/quagga/index.html) (<http://www.nongnu.org/quagga/index.html>).

F.2 Virtual IP Address Benefits

In spite of their simplicity, virtual IP addresses offer the following advantages over their physical counterparts:

- ♦ **Improves availability.** The virtual IP addresses are bound to virtual adapters instead of physical adapters. The host mask for the virtual adapter allows each cluster resource to have its own entry in the routing tables. The virtual IP addresses are resilient to physical interface failures because we know if a resource destination is reachable.
- ♦ **Improves mobility.** The virtual IP addresses are not tied to any IP address range of a physical network segment. A service can be moved between physical segments without the need to change the resource IP address.
- ♦ **Simplifies name resolution.** The association of cluster resource names to IP addresses can be maintained independently of the location of the cluster resource in the business continuity cluster.

These advantages exist because virtual IP addresses are purely virtual and are not bound to a physical network component. Each of these advantages is discussed in more detail below.

- ♦ [Section F.2.1, “High Availability,” on page 170](#)
- ♦ [Section F.2.2, “Unlimited Mobility,” on page 171](#)
- ♦ [Section F.2.3, “Automatic Name Resolution,” on page 171](#)

F.2.1 High Availability

Each cluster node 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. The virtual IP addresses of the cluster resources are highly available because each resource has its own entry in the routing tables of the LAN routers. This allows you to know whether a destination is reachable. However, with secondary IP addresses, you know only whether there is a route to a segment.

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 is always up. Also, because all other real links are forcibly made to act as intermediate links, their failures are easily worked around by the dynamic routing protocols.

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 physical adapter failures if the server has multiple adapters.

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

- ♦ A multi-homed server with a virtual IP address no longer needs to carry multiple DNS entries for its name in the naming system.
- ♦ If one of the subnets that a server interfaces to fails completely or is taken out of service for maintenance, the routing protocols can reroute the packets addressed to the virtual IP address through one of the other active subnets.

F.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 be transparently failed-over.

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, carry their virtual network to wherever they are moved. There is an implicit assumption here 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 it is required 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.

This unlimited mobility of virtual IP addresses is an advantage to network administrators, offering them 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, unless everyone accesses a service via its DNS name instead of its IP address, an IP address change can break the service for the IP address users. In contrast, if the service is hosted on a virtual IP address, the necessity of changing the DNS entries for the service is eliminated, and the service is not broken even for those who use the IP address instead of the DNS name.

F.2.3 Automatic Name Resolution

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.

F.3 Planning a Virtual IP Network Implementation

Consider the guidelines in this section when planning your virtual IP network implementation.

- ◆ [Section F.3.1, “Routing Protocol,” on page 172](#)
- ◆ [Section F.3.2, “LAN Routers,” on page 172](#)
- ◆ [Section F.3.3, “Internal Router,” on page 172](#)
- ◆ [Section F.3.4, “IP Addresses for BCC-Enabled Cluster Resources,” on page 172](#)
- ◆ [Section F.3.5, “Host Mask,” on page 173](#)

F.3.1 Routing Protocol

In theory, any state-of-the-art routing protocol could be used for the virtual IP network. This section describes how to set up the virtual router using the OSPF (Open Shortest Path First) routing protocol because it is a commonly used protocol.

For the OSPF routing protocol, define an OSPF Area ID to use for the BCC-enabled resources in a business continuity cluster. All nodes of every peer cluster in a given business continuity cluster need to use the same OSPF Area ID. When deploying multiple business continuity clusters, use a different OSPF Area ID for each one.

IMPORTANT: Do not use OSPF Area ID 0 for any of your business continuity clusters.

F.3.2 LAN Routers

For your LAN routers, you must define the OSPF Area ID to be used for each of your business continuity clusters. For guidelines about OSPF Area IDs, see [Section F.3.1, “Routing Protocol,” on page 172](#). The LAN routers are also where you define and handle the propagation of the routes to services that are using virtual IP addresses.

F.3.3 Internal Router

Do not use a cluster node as a general purpose router. Ensure that a default gateway is properly set up, and the default route of the cluster node is configured accordingly.

Redistribute any routing changes for the cluster resources to your routing infrastructure, but do not redistribute routing changes somewhere in the LAN/WAN to the cluster nodes.

F.3.4 IP Addresses for BCC-Enabled Cluster Resources

Define an IP address range for your BCC-enabled resources. The IP addresses must be unique across all peer clusters. When deploying multiple business continuity clusters, use a different IP address range for each of them.

For example, you can specify an IP address range for each peer cluster that accommodates the nodes and the master IP address of the cluster:

Cluster 1: 10.10.10.0/24
Cluster 2: 10.10.20.0/24
Cluster 3: 10.10.30.0/24
Cluster 4: 10.10.40.0/24

You can specify a different range to be used for all cluster resources in the business continuity cluster:

Resources: 10.10.50.nnn/32

F.3.5 Host Mask

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.

F.4 Configuring a Virtual Router with OSPF

Perform the following procedure for each node in every peer cluster:

- 1 Use YaST to add a Dummy Device to the node.

This creates a virtual network device named `dummy0`.

- 2 In YaST, use the Software Installer to install the Quagga Routing Suite package.
- 3 Edit the configuration files for the Quagga Routing Suite.

For information, see the [Quagga Routing Suite Documentation Web site \(http://www.nongnu.org/quagga/docs.html\)](http://www.nongnu.org/quagga/docs.html).

- 3a** In a text editor, configure `/etc/quagga/zebra.conf` to specify values for the following fields:

- ♦ Hostname
- ♦ Password (optional)
- ♦ Physical and virtual interface
- ♦ Log file (optional)

- 3b** In a text editor, configure `/etc/quagga/ospfd.conf` to specify values for the following fields:

- ♦ Hostname
- ♦ Password (optional)
- ♦ Physical and virtual interface
- ♦ OSPF area ID
- ♦ Log file (optional)

- 3c** In a text editor, configure `/etc/services`.

F.5 Configuring Virtual IP Addresses

After the appropriate virtual IP addresses and host masks have been determined (as described in [Section F.3, “Planning a Virtual IP Network Implementation,” on page 172](#)), you can enable virtual IP addresses in a business continuity cluster. A maximum of 256 virtual IP addresses can be bound. For every cluster resource in the business continuity cluster, you must modify the load and unload scripts to use the virtual IP address that you assign to that cluster resource.

For each cluster resource in the business continuity cluster, do the following:

- 1 Before you begin, the routers in a virtual IP address configuration must be running the routing protocol.

For information, see [Section F.4, “Configuring a Virtual Router with OSPF,” on page 173](#).

- 2 In the cluster resource load script, add the command to bind a virtual IP address for the cluster resource.

In iManager, use the Clusters plug-in to edit the resource load script to comment out the `add_secondary_ipaddress` line and to add the virtual IP address information for the `dummy0` adapter.

For example, the old and new lines are emphasized in the following sample load script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
exit_on_error nss /poolact=POOL1
exit_on_error ncpcon mount TVOL1=101

###exit_on_error add_secondary_ipaddress 10.10.0.101

exit_on_error ip addr add 10.50.0.101/32 dev dummy0

exit_on_error ncpcon bind --ncpservname=CL1-POOL1-VS --ipaddress=10.50.0.101
exit 0
```

- 3 In the cluster resource unload script, add the command to unbind the virtual IP address to the cluster resource unload script.

In iManager, use the Clusters plug-in to edit the resource unload script to comment out the `del_secondary_ipaddress` line and to add a line that deletes the virtual IP address information for the `dummy0` adapter.

For example, the old and new lines are emphasized in the following sample unload script:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
ignore_error ncpcon unbind --ncpservname=CL1-POOL1-VS --
ipaddress=10.50.0.101

###ignore_error del_secondary_ipaddress 10.10.0.101

ignore_error ip addr del 10.50.0.101/32 dev dummy0

ignore_error nss /pooldeact=POOL1
exit 0
```

- 4 In iManager, use the Clusters plug-in to offline the cluster resource, then online the cluster resource.

This activates the changes that you made to the cluster resource's load and unload scripts.

- 5** To verify that a virtual IP address is bound, enter `display secondary ipaddress` at a terminal console of the cluster node where the virtual IP address is assigned.
This displays all bound virtual IP addresses.
- 6** Repeat the process for the remaining cluster resources.

G

Upgrading to Identity Manager 4.7

Business Continuity Clustering supports using Identity Manager 4.7 on Open Enterprise Server (OES) 2018 Support Pack 1 (SP1).

IMPORTANT: For information about installing or upgrading Identity Manager 4.7, see the [NetIQ Identity Manager Setup Guide for Linux](#).

The following sections contain information about upgrading to Identity Manager 4.7 in an existing BCC environment:

- ♦ [Section G.1, “Requirements for Using Identity Manager,” on page 177](#)
- ♦ [Section G.2, “Upgrading to Identity Manager 4.7,” on page 177](#)
- ♦ [Section G.3, “Patching IDM 4.7 for OES 2018 SP1,” on page 178](#)

G.1 Requirements for Using Identity Manager

For information about the requirements for using Identity Manager with BCC, see [Section 4.9, “Identity Manager 4.7,” on page 43](#).

G.2 Upgrading to Identity Manager 4.7

The Identity Manager 4.7 Bundle Edition is available for OES 2018 SP1. Ensure that the Identity Manager 4.5.6 is installed and applied with the latest patches before you upgrade to version 4.7.

Repeat the following steps for the Identity Manager node in each peer cluster of the business continuity cluster:

- 1 Before you upgrade to Identity Manager 4.7, save the BCC driver configuration information, then stop the BCC drivers.
- 2 Upgrade the Identity Manager node to OES 2018 SP1. For information, see [Upgrading to OES 2018 SP1](#) in the [OES 2018 SP1: Installation Guide](#).
- 3 Upgrade the Identity Manager 4.5.6 to Identity Manager 4.7. For information, see [Section 5.2.2, “Installing and Configuring the Identity Manager Components Used by BCC,” on page 51](#).
- 4 Install the IDM 4.7 plug-ins using iManager. For information, see [Section 5.2.3, “Installing the Identity Manager Plug-in for iManager,” on page 54](#).
- 5 Start the BCC driver in Identity Manager. If BCC driver does not come up, ensure that the eDirectory is up and running on the Identity Manager node.

G.3 Patching IDM 4.7 for OES 2018 SP1

After you install or upgrade to Identity Manager 4.7, all maintenance patches for Identity Manager 4.7 must be applied on the IDM node in each of the peer clusters before you install or upgrade BCC on any of the nodes.

You can use the [Micro Focus Patch Finder \(http://download.novell.com/patch/finder/#familyId=7365\)](http://download.novell.com/patch/finder/#familyId=7365) to check for the Identity Manager 4.7 patches.



Upgrading to BCC 2.6 on OES 2018 SP1

You can use a rolling cluster update approach to upgrade from Business Continuity Clustering (BCC) 2.x to BCC 2.6 on your fully patched Open Enterprise Server (OES) 2018 Support Pack 1 (SP1) clusters.

- ◆ [Section H.1, “System Requirements for the BCC,” on page 179](#)
- ◆ [Section H.2, “Upgrading to BCC 2.6,” on page 179](#)

H.1 System Requirements for the BCC

In addition to the requirements in [Chapter 4, “Installation Requirements for BCC,” on page 35](#), do the following when you upgrade from an earlier release of BCC:

- ◆ Apply all patches for the existing OES versions (OES 11 SP2, OES 11 SP3, OES 2015, or OES 2015 SP1) on each node in every peer cluster.
- ◆ Apply all patches for Identity Manager on the IDM node in every peer cluster before you upgrade the operating systems to OES 2018 SP1 on any of the nodes. For information about upgrading the Identity Manager node, see [Appendix G, “Upgrading to Identity Manager 4.7,” on page 177](#).
- ◆ Apply all patches for OES 2018 SP1 on each node in every peer cluster before installing or upgrading to BCC 2.6.

H.2 Upgrading to BCC 2.6

When you upgrade the server to OES 2018 SP1, BCC is automatically upgraded to BCC 2.6 version. For more information, see [Upgrading to OES 2018 SP1](#) in the [OES 2018 SP1: Installation Guide](#).

- ◆ If **Novell Business Continuity Cluster** pattern is selected while upgrading the older OES versions to OES 2018 SP1, the BCC is automatically upgraded to BCC 2.6.

For example, if you are upgrading BCC 2.1 on OES 11 SP2 to BCC 2.6 on OES 2018 SP1, on the **Installation Settings** page, the following are listed under the **Update Options**:

- ◆ Product SUSE Linux Enterprise Server 11 SP3 will be updated to Open Enterprise Server 2018 SP1
- ◆ Product Novell Business Continuity Cluster 2.1 will be updated to Open Enterprise Server 2018 SP1
- ◆ Product Novell Open Enterprise Server 11 SP2 will be updated to Open Enterprise Server 2018 SP1
- ◆ If **Novell Business Continuity Cluster** pattern is not selected while upgrading the older OES versions to OES 2018 SP1, complete the OES upgrade and then manually install and configure the BCC. For more information on installing and configuring BCC, see [Section 5.4, “Installing and Configuring the Business Continuity Clustering Software,” on page 59](#).

Verify that BCC is working as expected by checking the BCC connection and resource status with iManager on every peer cluster.

Documentation Updates

This section contains information about documentation content changes made to the *Business Continuity Clustering Administration Guide* since the initial 2.0 release.

This document was updated on the following dates:

- [Section I.1, “January, 2019,” on page 181](#)
- [Section I.2, “September 21, 2013,” on page 181](#)

I.1 January, 2019

Modified eDirectory 9.1 links to point to the [NetIQ eDirectory 9.1 Web site \(http://www.netiq.com/documentation/edirectory-91/\)](http://www.netiq.com/documentation/edirectory-91/), iManager 3.1 links to point to the [NetIQ iManager 3.1 Website \(https://www.netiq.com/documentation/imanager-31/\)](https://www.netiq.com/documentation/imanager-31/), and Identity Manager 4.7 links to the [NetIQ Identity Manager 4.7 Website \(https://www.netiq.com/documentation/identity-manager-47/\)](https://www.netiq.com/documentation/identity-manager-47/).

Updates were made to the following sections. The changes are explained below.

Location	Change
Section 4.4, “Open Enterprise Server 2018 SP1,” on page 37	Updated this chapter to reflect the BCC 2.6 changes on OES 2018 SP1.
Section 4.9, “Identity Manager 4.7,” on page 43	Updated this chapter to reflect the changes in downloading and installing the Identity Manager 4.7.
Section 5.2.2, “Installing and Configuring the Identity Manager Components Used by BCC,” on page 51	Updated this section to reflect the Identity Manager 4.7 install and configure changes.
Section 5.4, “Installing and Configuring the Business Continuity Clustering Software,” on page 59	Updated this section to reflect the BCC 2.6 install and configure changes on OES 2018 SP1.
Section 5.5, “Installing the BCC Cluster Resource Template,” on page 62	Updated this section to reflect the BCC 2.6 changes on OES 2018 SP1.
Section 10.2.1, “Incorrect Status in the Cluster View Report After Enabling BCC on Peer Clusters,” on page 118	Added this issue in troubleshooting section.
Appendix G, “Upgrading to Identity Manager 4.7,” on page 177	Updated this chapter to reflect the latest Identity Manager 4.7 upgrade changes.
Appendix H, “Upgrading to BCC 2.6 on OES 2018 SP1,” on page 179	Updated this chapter to reflect the latest BCC 2.6 upgrade changes on OES 2018 SP1.

I.2 September 21, 2013

Modified Novell eDirectory 8.8 SP7 links to point to the [NetIQ eDirectory 8.8 SP7 Web site \(http://www.netiq.com/documentation/edir887/\)](http://www.netiq.com/documentation/edir887/).

Updates were made to the following sections. The changes are explained below.

Location	Change
"NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)"	This section was removed. It is not an issue in OES 11 and later.
Section 8.10, "Renaming a BCC-Enabled Pool or Volume," on page 103	This section is new.
