

Administration Guide for Novell Open Enterprise Server 2 Support Pack 1 for Linux

Novell® Business Continuity Clustering

1.2

September 22, 2010

www.novell.com



Legal Notices

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

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

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

Copyright © 2008–2010 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

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

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide	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	12
1.2.3 Stretch Clusters vs. Cluster of Clusters	13
1.3 Business Continuity Clustering	18
1.4 BCC Deployment Scenarios	19
1.4.1 Two-Site Business Continuity Cluster Solution	19
1.4.2 Multiple-Site Business Continuity Cluster Solution	20
1.4.3 Low-Cost Business Continuity Cluster Solution	21
1.5 Key Concepts	22
1.5.1 Business Continuity Clusters	22
1.5.2 Cluster Resources	22
1.5.3 Landing Zone	22
1.5.4 BCC Drivers for Identity Manager	22
2 What's New for BCC 1.2	23
2.1 BCC 1.2.0 Patch (January 2010)	23
2.1.1 BCC Engine	23
2.1.2 BCC Resource Driver Template for Identity Manager	23
2.2 Identity Manager 3.6.1 Support (June 2009)	24
2.3 BCC 1.2 for OES 2 SP1 Linux	24
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	29
3.3.5 IP Addresses for BCC-Enabled Cluster Resources	29
3.4 SAN Connectivity Guidelines	29
3.5 Storage Design Guidelines	30
3.6 eDirectory Design Guidelines	30
3.6.1 Object Location	30
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	31
3.6.5 Landing Zone	31
3.6.6 Naming Conventions for BCC-Enabled Resources	32
3.7 Cluster Design Guidelines	32

4	Installing Business Continuity Clustering	35
4.1	Requirements for BCC 1.2 for OES 2 SP1 Linux	35
4.1.1	Business Continuity Clustering License	36
4.1.2	Business Continuity Cluster Component Locations	36
4.1.3	OES 2 SP1 Linux	37
4.1.4	Novell Cluster Services 1.8.6 for Linux	37
4.1.5	Novell eDirectory 8.8	38
4.1.6	SLP	40
4.1.7	Identity Manager 3.6 Bundle Edition	40
4.1.8	Novell iManager 2.7.2	42
4.1.9	Storage-Related Plug-Ins for iManager 2.7.2	42
4.1.10	OpenWBEM	43
4.1.11	Shared Disk Systems	43
4.1.12	Mirroring Shared Disk Systems Between Peer Clusters	44
4.1.13	LUN Masking for Shared Devices	44
4.1.14	Link Speeds	44
4.1.15	Ports	45
4.1.16	Web Browser	45
4.2	Downloading the Business Continuity Clustering Software	46
4.3	Configuring a BCC Administrator User and Group	46
4.3.1	Creating the BCC Group and Administrative User	46
4.3.2	Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects	47
4.3.3	Adding the BCC Administrator User to the ncsgroup on Each Cluster Node	47
4.4	Installing and Configuring the Novell Business Continuity Clustering Software	48
4.4.1	Installing the Business Continuity Clustering RPMs	49
4.4.2	Configuring BCC Software	50
4.4.3	Installing the BCC Identity Manager Templates	50
4.5	Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software	51
4.5.1	Creating a YaST Auto-Configuration Profile	51
4.5.2	Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media	52
4.5.3	Installing and Configuring Business Continuity Clustering on Each Cluster Node	53
4.5.4	Removing the NFS Share from Your Server	53
4.5.5	Cleaning Up the Business Continuity Clustering Installation Source	53
4.6	What's Next	54
5	Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux	55
5.1	Preparing for Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux	55
5.2	Installing the BCC Patch on a Fully Patched OES 2 SP1 Linux BCC Cluster	56
5.3	Installing the BCC Patch Along With the OES 2 SP1 Linux Patches	56
6	Upgrading the Identity Manager Nodes to Identity Manager 3.6.1	59
6.1	Upgrading to 32-Bit Identity Manager 3.6.1	59
6.2	Upgrading to 64-Bit Identity Manager 3.6.1	59
7	Converting BCC Clusters from NetWare to Linux	61
7.1	Understanding the Conversion Process	61
7.2	Prerequisites for Converting from NetWare to Linux	61
7.3	Converting Clusters from NetWare to Linux	62
7.4	Deleting and Re-Creating the BCC Identity Manager Drivers	63
7.5	Finalizing the BCC Cluster Conversion	63

7.6	What's Next	63
8	Configuring the Identity Manager Drivers for BCC	65
8.1	Understanding the BCC Drivers	65
8.2	Prerequisites for Configuring the BCC Drivers for Identity Manager	70
8.2.1	Identity Manager	70
8.2.2	Novell eDirectory	71
8.2.3	Landing Zone Container	71
8.3	Configuring the BCC Drivers	71
8.4	Creating SSL Certificates	74
8.5	Enabling or Disabling the Synchronization of e-Mail Settings	74
8.6	Synchronizing Identity Manager Drivers	75
8.7	Preventing Synchronization Loops for Identity Manager Drivers	75
8.8	Changing the Identity Manager Synchronization Drivers	77
8.9	What's Next	78
9	Configuring BCC for Peer Clusters	79
9.1	Enabling Clusters for Business Continuity	79
9.2	Adding Peer Cluster Credentials	80
9.2.1	Using Console Commands to Add Credentials	80
9.2.2	Using iManager to Add Credentials	81
9.3	Adding Search-and-Replace Values to the Resource Replacement Script	82
9.4	Adding Storage Management Configuration Information	83
9.5	Configuring CIMOM Daemons to Bind to IP Addresses	86
9.6	Enabling Linux POSIX File Systems to Run on Secondary Clusters	86
9.7	Verifying BCC Administrator User Trustee Rights and Credentials	87
9.8	Disabling BCC for a Peer Cluster	88
9.9	What's Next	88
10	Managing a Business Continuity Cluster	89
10.1	Migrating a Cluster Resource to a Peer Cluster	89
10.1.1	Understanding BCC Resource Migration	89
10.1.2	Migrating Cluster Resources between Clusters	90
10.2	Bringing a Downed Cluster Back in Service	91
10.3	Changing Peer Cluster Credentials	91
10.4	Viewing the Current Status of a Business Continuity Cluster	92
10.4.1	Using iManager to View the Cluster Status	92
10.4.2	Using Console Commands to View the Cluster Status	92
10.5	Generating a Cluster Report	93
10.6	Resolving Business Continuity Cluster Failures	93
10.6.1	Storage-Based Mirroring Failure Types and Responses	93
10.6.2	Host-based Mirroring Failure Types and Responses	95
11	Configuring BCC for Cluster Resources	97
11.1	Requirements for Cluster Resources	97
11.1.1	LUNs for Cluster Pool Resources	97
11.1.2	Volumes for Cluster Pool Resources	97
11.1.3	Shared Disk Cluster Resources	98
11.1.4	Cluster Resources for OES 2 Linux Services	98
11.2	BCC-Enabling Cluster Resources	98

11.3	Configuring Search-and-Replace Values for an Individual Cluster Resource	99
11.4	Assigning Preferred Peer Clusters for the Resource	100
11.5	Assigning Preferred Nodes in Peer Clusters	101
11.6	Disabling BCC for a Cluster Resource	102
11.7	Changing the IP Address of a Cluster Resource	103
11.8	Deleting or Unsharing a BCC-Enabled Shared NSS Pool Resource	103
12	Troubleshooting Business Continuity Clustering	105
12.1	NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted	106
12.2	Identity Manager Plug-Ins Do Not Appear in iManager	106
12.3	Cluster Connection States	106
12.4	Driver Port Number Conflicts	107
12.5	Excluded Users	108
12.6	Security Equivalent User	108
12.7	SSL Certificates	109
12.8	Clusters Cannot Communicate	110
12.9	BCC Startup Flags	110
12.10	Identity Manager Drivers for Cluster Synchronization Do Not Start	111
12.11	Identity Manager Drivers Do Not Synchronize Objects from One Cluster to Another	111
12.12	Tracing Identity Manager Communications	112
12.13	Peer Cluster Communication Is Not Working	113
12.14	A Resource Does Not Migrate to a Peer Cluster	113
12.15	Resource Cannot Be Brought Online	113
12.16	Administration of Peer Clusters Is Not Functional	114
12.17	Resource Script Search-and-Replace Functions Do Not Work	114
12.18	Virtual NCP Server IP Addresses Won't Change	114
12.19	The IP Address, Virtual Server DN, or Pool Name Does Not Appear on the iManager Cluster Configuration Page	115
12.20	Blank Error String iManager Error Appears While Bringing a Resource Online	115
12.21	Mapping Drives in Login Scripts	115
12.22	Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable	116
12.23	BCC Error Codes	117
12.24	Clustered Pool Is Stuck in an eDirectory Synchronization State	118
13	Security Considerations	119
13.1	Security Features	119
13.2	Security Configuration	120
13.2.1	BCC Configuration Settings	120
13.2.2	Changing the NCS: BCC Settings Attributes in the BCC XML Configuration	121
13.2.3	Disabling SSL for Inter-Cluster Communication	122
13.2.4	Restricting the Network Address for Administration	124
13.3	General Security Guidelines	124
13.4	Security Information for Dependent Products	124
A	Console Commands for BCC	127
A.1	Using Console Commands	127
A.2	Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh	130

B	Setting Up Auto-Failover	131
B.1	Enabling Auto-Failover	131
B.2	Creating an Auto-Failover Policy	132
B.3	Refining the Auto-Failover Policy	132
B.4	Adding or Editing Monitor Configurations	134
C	Configuring Host-Based File System Mirroring for NSS Pools	135
C.1	Creating and Mirroring NSS Pools on Shared Storage	136
C.2	Creating NSS Volumes	138
C.3	Novell Cluster Services Configuration and Setup	138
C.4	Checking NSS Volume Mirror Status	138
D	Configuration Worksheet for the BCC Drivers for Identity Manager	141
D.1	Cluster1 to Cluster2	141
D.2	Cluster2 to Cluster1	143
E	Using Dynamic DNS with BCC 1.2	147
E.1	Requirements and Assumptions	147
E.1.1	Software Requirements	147
E.1.2	DNS Server	147
E.1.3	TSIG Keys	148
E.1.4	DNS Record Time-to-Live Values	148
E.2	Configuring the DNS Server for Dynamic DNS	148
E.2.1	Creating the TSIG Keys for DNS Server Authentication	148
E.2.2	Configuring the DNS Server with the Public Key	149
E.2.3	Configuring the DNS Server Zones	150
E.2.4	Testing the DNS Server	150
E.3	Configuring the Cluster Resources for Dynamic DNS	154
E.3.1	Modifying the BCC Load Script	155
E.3.2	Public and Private Keys	157
E.3.3	Testing the Perl Wrapper Script	158
E.4	Testing the Dynamic DNS Solution	159
F	Using Virtual IP Addresses with BCC 1.2	161
F.1	Understanding Internal Virtual IP Networks	161
F.1.1	Virtual Adapter	161
F.1.2	Host Mask	161
F.1.3	Internal Router	161
F.2	Virtual IP Address Benefits	162
F.2.1	High Availability	162
F.2.2	Unlimited Mobility	163
F.2.3	Automatic Name Resolution	163
F.3	Planning a Virtual IP Network Implementation	163
F.3.1	Routing Protocol	164
F.3.2	LAN Routers	164
F.3.3	Internal Router	164
F.3.4	IP Addresses for BCC-Enabled Cluster Resources	164
F.3.5	Host Mask	165
F.4	Configuring a Virtual Router with OSPF	165
F.5	Configuring Virtual IP Addresses	166

G	Removing Business Continuity Clustering Core Software	169
H	Documentation Updates	171
H.1	February 18, 2010	171
H.1.1	Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux	171
H.2	February 7, 2010 (BCC 1.2.0 Patch for OES 2 SP1 Linux)	171
H.2.1	Updating (Patching) BCC 1.2 on OES 2 SP1 Linux	171
H.2.2	What's New for BCC 1.2	172
H.3	January 7, 2010	172
H.3.1	Configuring BCC for Peer Clusters	172
H.3.2	Console Commands for BCC	172
H.3.3	Installing Business Continuity Clustering	173
H.3.4	Managing a Business Continuity Cluster	173
H.3.5	Removing Business Continuity Clustering Core Software	173
H.3.6	Troubleshooting Business Continuity Clustering	173
H.3.7	Upgrading Identity Manager Nodes to Identity Manager 3.6.1	173
H.3.8	Using Virtual IP Addresses with BCC 1.2	174
H.4	September 9, 2009	174
H.4.1	Configuring BCC for Cluster Resources	174
H.4.2	Converting BCC Clusters from NetWare to Linux	174
H.5	August 14, 2009	174
H.5.1	Console Commands for BCC	175
H.5.2	Installing Business Continuity Clustering	175
H.5.3	Upgrading Identity Manager Nodes to Identity Manager 3.6.1	175
H.5.4	What's New for BCC 1.2	175
H.6	May 11, 2009	175
H.6.1	Converting BCC Clusters from NetWare to Linux	175
H.7	April 28, 2009	176
H.7.1	Converting BCC Clusters from NetWare to Linux	176

About This Guide

This guide describes how to install, configure, and manage Novell® Business Continuity Clustering 1.2 for Novell Open Enterprise Server (OES) 2 Support Pack 1 (SP1) for Linux servers in combination with Novell Cluster Services™ 1.8.6 for Linux clusters (the version released in OES 2 SP1 Linux).

- ◆ Chapter 1, “Overview of Business Continuity Clustering,” on page 11
- ◆ Chapter 2, “What’s New for BCC 1.2,” on page 23
- ◆ Chapter 3, “Planning a Business Continuity Cluster,” on page 27
- ◆ Chapter 4, “Installing Business Continuity Clustering,” on page 35
- ◆ Chapter 5, “Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux,” on page 55
- ◆ Chapter 6, “Upgrading the Identity Manager Nodes to Identity Manager 3.6.1,” on page 59
- ◆ Chapter 7, “Converting BCC Clusters from NetWare to Linux,” on page 61
- ◆ Chapter 8, “Configuring the Identity Manager Drivers for BCC,” on page 65
- ◆ Chapter 9, “Configuring BCC for Peer Clusters,” on page 79
- ◆ Chapter 10, “Managing a Business Continuity Cluster,” on page 89
- ◆ Chapter 11, “Configuring BCC for Cluster Resources,” on page 97
- ◆ Chapter 12, “Troubleshooting Business Continuity Clustering,” on page 105
- ◆ Chapter 13, “Security Considerations,” on page 119
- ◆ Appendix A, “Console Commands for BCC,” on page 127
- ◆ Appendix B, “Setting Up Auto-Failover,” on page 131
- ◆ Appendix C, “Configuring Host-Based File System Mirroring for NSS Pools,” on page 135
- ◆ Appendix D, “Configuration Worksheet for the BCC Drivers for Identity Manager,” on page 141
- ◆ Appendix E, “Using Dynamic DNS with BCC 1.2,” on page 147
- ◆ Appendix F, “Using Virtual IP Addresses with BCC 1.2,” on page 161
- ◆ Appendix G, “Removing Business Continuity Clustering Core Software,” on page 169
- ◆ Appendix H, “Documentation Updates,” on page 171

Audience

This guide is intended for anyone involved in installing, configuring, and managing Novell Cluster Services for Linux in combination with Novell 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, or go to www.novell.com/documentation/feedback.html (<http://www.novell.com/documentation/feedback.html>) and enter your comments there.

Documentation Updates

The latest version of this *Novell Business Continuity Clustering 1.2 Administration Guide for Linux* is available on the [Business Continuity Clustering Documentation Web site](http://www.novell.com/documentation/bcc/) (<http://www.novell.com/documentation/bcc/>) under BCC 1.2.0 for OES 2 SP1 Linux.

Additional Documentation

For information about BCC 1.2 SP1 for OES 2 SP2 Linux, see:

- ♦ *BCC 1.2 SP1 Administration Guide for OES 2 SP2 Linux* (http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html)
- ♦ *OES 2 SP2: Novell Cluster Services 1.8.7 for Linux Administration Guide* (http://www.novell.com/documentation/oes2/clus_admin_lx/data/h4hgu4hs.html)
- ♦ Identity Manager 3.6.x Documentation Web site (<http://www.novell.com/documentation/idm36/>)

For information about BCC 1.2 for OES 2 SP1 Linux, see:

- ♦ *BCC 1.2 Administration Guide for Linux* (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html)
- ♦ *OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide* (http://www.novell.com/documentation/oes2/clus_admin_lx/data/h4hgu4hs.html)

Identity Manager 3.6.x Documentation Web site (<http://www.novell.com/documentation/idm36/>)

For information about BCC 1.1 SP2 for NetWare[®] 6.5 SP8, see:

- ♦ *BCC 1.1 SP2 Administration Guide for NetWare 6.5 SP8* (http://www.novell.com/documentation/bcc/bcc11_admin_nw/data/bktitle.html)
- ♦ *NW65 SP8: Novell Cluster Services 1.8.5 Administration Guide* (http://www.novell.com/documentation/nw65/clus_admin_nw/data/h4hgu4hs.html)
- ♦ Identity Manager 3.5.1 Documentation Web site (<http://www.novell.com/documentation/idm35/>)

For information about OES 2 Linux, see the [OES 2 Documentation Web site](http://www.novell.com/documentation/oes2/) (<http://www.novell.com/documentation/oes2/>).

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

Overview of Business Continuity Clustering

1

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

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

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

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

- ◆ [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 18](#)
- ◆ [Section 1.4, “BCC Deployment Scenarios,” on page 19](#)
- ◆ [Section 1.5, “Key Concepts,” on page 22](#)

1.1 Disaster Recovery Implications

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

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

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

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

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

1.2 Disaster Recovery Implementations

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

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

1.2.1 LAN-Based versus Internet-Based Applications

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

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

1.2.2 Host-Based versus Storage-Based Data Mirroring

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

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

[Table 1-1](#) compares the benefits and limitations of host-based and storage-based mirroring.

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

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

1.2.3 Stretch Clusters vs. Cluster of Clusters

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

Novell Business Continuity Clustering automates some of the configuration and processes used in a cluster of clusters. For information, see [Section 1.3, “Business Continuity Clustering,”](#) on page 18.

- ♦ [“Stretch Clusters”](#) on page 13
- ♦ [“Cluster of Clusters”](#) on page 14
- ♦ [“Comparison of Stretch Clusters and Cluster of Clusters”](#) on page 15
- ♦ [“Evaluating Disaster Recovery Implementations for Clusters”](#) on page 17

Stretch Clusters


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

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

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

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

Figure 1-1 *Stretch Cluster*



Cluster of Clusters


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

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

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

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

Figure 1-2 Cluster of Clusters



Comparison of Stretch Clusters and Cluster of Clusters

[Table 1-2](#) compares the capabilities of a stretch cluster and a cluster of clusters.

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

Capability	Stretch Cluster	Cluster of Clusters
Number of clusters	One	Two or more
Number of geographically separated data centers	Two	Two or more
eDirectory trees	Single tree only; requires the replica ring to span data centers.	One or multiple trees

Capability	Stretch Cluster	Cluster of Clusters
eDirectory Organizational Units (OUs)	<p>Single OU container for all nodes.</p> <p>As a best practice, place the cluster container in an OU separate from the rest of the tree.</p>	<p>Each cluster can be in a different OU. Each cluster is in a single OU container.</p> <p>As a best practice, place each cluster container in an OU separate from the rest of the tree.</p>
IP subnet	<p>IP addresses for nodes and cluster resources must be in a single IP subnet.</p> <p>Because the subnet spans multiple locations, you must ensure that your switches handle gratuitous ARP (Address Resolution Protocol).</p>	<p>IP addresses in a given cluster are in a single IP subnet. Each cluster can use the same or different IP subnet.</p> <p>If you use the same subnet for all clusters in the cluster of clusters, you must ensure that your switches handle gratuitous ARP.</p>
SBD partition	<p>A single SBD is mirrored between two sites by using host-based mirroring, which limits the distance between data centers to 10 km.</p>	<p>Each cluster has its own SBD.</p> <p>Each cluster can have an on-site mirror of its SBD for high availability.</p> <p>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.</p>
Failure of the site interconnect if using host-based mirroring	<p>LUNs might become primary at both locations (split brain problem).</p>	<p>Clusters continue to function independently.</p>
Storage subsystem	<p>Each cluster accesses only the storage subsystem on its own site.</p>	<p>Each cluster accesses only the storage subsystem on its own site.</p>
<p>Data-block replication between sites</p> <p>For information about data replication solutions, see Section 1.2.2, "Host-Based versus Storage-Based Data Mirroring," on page 12.</p>	<p>Yes; typically uses storage-based mirroring, but host-based mirroring is possible for distances up to 10 km.</p>	<p>Yes; typically uses storage-based mirroring, but host-based mirroring is possible for distances up to 10 km.</p>
Clustered services	<p>A single service instance runs in the cluster.</p>	<p>Each cluster can run an instance of the service.</p>
Cluster resource failover	<p>Automatic failover to preferred nodes at the other site.</p>	<p>Manual failover to preferred nodes on one or multiple clusters (multiple-site fan-out failover).</p> <p>Failover requires additional configuration.</p>

Capability	Stretch Cluster	Cluster of Clusters
Cluster resource configurations	Configured for a single cluster	Configured for the primary cluster that hosts the resource, then the configuration is manually replicated to the peer clusters.
Cluster resource configuration synchronization	Controlled by the master node	Manual process that can be tedious and error-prone.
Failover of cluster resources between clusters	Not applicable	Manual management of the storage systems and the cluster.
Link latency between sites	Can cause false failovers. The cluster heartbeat tolerance between master and slave must be increased to as high as 30 seconds. Monitor cluster heartbeat statistics, then tune down as needed.	Each cluster functions independently in its own geographical site.

Evaluating Disaster Recovery Implementations for Clusters

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

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

	Stretch Cluster	Cluster of Clusters
Advantages	<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.

	Stretch Cluster	Cluster of Clusters
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.
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 2 Linux if the SANs are not SMI-S compliant.

1.3 Business Continuity Clustering

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

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

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

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

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

- ◆ Supports up to four clusters with up to 32 nodes each.
- ◆ Integrates with shard storage hardware devices to automate the failover process through standards-based mechanisms such as SMI-S.
- ◆ Uses 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 Novell iManager management tool. iManager is used for the configuration and monitoring of the overall system and for the individual resources.

1.4 BCC Deployment Scenarios

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

- ◆ [Section 1.4.1, “Two-Site Business Continuity Cluster Solution,” on page 19](#)
- ◆ [Section 1.4.2, “Multiple-Site Business Continuity Cluster Solution,” on page 20](#)
- ◆ [Section 1.4.3, “Low-Cost Business Continuity Cluster Solution,” on page 21](#)

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 either with iSCSI-based block replication, or host-based mirroring. In either case, snapshot technology can allow for asynchronous replication over long distances. However, the low-cost solution does not necessarily have the performance associated with higher-end Fibre Channel storage arrays.

1.5 Key Concepts

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

- ♦ [Section 1.5.1, “Business Continuity Clusters,” on page 22](#)
- ♦ [Section 1.5.2, “Cluster Resources,” on page 22](#)
- ♦ [Section 1.5.3, “Landing Zone,” on page 22](#)
- ♦ [Section 1.5.4, “BCC Drivers for Identity Manager,” on page 22](#)

1.5.1 Business Continuity Clusters

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

1.5.2 Cluster Resources

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

1.5.3 Landing Zone

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

1.5.4 BCC Drivers for Identity Manager

Business Continuity Clustering requires a special Identity Manager driver that uses an Identity Vault to synchronize the cluster resource configuration information between the peer clusters. If the peer clusters are in different eDirectory trees, an additional BCC driver helps synchronize user information between the trees. For information, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,” on page 65](#).

What's New for BCC 1.2

2

This section describes the changes and enhancements that were made to Novell® Business Continuity Clustering (BCC) 1.2 for Novell Open Enterprise Server (OES) 2 Support Pack 1 (SP1) since the initial release of BCC 1.2.

- ♦ [Section 2.1, “BCC 1.2.0 Patch \(January 2010\),” on page 23](#)
- ♦ [Section 2.2, “Identity Manager 3.6.1 Support \(June 2009\),” on page 24](#)
- ♦ [Section 2.3, “BCC 1.2 for OES 2 SP1 Linux,” on page 24](#)

2.1 BCC 1.2.0 Patch (January 2010)

In January 2010, a BCC 1.2.0 patch is available through the OES 2 SP1 Linux patch channel (oes2sp1-January-2010-Scheduled-Maintenance-6749). For information about applying the patch, see [Chapter 5, “Updating \(Patching\) BCC 1.2.0 on OES 2 SP1 Linux,” on page 55](#).

The major changes for BCC 1.2.0 are described in the following sections:

- ♦ [Section 2.1.1, “BCC Engine,” on page 23](#)
- ♦ [Section 2.1.2, “BCC Resource Driver Template for Identity Manager,” on page 23](#)

2.1.1 BCC Engine

The BCC 1.2.0 patch includes the following major bug fixes for the BCC engine:

- ♦ Improves the update process to wait for the `adminfsd` and `bccd` daemons to gracefully stop before running the BCC install scripts. (Bug 561055)
 - ♦ Modified the `novell-bcc init.d` script to wait up to 15 seconds when `bccd` is stopped. Typically, the wait is less than 10 seconds.
 - ♦ Modified the post-install script of the Novell BCC specification file to wait up to 5 seconds when `adminfsd` is stopped. Typically, the wait is about 1 second.
- ♦ Improves memory management functions that might cause the `bccd` daemon to die. Overall, the code was simplified and clarified so that the shared memory functions now do exactly what their names describe. The `bccd` daemon was modified to generate unique keys to verify that the keys are not in use by other processes, then to use the verified unique keys for its processing threads. (Bug 553527)
- ♦ Improves the detection and handling of `No Memory` exceptions. In addition, if an exception is not caught and handled where it occurs, the engine's main thread detects the exception and gracefully shuts itself down. (Bug 428161)

2.1.2 BCC Resource Driver Template for Identity Manager

The BCC 1.2.0 patch for OES 2 SP1 Linux includes a new BCC resource driver template for Identity Manager that offers the following new feature and bug fixes:

- ♦ Uses a newer policy linking format so that you are no longer prompted to update the driver in iManager. (New)

- ◆ Adds the host resource name and Novell Distributed File Services (DFS) GUID attributes to the Volume objects that are synchronized for a BCC-enabled volume resource. (Bug 535127)
- ◆ No longer creates duplicate NCP Server, Volume, and Pool objects when the landing zone is not the same location as the cluster server's container. (Found while debugging Bug 537981)
- ◆ Ensures that a volume resource's link to the virtual NCP Server object is updated to point to the cluster where the resource is mounted. (Found while debugging Bug 537981)
- ◆ The event for the `IsClusterEnabled` policy in a BCC resource driver now allows a resource's peer list to be synchronized to the peer clusters, even if the cluster is disabled, if the current cluster's name is being removed from the peer list. Only this specific change is allowed; other changes to a resource are dropped (vetoed) by the driver after a cluster is disabled. (Bug 434243)

The new BCC resource driver template is compatible with the following combinations of Identity Manager and operating systems:

Identity Manager	Operating System
Identity Manager 3.6 (32-bit)	OES 2 SP1 Linux (32-bit)
Identity Manager 3.6.1 (32-bit or 64-bit)	OES 2 SP1 Linux (32-bit or 64-bit)
Identity Manager 3.6.1 (32-bit or 64-bit)	OES 2 SP2 Linux (32-bit or 64-bit)
Identity Manager 3.5.x	NetWare® 6.5 SP8

The new BCC resource driver template is not automatically applied to existing drivers. You can continue to use your existing BCC resource drivers, or you can re-create the BCC resource drivers with the new template in order to take advantage of the changes it offers. We recommend that you re-create the drivers with the new template, but it is not required.

2.2 Identity Manager 3.6.1 Support (June 2009)

In June 2009, Identity Manager 3.6.1 was released to provide support for the 64-bit OES 2 SP1 Linux operating system. Previously, Identity Manager required a 32-bit operating system, even with 64-bit hardware. This means the Identity Manager node in a BCC peer cluster can now be installed on a 64-bit operating system. Updating to Identity Manager 3.6.1 is needed only for 64-bit support, or to take advantage of bug fixes that might be offered in 3.6.1.

For information about upgrading from Identity Manager 3.6 to Identity Manager 3.6.1 in a BCC environment, see [Chapter 6, “Upgrading the Identity Manager Nodes to Identity Manager 3.6.1,”](#) on page 59.

2.3 BCC 1.2 for OES 2 SP1 Linux

BCC 1.2 for OES 2 SP1 Linux provides the following enhancements and changes over BCC 1.1 SP2 for NetWare® 6.5 SP8:

- ◆ Support for OES 2 SP1 Linux
- ◆ Support for Novell Cluster Services™ 1.8.6 for Linux
- ◆ Support for Identity Manager 3.6 (32-bit). A 64-bit update is planned.
- ◆ Support for 32-bit and 64-bit architectures

- ◆ Support for Novell eDirectory™ 8.8
- ◆ Support for Novell iManager 2.7.2
- ◆ Preferred node failover between clusters
- ◆ Enterprise data center capabilities
 - ◆ Geographical failover of virtual machines as cluster resources
 - ◆ Full support for CIM management in tools (requires OpenWBEM)

Planning a Business Continuity Cluster

3

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

- ◆ [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 29](#)
- ◆ [Section 3.5, “Storage Design Guidelines,” on page 30](#)
- ◆ [Section 3.6, “eDirectory Design Guidelines,” on page 30](#)
- ◆ [Section 3.7, “Cluster Design Guidelines,” on page 32](#)

3.1 Determining Design Criteria

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

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

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

3.2 Best Practices

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

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

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

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

3.3 LAN Connectivity Guidelines

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

- ◆ [Section 3.3.1, “VLAN,” on page 28](#)
- ◆ [Section 3.3.2, “Channel Bonding,” on page 28](#)
- ◆ [Section 3.3.3, “IP Addresses,” on page 29](#)
- ◆ [Section 3.3.4, “Name Resolution,” on page 29](#)
- ◆ [Section 3.3.5, “IP Addresses for BCC-Enabled Cluster Resources,” on page 29](#)

3.3.1 VLAN

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

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

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

3.3.2 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 make sure that a node is not cast out of the cluster when a broken link is restored.

Consider connecting cluster nodes to access switches for fault tolerance.

3.3.3 IP Addresses

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

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

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

3.3.4 Name Resolution

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

3.3.5 IP Addresses for BCC-Enabled Cluster Resources

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

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

For an example of configuring a dynamic transformation by using DNS, see [Appendix E, "Using Dynamic DNS with BCC 1.2,"](#) on page 147.

3.4 SAN Connectivity Guidelines

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

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

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

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

3.5 Storage Design Guidelines

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

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

3.6 eDirectory Design Guidelines

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

- ♦ [Section 3.6.1, “Object Location,” on page 30](#)
- ♦ [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 31](#)
- ♦ [Section 3.6.5, “Landing Zone,” on page 31](#)
- ♦ [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 cluster-enabling NSS pools the default naming conventions used by NSS are:


Cluster Resource: `poolname_SERVER`
Cluster-Enabled Pool: `clustername_poolname_POOL`
Cluster-Enabled Volume: `clustername_volumename`
Virtual Server: `clustername_poolname_SERVER`

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

Cluster Resource: `poolname_SERVER`
Cluster-Enabled Pool: `BCC_poolname_POOL`
Cluster-Enabled Volume: `BCC_volumename`
Virtual Server: `BCC_poolname_SERVER`

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

Figure 3-2 Cluster Resource Identity in Two Clusters



3.7 Cluster Design Guidelines

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

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

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

When existing clusters are configured and enabled within the same business continuity cluster, the volume IDs for the existing shared volumes might also share the same volume IDs. To resolve this conflict, manually edit the load script for each volume that has been enabled for business continuity and change the volume IDs to unique values for each volume in the business continuity cluster.

- ◆ BCC configuration should consider the configuration requirements for each of the services supported across all peer clusters.
- ◆ Create failover matrixes for each cluster resource so that you know what service is supported and which nodes are the preferred nodes for failover within the same cluster and among the peer clusters.

Installing Business Continuity Clustering

4

This section describes how to install, set up, and configure Novell® Business Continuity Clustering 1.2 for Novell Open Enterprise Server (OES) 2 SP1 Linux for your specific needs.

- ◆ [Section 4.1, “Requirements for BCC 1.2 for OES 2 SP1 Linux,” on page 35](#)
- ◆ [Section 4.2, “Downloading the Business Continuity Clustering Software,” on page 46](#)
- ◆ [Section 4.3, “Configuring a BCC Administrator User and Group,” on page 46](#)
- ◆ [Section 4.4, “Installing and Configuring the Novell Business Continuity Clustering Software,” on page 48](#)
- ◆ [Section 4.5, “Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software,” on page 51](#)
- ◆ [Section 4.6, “What’s Next,” on page 54](#)

4.1 Requirements for BCC 1.2 for OES 2 SP1 Linux

The requirements in this section must be met prior to installing Novell Business Continuity Clustering 1.2 for Linux software.

- ◆ [Section 4.1.1, “Business Continuity Clustering License,” on page 36](#)
- ◆ [Section 4.1.2, “Business Continuity Cluster Component Locations,” on page 36](#)
- ◆ [Section 4.1.3, “OES 2 SP1 Linux,” on page 37](#)
- ◆ [Section 4.1.4, “Novell Cluster Services 1.8.6 for Linux,” on page 37](#)
- ◆ [Section 4.1.5, “Novell eDirectory 8.8,” on page 38](#)
- ◆ [Section 4.1.6, “SLP,” on page 40](#)
- ◆ [Section 4.1.7, “Identity Manager 3.6 Bundle Edition,” on page 40](#)
- ◆ [Section 4.1.8, “Novell iManager 2.7.2,” on page 42](#)
- ◆ [Section 4.1.9, “Storage-Related Plug-Ins for iManager 2.7.2,” on page 42](#)
- ◆ [Section 4.1.10, “OpenWBEM,” on page 43](#)
- ◆ [Section 4.1.11, “Shared Disk Systems,” on page 43](#)
- ◆ [Section 4.1.12, “Mirroring Shared Disk Systems Between Peer Clusters,” on page 44](#)
- ◆ [Section 4.1.13, “LUN Masking for Shared Devices,” on page 44](#)
- ◆ [Section 4.1.14, “Link Speeds,” on page 44](#)
- ◆ [Section 4.1.15, “Ports,” on page 45](#)
- ◆ [Section 4.1.16, “Web Browser,” on page 45](#)

4.1.1 Business Continuity Clustering License

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

4.1.2 Business Continuity Cluster Component Locations

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

Figure 4-1 Business Continuity Cluster Component Locations




Figure 4-1 uses the following abbreviations:

BCC: Novell Business Continuity Clustering 1.2 for Linux

eDir: [Novell eDirectory 8.8](#)

IDM: [Identity Manager 3.6.x](#)

iManager: [Novell iManager 2.7.2](#)

NCS: [Novell Cluster Services 1.8.6 for Linux](#)

OES Linux: [Novell Open Enterprise Server 2 SP1 for Linux](#)

4.1.3 OES 2 SP1 Linux

Novell Open Enterprise Server (OES) 2 Support Pack 1 (SP1) for Linux must be installed and running on every node in each peer cluster that will be part of the business continuity cluster.

See the *OES 2 SP1: Linux Installation Guide* (http://www.novell.com/documentation/oes2/inst_oes_lx/data/front.html) for information on installing and configuring OES 2 SP1 Linux.

IMPORTANT: If you use Identity Manager 3.6, each node in every peer cluster where Identity Manager 3.6 is installed must be running the 32-bit version of OES 2 SP1 Linux because Identity Manager 3.6 supports only 32-bit operating systems. Identity Manager 3.6.1 supports the 64-bit version of OES 2 SP1 Linux.

4.1.4 Novell Cluster Services 1.8.6 for Linux

You need two to four clusters with Novell Cluster Services™ 1.8.6 (the version that ships with OES 2 SP1) installed and running on each node in the cluster.

See the *OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide* for information on installing, configuring, and managing Novell Cluster Services.

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

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

Cluster Names

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

Storage

The storage requirements for Novell Business Continuity Clustering software are the same as for Novell Cluster Services. For more information, see the following in the *OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide*:

- ♦ “Hardware Requirements”
- ♦ “Shared Disk System Requirements”
- ♦ “LUN Masking”

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

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

eDirectory

The recommended configuration is to have each cluster in the same eDirectory tree but in different OUs (Organizational Units). BCC 1.2 for Linux supports only a single-tree setup.

Peer Cluster Credentials

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

4.1.5 Novell eDirectory 8.8

Novell eDirectory 8.8 is supported with Business Continuity Clustering 1.2. See the [eDirectory 8.8 documentation \(http://www.novell.com/documentation/edir88/\)](http://www.novell.com/documentation/edir88/) for more information.

- ♦ “eDirectory Containers for Clusters” on page 38
- ♦ “Rights Needed for Installing BCC” on page 38
- ♦ “Rights Needed for Individual Cluster Management” on page 39
- ♦ “Rights Needed for BCC Management” on page 39
- ♦ “Rights Needed for Identity Manager” on page 39

eDirectory Containers for Clusters

Each of the clusters that you want to add to a business continuity cluster should reside in its own OU level container. Each OU should reside in a different eDirectory partition.

As a best practice for each of the peer clusters, put its Server objects, Cluster object, Driver objects, and Landing Zone in a the same eDirectory partition.

For example:

ou=cluster1	ou=cluster2	ou=cluster3
ou=cluster1LandingZone	ou=cluster2LandingZone	ou=cluster3LandingZone
cn=cluster1	cn=cluster2	cn=cluster3
c1server35	c2server135	c3server235
c1server36	c2server136	c3server236
c1server37 (IDM node with read/write access to ou=cluster1)	c3server137 (IDM node with read/write access to ou=cluster2)	c3server237 (IDM node with read/write access to ou=cluster3)
cluster1BCCDrivers	cluster2BCCDrivers	cluster3BCCDrivers
♦ c1toc2BCCDriver	♦ c2toc1BCCDriver	♦ c3toc1BCCDriver
♦ c1toc3BCCDriver		

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 who installs BCC must have the eDirectory credentials necessary to extend the schema.

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

Rights Needed for Individual Cluster Management

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

Rights Needed for BCC Management

Before you install BCC, create the BCC Administrator user and group identities in eDirectory to use when you manage the BCC. For information, see [Section 4.3, “Configuring a BCC Administrator User and Group,”](#) on page 46.

Rights Needed for Identity Manager

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

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

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

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

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

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

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

4.1.6 SLP

You must have SLP (Server Location Protocol) set up and configured properly on each server node in every cluster. Typically, SLP is installed as part of the eDirectory installation and setup when you install the server operating system for the server. For information, see “[Implementing the Service Location Protocol](http://www.novell.com/documentation/edir88/edir88/data/ba51b4b.html)” (<http://www.novell.com/documentation/edir88/edir88/data/ba51b4b.html>) in the *Novell eDirectory 8.8 Administration Guide*.

4.1.7 Identity Manager 3.6 Bundle Edition

The Identity Manager 3.6 Bundle Edition is required for synchronizing the configuration of the peer clusters in your business continuity cluster. It is not involved in other BCC management operations such as migrating cluster resources within or across peer clusters.

Beginning in June 2009, Identity Manager 3.6.1 is also supported for BCC 1.2 on OES 2 SP1 Linux. For information about upgrading to 3.6.1 in an existing BCC environment, see [Chapter 6](#), “[Upgrading the Identity Manager Nodes to Identity Manager 3.6.1](#),” on page 59.

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

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

- ♦ “[Downloading the Bundle Edition](#)” on page 40
- ♦ “[Credential for Drivers](#)” on page 41
- ♦ “[Identity Manager Engine and eDirectory Driver](#)” on page 41
- ♦ “[Identity Manager Driver for eDirectory](#)” on page 41
- ♦ “[Identity Manager Management Utilities](#)” on page 41

Downloading the Bundle Edition

The bundle edition is a limited release of Identity Manager 3.6 or 3.6.1 for OES 2 SP1 Linux that allows you to use the Identity Manager software, the eDirectory driver, and the Identity Manager management tools for Novell iManager 2.7.2. BCC driver templates are applied to the eDirectory driver to create BCC-specific drivers that automatically synchronize BCC configuration information between the Identity Manager nodes in peer clusters.

To download Identity Manager, go to the [Novell Downloads Web site](http://download.novell.com/) (<http://download.novell.com/>).

IMPORTANT: Identity Manger 3.6 requires a 32-bit OES 2 SP1 Linux operating system, even on 64-bit hardware. Identity Manager 3.6.1 additionally supports the 64-bit OES 2 SP1 Linux operating system.

Credential for Drivers

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

Identity Manager Engine and eDirectory Driver

BCC requires Identity Manager 3.6 or later to run on one node in each of the clusters that belong to the business continuity cluster. (Identity Manager was formerly called DirXML[®].) Identity Manager should not be set up as clustered resource. Each Identity Manager 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.

IMPORTANT: Identity Manager 3.6 is 32-bit and is supported only on the 32-bit OES 2 SP1 Linux operating system, even on 64-bit hardware. Identity Manager 3.6.1 additionally supports the 64-bit OES 2 SPI Linux operating system.

For installation instructions, see the *Identity Manager 3.6.x Installation Guide* (http://www.novell.com/documentation/idm36/idm_install/data/front.html).

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

Identity Manager Driver for eDirectory

On the same node where you install the Identity Manager engine, install one instance of the Identity Manager driver for eDirectory.

For information about installing the Identity Manager driver for eDirectory, see *Identity Manager 3.6.x Driver for eDirectory: Implementation Guide* (<http://www.novell.com/documentation/idm36drivers/edirectory/data/bktitle.html>).

Identity Manager Management Utilities

The Identity Manager management utilities must be installed on the same server as Novell iManager. The Identity Manager utilities and iManager can be installed on a cluster node, but installing them on a non-cluster node is the recommended configuration. For information about iManager requirements for BCC, see [Section 4.1.8, “Novell iManager 2.7.2,” on page 42](#).

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

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

```
rcndsd restart
```

4.1.8 Novell iManager 2.7.2

Novell iManager 2.7.2 (the version released with OES 2 SP1 Linux) must be installed and running on a server in the eDirectory tree where you are installing Business Continuity Clustering software. You need to install the BCC plug-in, the Clusters plug-in, and the Storage Management plug-in in order to manage the BCC in iManager. As part of the install process, you must also install plug-ins for the Identity Manager role that are management templates for configuring a business continuity cluster. The templates are in the `novellbusiness-continuity-cluster-idm.rpm` module.

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

The Identity Manager management utilities must be installed on the same server as iManager. You can install iManager and the Identity Manager utilities on a cluster node, but installing them on a non-cluster node is the recommended configuration. For information about Identity Manager requirements for BCC, see [Section 4.1.7, “Identity Manager 3.6 Bundle Edition,” on page 40](#).

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

4.1.9 Storage-Related Plug-Ins for iManager 2.7.2

The Clusters plug-in (`ncsmgmt.npm`) has been updated from the release in OES 2 SP1 Linux to provide support for this release of Business Continuity Clustering. You must install the Clusters plug-in and the Storage Management plug-in (`storagemgmt.npm`).

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

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

The Novell Storage Related Plug-ins for iManager 2.7x are available as a zipped download file on the [Novell Downloads](http://download.novell.com/) (<http://download.novell.com/>) Web site.

To install or upgrade the Clusters plug-in:

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

This step is necessary for upgrades only if you did not uninstall and reinstall the storage-related plug-ins as part of the upgrade process.

- 2 Copy the new `.npm` files into the iManager plug-ins location, manually overwriting the older version of the plug-in in the packages folder with the newer version of the plug-in.
- 3 In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.

- 4 Restart Tomcat by entering the following command at a terminal console prompt:

```
rcnovell-tomcat5 restart
```

- 5 Restart Apache by entering the following command at a terminal console prompt.

```
rcapache2 restart
```

4.1.10 OpenWBEM

OpenWBEM must be running and configured to start using `chkconfig`. For information, see the *OES 2: OpenWBEM Services Administration Guide* (http://www.novell.com/documentation/oes2/mgmt_openwbem_lx_nw/data/front.html).

The CIMOM daemons on all nodes in the business continuity cluster must be configured to bind to all IP addresses on the server. For information, see [Section 9.5, “Configuring CIMOM Daemons to Bind to IP Addresses,”](#) on page 86.

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

Beginning in OES 2, the Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a username and password) between iManager and the `_admin` volume on the OES 2 server that you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

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

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

To check the status of WBEM:

- 1 Log in as the `root` user in a terminal console, then enter

```
rcowcimomd status
```

To start WBEM:

- 1 Log in as the `root` user in a terminal console, then enter

```
rcowcimomd start
```

4.1.11 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 System Requirements](#)” in the *OES 2 SPI: Novell Cluster Services 1.8.6 for Linux Administration Guide*.

In addition to the shared disks in an original cluster, you need additional shared disk storage in the other peer clusters to mirror the data between sites as described in [Section 4.1.12, “Mirroring Shared Disk Systems Between Peer Clusters,”](#) on page 44.

4.1.12 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 C, “Configuring Host-Based File System Mirroring for NSS Pools,”](#) on page 135.

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

4.1.13 LUN Masking for Shared Devices

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

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

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

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

4.1.14 Link Speeds

For real-time mirroring, link latency is the essential consideration. For best performance, the link speeds should be at least 1 GB per second, and the links should be dedicated.

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

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

4.1.15 Ports

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

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

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

4.1.16 Web Browser

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

- ◆ “Web Browser Language Setting” on page 45
- ◆ “Web Browser Character Encoding Setting” on page 45

Web Browser Language Setting

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

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

Web Browser Character Encoding Setting

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

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

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

4.2 Downloading the Business Continuity Clustering Software

Before you install Novell Business Continuity Clustering, download and copy the software to a directory on your workstation. To download Novell Business Continuity Clustering 1.2 for Linux, go to the [Novell Downloads Web site \(http://download.novell.com\)](http://download.novell.com) and select the Business Continuity Clustering product.

The Business Continuity Clustering installation program and software is downloaded as an ISO image. There are a few installation options for using the ISO image on each Linux server that will be part of the business continuity cluster:

- ♦ Create a CD from the ISO, mount the CD on the server, and add the CD as a local installation source.
- ♦ Copy the ISO to the server and add the ISO as a local installation source.
- ♦ Copy the ISO file to a location where it can be used for an HTTP installation source.

To use one of these package installation methods, follow the instructions in [Section 4.4, “Installing and Configuring the Novell Business Continuity Clustering Software,”](#) on page 48.

You can also follow the instructions in [Section 4.5, “Using a YaST Auto-Configuration File to Install and Configure Business Continuity Clustering Software,”](#) on page 51 to install from the network without using a CD or copying the ISO to each server.

4.3 Configuring a BCC Administrator User and Group

During the install, 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 4.3.1, “Creating the BCC Group and Administrative User,”](#) on page 46
- ♦ [Section 4.3.2, “Assigning Trustee Rights for the BCC Administrator User to the Cluster Objects,”](#) on page 47
- ♦ [Section 4.3.3, “Adding the BCC Administrator User to the ncsgroup on Each Cluster Node,”](#) on page 47

4.3.1 Creating the BCC Group and Administrative User

For Linux, ensure that the BCC Administrator user is a Linux-enabled user by creating a BCC Group object (`bccgroup`), adding the BCC Administrator user to the group, then Linux-enabling the group by configuring Linux User Management (LUM) for the group. You also add the Linux nodes (Node objects) of each node in every cluster in the BCC to this BCC group.

IMPORTANT: Having a LUM-enabled BCC group and user and adding the Linux cluster nodes to the BCC group is necessary for inter-cluster communication to function properly.

Prior to installing and configuring Business Continuity Clustering software, you must complete the following tasks:

- 1 Create a BCC group and name the group, such as `bccgroup`.

IMPORTANT: The name you specify must be in all lowercase.

- 2 Create a BCC Administrator user (`bccadmin`) and add that user to the BCC group (`bccgroup`) before LUM enabling the group.
- 3 Enable the `bccgroup` for Linux by using Linux User Management.

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

- ♦ Select *LUM enable all users in group* option.
- ♦ Add all Linux nodes (Node objects) in the cluster to the `bccgroup`.

For information about LUM-enabling groups, see “[Managing User and Group Objects in eDirectory](#)” in the *OES 2 SPI: Novell Linux User Management Technology Guide*.

LUM-enabling the `bccgroup` automatically enables all users in that group for Linux.

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

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

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

4.3.3 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.

4.4 Installing and Configuring the Novell Business Continuity Clustering Software

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

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

The Business Continuity Clustering for Linux installation installs to only one server at a time. You must install it on each node of each cluster that you want to be part of a business continuity cluster.

- ◆ Install the BCC-specific Identity Manager templates for iManager on an OES 2 Linux server where you have installed iManager.

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

- ◆ Uninstall Business Continuity Clustering software.

IMPORTANT: Before you begin, make sure your setup meets the requirements specified in [Section 4.1, “Requirements for BCC 1.2 for OES 2 SP1 Linux,” on page 35](#). The BCC Administrator user and group must already be configured as specified in [Section 4.3, “Configuring a BCC Administrator User and Group,” on page 46](#).

You must 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.

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

- ♦ [Section 4.4.1, “Installing the Business Continuity Clustering RPMs,” on page 49](#)
- ♦ [Section 4.4.2, “Configuring BCC Software,” on page 50](#)
- ♦ [Section 4.4.3, “Installing the BCC Identity Manager Templates,” on page 50](#)

4.4.1 Installing the Business Continuity Clustering RPMs

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

- 1** Log in to the server as the `root` user.
- 2** Set up the Business Continuity Clustering ISO file as an installation source.
 - 2a** Copy the Business Continuity Clustering ISO file that you downloaded in [Section 4.2, “Downloading the Business Continuity Clustering Software,” on page 46](#) to a local directory.

If you use a CD or the HTTP method, modify the following steps as needed.
 - 2b** Use one of the following methods to open the Add-On Product page.
 - ♦ In YaST, select *Software > Installation Source*, then click *Add*.
 - ♦ At a terminal console prompt, enter

```
yast2 add-on
```
 - 2c** On the Add-On Product Media page, select *Local*, then click *Next*.
 - 2d** Select *ISO Image*, then browse to locate and select the file, click *Open*, then click *Next*.
 - 2e** Read the Business Continuity Clustering License Agreement, click *Yes* to accept it, then click *Next*.

When the installation source is added, it appears in the list on the Add-On Product Media page.
 - 2f** On the Add-Product Media page, click *Finish*.

The new installation source is synchronized with the Software Updater. The *Software Updater* icon in the Notification area indicates that updates are ready.
- 3** In YaST, select *Software > Software Manager*, search for BCC, select the three `.rpm` files, then click *Accept* to install the packages.

```
novell-business-continuity-cluster.rpm
novell-business-continuity-cluster-idm.rpm
yast-novell-bcc.rpm
```

You can also double-click the Software Updater icon in the Notification area, then select the BCC `.rpm` files.
- 4** After the packages are installed, exit YaST.
- 5** Continue with [Section 4.4.2, “Configuring BCC Software,” on page 50](#).

4.4.2 Configuring BCC Software

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

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

```
yast2 novell-bcc
```
- 3 When prompted to *Install Core Business Continuity Clustering Software* and *Configure Core Software*, click *Yes* to install and configure the BCC software.
- 4 Click *Continue* if you are prompted to configure LDAP.
- 5 In the eDirectory Tree dialog box, specify the eDirectory Administrator user password that you used when you installed the operating system, then click *OK*.
- 6 Specify the typeful fully distinguished eDirectory name for the cluster where the server is currently a member.
- 7 Select *Start BCC services now* to start the software immediately following the configuration process.

If you deselect the check box, you must manually start BCC services later by entering the following at the terminal console prompt:

```
rcnovell-bcc start
```
- 8 Specify the *Directory Server Address* by selecting the IP addresses of the master eDirectory server and the local server.
- 9 Accept or change the eDirectory Administrator user name and specify the Administrator user's password.
- 10 Click *Next*.
- 11 Review your setup on the Novell Business Continuity Clustering Configuration Summary page, then click *Next* to install the BCC software.
- 12 Click *Finish* to save the BCC configuration and exit the tool.
- 13 Verify that the BCC software is running on the server by entering the following at a terminal console prompt:

```
rcnovell-bcc status
```
- 14 Continue with [Section 4.4.3, "Installing the BCC Identity Manager Templates,"](#) on page 50.

4.4.3 Installing the BCC Identity Manager Templates

The Identity Manager management templates for Business Continuity Clustering consist of XML templates for the eDirectory driver that are used for synchronizing objects between peer clusters. The templates are required to configure your business continuity cluster. You must have previously installed iManager on the server where you plan to install the templates.

You should install the management templates on the same server where you installed the Identity Manager Management utilities. For information, see [Section 4.1.7, "Identity Manager 3.6 Bundle Edition,"](#) on page 40.

IMPORTANT: If you have not already done so, install the BCC RPMs on this server as described in [Section 4.4.1, “Installing the Business Continuity Clustering RPMs,”](#) on page 49.

Perform the following tasks on the iManager server in each tree that belongs to the business continuity cluster:

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

```
yast2 novell-bcc
```
- 3 When prompted, deselect the *Install Core Business Continuity Clustering Software* and *Configure Core Software* option and select the *Install Identity Manager Templates* option, then click *Next*.

Selecting the *Install Identity Manager Templates* option installs the iManager plug-ins on this Linux server.

4.5 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 4.5.1, “Creating a YaST Auto-Configuration Profile,”](#) on page 51
- ♦ [Section 4.5.2, “Setting Up an NFS Server to Host the Business Continuity Clustering Installation Media,”](#) on page 52
- ♦ [Section 4.5.3, “Installing and Configuring Business Continuity Clustering on Each Cluster Node,”](#) on page 53
- ♦ [Section 4.5.4, “Removing the NFS Share from Your Server,”](#) on page 53
- ♦ [Section 4.5.5, “Cleaning Up the Business Continuity Clustering Installation Source,”](#) on page 53

4.5.1 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`.

4.5.2 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 you created 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, you could enter commands similar to the following:

```
mkdir /tmp/bcc_install
```

```
cp -r /media/dvdrecorder /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`. Then continue with [Step 3](#) below.
- 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, click *OK*, then click *Finish*.

4.5.3 Installing and Configuring Business Continuity Clustering on Each Cluster Node

You must install BCC 1.2 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 52](#), click *OK*, then *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 51](#).

- 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*.

4.5.4 Removing the NFS Share from Your Server

You can optionally remove the Business Continuity Clustering installation directory as an NFS share. After a successful install, it is needed only if you re-install 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*.

4.5.5 Cleaning Up the Business Continuity Clustering Installation Source

Clean up the Business Continuity Clustering installation source by opening a terminal console and running one of the commands below, depending on which method you chose in [Step 1 on page 52](#).

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

or

```
umount /mnt/iso
```

4.6 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 8, “Configuring the Identity Manager Drivers for BCC,” on page 65](#)
- ♦ If you are adding a new cluster to an existing business continuity cluster, follow the instructions in [“Synchronizing Identity Manager Drivers” on page 75](#) to synchronize the BCC Identity Manager drivers across all the clusters.
- ♦ [Chapter 9, “Configuring BCC for Peer Clusters,” on page 79](#)
- ♦ [Chapter 11, “Configuring BCC for Cluster Resources,” on page 97](#)

Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux

5

Beginning in January 2010, patches are available for Novell® Business Continuity Clustering (BCC) 1.2.0 in the Novell Open Enterprise Server (OES) 2 SP1 Linux patch channel.

BCC administrators can use a rolling update approach to download and install the BCC 1.2.0 patch for each node in every peer cluster in the business continuity cluster. The BCC patch can be installed on a fully patched OES 2 SP1 Linux cluster, or it can be installed when the OES 2 SP1 Linux patches are applied in the cluster. Each of these update approaches are discussed in this section.

- ♦ [Section 5.1, “Preparing for Updating \(Patching\) BCC 1.2.0 on OES 2 SP1 Linux,” on page 55](#)
- ♦ [Section 5.2, “Installing the BCC Patch on a Fully Patched OES 2 SP1 Linux BCC Cluster,” on page 56](#)
- ♦ [Section 5.3, “Installing the BCC Patch Along With the OES 2 SP1 Linux Patches,” on page 56](#)

5.1 Preparing for Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux

Updating (patching) BCC 1.2.0 for OES 2 SP1 Linux assumes that you have installed BCC 1.2 on an existing business continuity cluster. The BCC cluster environment must meet the system requirements described in [Section 4.1, “Requirements for BCC 1.2 for OES 2 SP1 Linux,” on page 35](#).

You must also prepare the cluster nodes for the OES 2 patch process as described in “[Updating \(Patching\) an OES 2 SP1 Linux Server](#)” in the *OES 2 SP1: Linux Installation Guide*. If BCC 1.2 is installed on an OES 2 SP1 Linux server, the BCC 1.2.0 patch files are pulled down from the OES 2 SP1 Linux patch channel in the maintenance patch for OES 2 SP1 Linux.

The following BCC 1.2 patches are available for OES 2 SP1 Linux:

- ♦ **OES2 SP1 January 2010 Scheduled Maintenance 20100130:** The BCC 1.2.0 patch includes the following RPM files:

```
novell-business-continuity-cluster-1.2.0.765-0.4.1.i586.rpm
```

```
novell-business-continuity-cluster-idm-1.2.0.759-0.4.1.noarch.rpm
```

The features and changes included in the BCC1.2.0 patch are described in [Section 2.1, “BCC 1.2.0 Patch \(January 2010\),” on page 23](#).

The patch includes a new [BCC resource driver template for Identity Manager](#). We recommend that you delete and re-create the BCC drivers to use the new BCC driver template, but it is not required.

5.2 Installing the BCC Patch on a Fully Patched OES 2 SP1 Linux BCC Cluster

Use the procedure in this section to apply the BCC patch to a fully patched OES 2 SP1 Linux BCC cluster. In this scenario, it is recommended, but not required, that you migrate the cluster resources to a different node before installing the BCC patch. The BCC installation process does not affect the cluster resources that are already in place. It is not necessary to reboot the server in order for the BCC code changes to be applied.

- 1 On the Identity Manager node in every peer cluster, save the BCC driver configuration information, then stop the BCC drivers.
- 2 On one peer cluster, use a rolling update approach to install the BCC 1.2.0 patch:
 - 2a (Optional) On one of the nodes in the cluster, migrate its cluster resources to another node in the cluster.
 - 2b Install the BCC 1.2.0 patch on the node (where cluster resources are not running).
It is not necessary to reboot the server in order for the BCC code changes to be applied.
 - 2c If you migrated cluster resources in [Step 2a](#), migrate them back to the updated node.
 - 2d Repeat [Step 2a](#) through [Step 2c](#) for the remaining nodes in the cluster.
- 3 Repeat [Step 2](#) on one peer cluster at a time until the BCC 1.2.0 patch has been applied to each node in every peer cluster in the business continuity cluster.
- 4 After all nodes in every peer cluster are updated, do one of the following for the BCC drivers for Identity Manager:
 - ♦ **Re-Create BCC Drivers (Recommended):** On the Identity Manager node in every peer cluster, delete the old BCC drivers, then re-create the BCC drivers with the new BCC template.
For information, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,” on page 65](#).
 - ♦ **Use the Old BCC Drivers:** On the Identity Manager node in every peer cluster, start the old BCC drivers. The new feature and bug fixes for the template will not be available.
- 5 Restart Tomcat by entering the following command at a terminal console prompt:

```
rcnovell-tomcat5 restart
```
- 6 Verify that BCC is working as expected by checking the BCC connection and resource status with iManager on every peer cluster.

5.3 Installing the BCC Patch Along With the OES 2 SP1 Linux Patches

Use the procedure in this section to apply the BCC patch along with the OES 2 SP1 Linux patches during a rolling upgrade of the cluster nodes. In this scenario, it is required that you migrate the cluster resources to a different node before installing the OES 2 SP1 Linux patches and the BCC patch. Although it is not necessary to reboot the server in order for the BCC changes to be applied, other OES 2 SP1 Linux patches might require the server to be rebooted.

- 1 On the Identity Manager node in every peer cluster, save the BCC driver configuration information, then stop the BCC drivers.

- 2 On one peer cluster, use a rolling update approach to install the BCC 1.2.0 patch:
 - 2a On the Identity Manager node in the cluster, apply the OES 2 SP1 Linux patches, then reboot the server if you are prompted to do so.
 - 2b On one of the nodes in the cluster, migrate its cluster resources to another node in the cluster.
 - 2c On the node, install the OES 2 SP1 Linux patches.
 - 2d On the node, install the BCC 1.2.0 patch.
 - 2e If you are prompted to do so, reboot the updated node.
 - 2f Migrate the cluster resources back to the updated node.
 - 2g Repeat [Step 2b](#) through [Step 2f](#) for the remaining nodes in the cluster.
- 3 Repeat [Step 2](#) on one peer cluster at a time until the OES 2 SP1 Linux patches and the BCC 1.2.0 patch have been applied to each node in every peer cluster in the business continuity cluster.
- 4 After all nodes in every peer cluster are updated, do one of the following for the BCC drivers for Identity Manager:
 - ♦ **Re-Create BCC Drivers (Recommended):** On the Identity Manager node in every peer cluster, delete the old BCC drivers, then re-create the BCC drivers with the new BCC template.
For information, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,” on page 65](#).
 - ♦ **Use the Old BCC Drivers:** On the Identity Manager node in every peer cluster, start the old BCC drivers. The new feature and bug fixes for the template will not be available.
- 5 Restart Tomcat by entering the following command at a terminal console prompt:

```
rcnovell-tomcat5 restart
```
- 6 Verify that BCC is working as expected by checking the BCC connection and resource status with iManager on every peer cluster.

Upgrading the Identity Manager Nodes to Identity Manager 3.6.1

Beginning in June 2009, Novell® Business Continuity Clustering 1.2 supports using Identity Manager 3.6.1 (32-bit and 64-bit) on Novell Open Enterprise Server (OES) 2 SP1 Linux. Updating to Identity Manager 3.6.1 is needed only for 64-bit support, or to take advantage of bug fixes that might be offered in 3.6.1. You can upgrade from Identity Manager 3.6 (which is 32-bit only) to Identity Manager 3.6.1 on 32-bit or 64-bit platforms.

IMPORTANT: For information about installing or upgrading Identity Manager 3.6.1, see the *Novell Identity Manager 3.6.1 Installation Guide* (http://www.novell.com/documentation/idm36/idm_install/data/be114ik.html).

The following sections contain information about upgrading to Identity Manager 3.6.1 in an existing BCC environment:

- ♦ [Section 6.1, “Upgrading to 32-Bit Identity Manager 3.6.1,” on page 59](#)
- ♦ [Section 6.2, “Upgrading to 64-Bit Identity Manager 3.6.1,” on page 59](#)

6.1 Upgrading to 32-Bit Identity Manager 3.6.1

On a 32-bit OES 2 SP1 Linux operating system, you can install the 32-bit version of Identity Manager 3.6.1 to automatically upgrade Identity Manager to the latest version. You must stop the BCC drivers, but it is not necessary to re-create them.

Repeat the following steps for the Identity Manager node in each peer cluster of the business continuity cluster:

- 1 Before you upgrade to 32-bit Identity Manager 3.6.1, stop the BCC drivers.
- 2 Install 3.6.1 on a 32-bit OES 2 SP1 Linux operating system to upgrade Identity Manager from 3.6 to 3.6.1.
- 3 Restart the BCC drivers in Identity Manager.

6.2 Upgrading to 64-Bit Identity Manager 3.6.1

There is no in-place upgrade of the OES 2 SP1 operating system from 32-bit to 64-bit. The upgrade from 32-bit Identity Manager 3.6 to 64-bit Identity Manager 3.6.1 requires a rebuild of the 64-bit cluster node to install the 64-bit OES 2 SP1 Linux operating system and the Identity Manager and iManager software components. You must re-create the BCC drivers.

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 3.6.1, save the BCC driver configuration information, then stop the BCC drivers.

- 2** On a 64-bit machine, reinstall the operating system with the 64-bit OES 2 SP1 Linux, then install Identity Manager 3.6.1 and iManager 2.7.2 on the system as described in [Section 4.1.7, “Identity Manager 3.6 Bundle Edition,”](#) on page 40.
- 3** Re-create the BCC drivers in Identity Manager.
For information about creating drivers, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,”](#) on page 65.

Converting BCC Clusters from NetWare to Linux

7

Novell® Business Continuity Clustering (BCC) 1.2 for Novell Open Enterprise Server (OES) 2 Support Pack 1 for Linux supports conversion of the clusters in a business continuity cluster from BCC 1.1 SP2 for NetWare® 6.5 SP8 (same as OES 2 SP1 NetWare).

- ♦ [Section 7.1, “Understanding the Conversion Process,” on page 61](#)
- ♦ [Section 7.2, “Prerequisites for Converting from NetWare to Linux,” on page 61](#)
- ♦ [Section 7.3, “Converting Clusters from NetWare to Linux,” on page 62](#)
- ♦ [Section 7.4, “Deleting and Re-Creating the BCC Identity Manager Drivers,” on page 63](#)
- ♦ [Section 7.5, “Finalizing the BCC Cluster Conversion,” on page 63](#)
- ♦ [Section 7.6, “What’s Next,” on page 63](#)

7.1 Understanding the Conversion Process

When you convert a business continuity cluster from NetWare to Linux, the tasks are similar to those you have performed before, but the sequence is different. Make sure you understand the process before you begin.

First, you convert all nodes in the business continuity cluster from NetWare to Linux. Convert the nodes one node at a time for only one peer cluster at a time. Do not finalize the cluster conversion for the peer cluster at this time. Repeat the NetWare-to-Linux conversion of the nodes for each peer cluster so that all of the nodes in every peer cluster are running Linux, but none of the cluster conversions have been finalized. For step-by-step procedures, see [Section 7.3, “Converting Clusters from NetWare to Linux,” on page 62](#).

After each node in every peer cluster is running Linux, you must delete and re-create all instances of the BCC drivers for Identity Manager in each of the peer clusters. For step-by-step procedures, see [Section 7.4, “Deleting and Re-Creating the BCC Identity Manager Drivers,” on page 63](#).

After all nodes are running Linux and the BCC drivers have been re-created, you are ready to finalize the BCC cluster conversion. Finalize the conversion on one peer cluster at a time. For step-by-step procedures, see [Section 7.5, “Finalizing the BCC Cluster Conversion,” on page 63](#).

7.2 Prerequisites for Converting from NetWare to Linux

Before you can upgrade from servers running BCC 1.0 (NetWare only) or BCC 1.1 SP1 for NetWare to BCC 1.2 for OES 2 SP1 Linux, you must upgrade the operating system and Novell Cluster Services™ on each server in every cluster to NetWare 6.5 SP8 (OES 2 SP1 NetWare), then upgrade the BCC software in the clusters to BCC 1.1 SP2 for NetWare.

For information about upgrading BCC in NetWare clusters, see [“Upgrading Business Continuity Clustering for NetWare”](#) in the *BCC 1.1 SP2: Administration Guide for NetWare 6.5 SP8*.

IMPORTANT: Every cluster node in each cluster in your business continuity cluster must be upgraded to BCC 1.1 SP2 for NetWare before you begin converting the NetWare clusters to OES 2 SP1 Linux.

Make sure your hardware and shared storage meet the requirements for BCC 1.2 for OES 2 SP1 Linux. For information, see [Section 4.1, “Requirements for BCC 1.2 for OES 2 SP1 Linux,”](#) on page 35.

7.3 Converting Clusters from NetWare to Linux

To convert a business continuity cluster from BCC 1.1 SP2 for NetWare to BCC 1.2 for Linux, you must convert all of the nodes in every peer cluster from NetWare 6.5 SP8 (OES 2 SP1 NetWare) to OES 2 SP1 Linux. Consider the following caveats as you convert the clusters to Linux:

- ♦ Clusters containing both NetWare and Linux servers are supported in a BCC only as a temporary means to convert a cluster from NetWare to Linux.
- ♦ Part of the temporary conversion state includes a restriction that only one mixed cluster can exist in your business continuity cluster at a time. For example, Cluster A can have both NetWare and Linux nodes, but Cluster B cannot. All nodes in Cluster B must be either NetWare or Linux.
- ♦ The same restrictions that apply to migrating or failing over resources between nodes within a mixed cluster also apply to migrating or failing over resources between clusters in a mixed BCC. You can migrate or fail over only NSS pool/volume resources that were originally created on NetWare between the peer clusters in a mixed BCC.

On one cluster at a time, do the following:

- 1 Convert NetWare cluster nodes to Linux by following the instructions in “[Converting NetWare 6.5 Clusters to OES 2 Linux](#)” in the *OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide*.

IMPORTANT: Do not perform the step to finalize the cluster conversion until after all of the nodes in every peer cluster in the business continuity cluster have been converted to Linux, and the BCC drivers for Identity Manager have been re-created.

- 2 Add a Read/Write replica of the cluster’s eDirectory™ partition to the new Linux nodes.
- 3 Stop the Identity Manager drivers, then remove the old NetWare Identity Manager node from the cluster.
- 4 Install Identity Manager 3.6x on one Linux node in each cluster.

IMPORTANT: Identity Manager 3.6 supports only a 32-bit OES 2 Linux operating system, even on 64-bit hardware. Identity Manager 3.6.1 supports the 64-bit OES 2 SP1 Linux operating system.

- 5 Configure the BCC Administrator user and group for the Linux cluster.

On Linux, the BCC Administrator user must be Linux-enabled with Linux User Management. The user must also be added to the Novell Cluster Services administration group (such as ncsgroup) on each cluster node. Follow the steps outlined in [Section 4.3, “Configuring a BCC Administrator User and Group,”](#) on page 46 to do the following:

- 5a Create the bccgroup and add the BCC Administrator to it.

- 5b** LUM-enable the `bccgroup` and add all the Linux nodes in the cluster to it.
 - 5c** Assign trustee rights for the BCC Administrator user to the cluster objects.
 - 5d** Add the BCC Administrator user to the `ncsgroup` on each cluster node.
- 6** (Conditional) If you LUM-enabled the Linux/UNIX Workstation object for a server after BCC is running, restart BCC on that server before continuing to the next node in the cluster.

7.4 Deleting and Re-Creating the BCC Identity Manager Drivers

After all of the nodes in every peer cluster has been successfully converted to Linux, you must re-create the BCC-specific Identity Manager drivers before finalizing the conversion. For information about re-creating the drivers, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,”](#) on page 65.

7.5 Finalizing the BCC Cluster Conversion

Normally, when converting a NetWare cluster to Linux, you need to run the `cluster convert` command after every node in the cluster has been converted to Linux in order to finalize the conversion. When converting a business continuity cluster from NetWare to Linux, do not run the `cluster convert` command after you convert the nodes in a given peer cluster. You must wait until all of the nodes in every peer cluster in the business continuity cluster has been upgraded to the latest NetWare version and converted to Linux, and you have re-created the BCC drivers. Then you can run the `cluster convert` command one cluster at a time to finalize the BCC cluster conversion. See “[Finalizing the Cluster Conversion](#)” in the *OES 2 SPI: Novell Cluster Services 1.8.6 for Linux Administration Guide*.

7.6 What’s Next

After the BCC is upgraded, continue with [Chapter 9, “Configuring BCC for Peer Clusters,”](#) on page 79.

Configuring the Identity Manager Drivers for BCC

8

Novell® Business Continuity Clustering (BCC) software provides two drivers for Identity Manager that are used to synchronize cluster resources and User objects 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 D, “Configuration Worksheet for the BCC Drivers for Identity Manager,”](#) on page 141.

- ◆ [Section 8.1, “Understanding the BCC Drivers,”](#) on page 65
- ◆ [Section 8.2, “Prerequisites for Configuring the BCC Drivers for Identity Manager,”](#) on page 70
- ◆ [Section 8.3, “Configuring the BCC Drivers,”](#) on page 71
- ◆ [Section 8.4, “Creating SSL Certificates,”](#) on page 74
- ◆ [Section 8.5, “Enabling or Disabling the Synchronization of e-Mail Settings,”](#) on page 74
- ◆ [Section 8.6, “Synchronizing Identity Manager Drivers,”](#) on page 75
- ◆ [Section 8.7, “Preventing Synchronization Loops for Identity Manager Drivers,”](#) on page 75
- ◆ [Section 8.8, “Changing the Identity Manager Synchronization Drivers,”](#) on page 77
- ◆ [Section 8.9, “What’s Next,”](#) on page 78

8.1 Understanding the BCC Drivers

Business Continuity Clustering provides two templates that are used with the eDirectory driver in Identity Manager to create the BCC drivers:

- ◆ **Cluster Resource Synchronization:** 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.
- ◆ **User Object Synchronization:** A set of policies, filters, and objects that synchronize User objects between any any two trees (or partitions) that contain the clusters in the business continuity cluster. Typically, this template is used to configure drivers when the clusters in your business continuity cluster are in different eDirectory™ trees.

IMPORTANT: Using two eDirectory trees is not supported for BCC on Linux.

You might also need to set up User Object Synchronization drivers between clusters if you put User objects in a different eDirectory partition than is used for the Cluster objects. This is not a recommended configuration; however, it is explained below for completeness.

Both the *Cluster Resource Synchronization* driver and the *User Object Synchronization* driver can be added to the same driver set. The driver set can also contain multiple instances of a given driver. For example, you have an 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 and between any two trees. The default port in the Cluster Resource Synchronization template is 2002. The default port in the User Object Synchronization template is 2001. You can use any ports that are unique for each instance of a driver, and that are not otherwise allocated. Make sure the ports are not blocked by the firewall. Examples of port assignments are shown in the tables below.

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.

For example, 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 8-1](#).

Table 8-1 *Single-Tree 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 8-2](#).

Table 8-2 *Driver Set Summary for a Single-Tree, 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

If the clusters are in different trees, or if the User objects are in a separate eDirectory partition than Cluster objects, you also need to install an instance of the User Object Synchronization driver on a different port, as shown in [Table 8-3](#) and [Table 8-4](#).

Table 8-3 *Two-Cluster Driver Set Example with User Object Synchronization*

Cluster Resource and User Object	Subscriber Node	
	Cluster One	Cluster Two
Cluster One	Not applicable	CR, port 2002 UO, port 2001
Cluster Two	CR, port 2002 UO, port 2001	Not applicable

Table 8-4 *Driver Set Summary for a Two-Cluster Business Continuity Cluster with User Object Synchronization*

Driver Instance	Driver Set for Cluster One	Driver Set for Cluster Two
Cluster Resource	C1 to C2, port 2002	C2 to C1, port 2002
User Object	C1 to C2, port 2001	C2 to C1, port 2001

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 8-5](#) and [Table 8-6](#).

Table 8-5 *Single-Tree Three-Cluster Driver Set Example*

Cluster Resource	Subscriber 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 8-6 *Driver Set Summary for a Single-Tree, 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		

If one of the clusters is in a different tree, or if the User objects are in a separate eDirectory partition, you also need to install an instance of the User Object Synchronization driver on a different port for the two nodes that communicate across the tree (or across the partitions). [Table 8-7](#) shows Cluster One and Cluster Two in Tree A (or User_PartitionA) and Cluster Three in Tree B (or User_PartitionB). The User Object Synchronization driver has been set up for Cluster One and Cluster Three to communicate across the trees (or across the partitions).

Table 8-7 *Three-Cluster Driver Set Example with User Object Synchronization*

Cluster Resource and User Object	Subscriber Node			
	Publisher Node	Cluster One	Cluster Two	Cluster Three
Cluster One (master node)		Not applicable	CR, port 2002	CR, port 2003 UO, port 2001
Cluster Two		CR, port 2002	Not applicable	No channel
Cluster Three (master node in the second partition)		CR, port 2003	No channel	Not applicable

You install the drivers on each cluster, with multiple instances needed only where the master cluster talks to multiple clusters and across trees, as shown in [Table 8-8](#).

Table 8-8 *Driver Set Summary for a Three-Cluster Business Continuity Cluster with User Object Synchronization*

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		
User Object	C1 to C3, port 2001		C3 to C1, port 2001

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 8-9](#).

Table 8-9 *Single-Tree 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 8-10](#).

Table 8-10 *Driver Set Summary for a Single-Tree, 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			

In the four-cluster business continuity cluster, you can set up the fourth node to talk to any one of the other three, making sure to avoid a configuration that results in a synchronization loop. This might be desirable if Cluster One and Cluster Two are in one tree (or user object partition), and Cluster Three and Cluster Four are in a second tree (or user object partition). In this case, you could set up channels for the Cluster Resource Synchronization driver between Cluster One and Cluster Two, between Cluster One and Cluster Three, and between Cluster Three and Cluster Four. You must assign a separate port for each of these channels, as shown in [Table 8-11](#). You also need to install an instance of the User Object Synchronization driver on a different port between the two clusters that communicate across the two trees (or across the two User object partitions).

Table 8-11 *Four-Cluster Driver Set Example with User Object Synchronization*

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	No channel	
Cluster Two	CR, port 2002	Not applicable	No channel	UO, port 2001	No channel

Cluster Resource	Subscriber Node			
Publisher Node	Cluster One	Cluster Two	Cluster Three	Cluster Four
Cluster Three (master node in the second partition)	CR, port 2003 UO, port 2001	No channel	Not applicable	CR, port 2004
Cluster Four	No channel	No channel	CR, port 2004	Not applicable

You install the drivers on each cluster, with multiple instances needed only where the master cluster talks to multiple clusters and across trees, as shown in [Table 8-12](#).

Table 8-12 Driver Set Summary for a Four-Cluster Business Continuity Cluster with User Object Synchronization

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 C3, port 2004
Cluster Resource	C1 to C3, port 2003		C3 to C4, port 2004	
User Object	C1 to C3, port 2001		C3 to C1, port 2001	

8.2 Prerequisites for Configuring the BCC Drivers for Identity Manager

- ◆ [Section 8.2.1, “Identity Manager,” on page 70](#)
- ◆ [Section 8.2.2, “Novell eDirectory,” on page 71](#)
- ◆ [Section 8.2.3, “Landing Zone Container,” on page 71](#)

8.2.1 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 4.1.7, “Identity Manager 3.6 Bundle Edition,” on page 40](#).

Identity Manager plug-ins for iManager require that eDirectory is running and working properly on the master eDirectory replica in the tree.

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

During the setup, you will make the IDM Driver object security equivalent to an existing User object. The IDM Driver object must have sufficient rights to any object it reads or writes in the following containers:

- ◆ The Identity Manager driver set container.

- ◆ The container where the Cluster object resides.
- ◆ The container where the server objects reside.

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

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

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

- ◆ In a multiple-partition business continuity cluster, the container where the User objects reside that need to be synchronized between the eDirectory partitions.

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

IMPORTANT: If you choose to include User object synchronization, exclude the Admin User object from being synchronized.

8.2.2 Novell 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 4.1.5, “Novell eDirectory 8.8,” on page 38](#).

8.2.3 Landing Zone Container

The landing zone that you specify for drivers must already exist. You can optionally create a separate container in eDirectory specifically for these cluster pool and volume objects.

8.3 Configuring the BCC Drivers

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your administrator username and password, specify the tree where you want to log in, then click *Login*.
- 3 In iManager, click *Identity Manager > Identity Manager Overview*.
- 4 Browse to select the Identity Manager server in this cluster that the driver set is associated with.
This is the node in the cluster where you installed the Identity Manager engine and eDirectory driver.
- 5 If a BCC driver set does not exist for this cluster, create it now.
 - 5a On the Identity Manager Overview page, click *Driver Sets > New*.
 - 5b Type the name of the driver set you want to create for this cluster.

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

- 5c** 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*

- 5d** Deselect (disable) the *Create a new partition on this driver set* option, then click *Next*.

- 6** On the Driver Set Overview page, click *Drivers > Add Driver* from the drop-down menu.

- 7** Verify that the driver set for the cluster is specified in an existing driver set text box, then click *Next*.

If the driver set does not exist, go to [Step 5](#) and create it.

- 8** Browse to select the server in this cluster that has Identity Manager installed on it, then click *Next*.

- 9** Open the *Show* drop-down menu and select *All Configurations*.

- 10** Select one of the BCC preconfigured driver template files from the *Configurations* drop-down menu, then click *Next*.

- ◆ To create a cluster resource synchronization driver instance, select the `BCCClusterResourceSynchronization.xml` file.
- ◆ To create a user object synchronization driver instance, select the `UserObjectSynchronization.xml` file.

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

Each field contains an example of the type of information that should go into the field.

Descriptions of the information required are also included with each field.

- ◆ **Driver name for this driver instance:** Specify a unique name for this driver 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 *Cluster1toCluster2 BCC Sync*.

If you use both preconfigured templates, you must specify different driver names for each of the driver instances that represent that same connection. For example, *Cluster1toCluster2 BCCCR Sync* and *Cluster1toCluster2 BCCUO Sync*.

- ◆ **Name of SSL Certificate to use:** Specify a unique name for the certificate such as *BCC Cluster Sync*. The certificate is created later in the configuration process in [“Creating SSL Certificates” on page 74](#), after you have created the driver instance.

In a single tree configuration, if you specify the SSL CertificateDNS certificate that was created when you installed OES 2 on the Identity Manager node, you do not need to create an additional SSL certificate later.

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

- ◆ **IP address or DNS name of 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 *servername.cluster2.clusters.siteB.example*.
- ◆ **Port number for this connection:** You must specify unique port numbers for each driver instance for a given connection between two clusters. The default port number is 2002 for the cluster resource synchronization and 2001 for the user object 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 Cluster 2 to Cluster 1 resource synchronization driver instance for the peer driver you create on Cluster2.

- ♦ **Full Distinguished Name (DN) of the cluster this driver services:** For example, *cluster1.clusters.siteA.example*.
- ♦ **Fully Distinguished Name (DN) of the landing zone container:** Specify the context of the container where the cluster pool and volume objects in the other cluster are placed when they are synchronized to this cluster.

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

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

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

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

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

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

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

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

- ♦ In a multiple-partition business continuity cluster, the container where the User objects reside that need to be synchronized between the eDirectory partitions.

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

IMPORTANT: If you choose to include User object synchronization, exclude the Admin User object from being synchronized.

12a Click *Define Security Equivalences*, then click *Add*.

12b Browse to and select the desired User object, then click *OK*.

12c Click *Next*, and then click *Finish*.

13 Repeat [Step 1](#) through [Step 12](#) above on the peer clusters in your business continuity cluster.

This includes creating a new driver and driver set for each cluster. Remember that you create the User Synchronization only on the peer clusters that are actually communicating with each other between the partitions.

14 After you have configured the BCC IDM drivers on every node in each cluster, you must upgrade the drivers to the Identity Manager 3.6x architecture.

Do the follow to upgrade each BCC driver set you created in “[Configuring the BCC Drivers](#)” on page 71

14a In iManager, click *Identity Manager*, then click *Identity Manager Overview*.

14b Search for the driver sets that you have added, then click the driver set link to bring up the *Driver Set Overview*.

14c Click the red *Cluster Sync* icon, and you should be prompted to upgrade the driver.

8.4 Creating SSL Certificates

If SSL certificates are not present or have not been created, Identity Manager drivers might not start or function properly. We recommend using SSL certificates for encryption and secure information transfer between clusters and the Identity Manager vault. Create separate certificates for the Cluster Resource Synchronization driver and the User Synchronization driver.

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

To create an SSL certificate:

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

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

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

3 In *Roles and Tasks*, click *Identity Manager Overview*, then click *NDS-to-NDS Driver Certificates*.

4 Specify the requested driver information for this cluster, then click *Next*.

You must specify the driver name (including the context) you supplied in [Step 11 on page 72](#) for this cluster. Use the following format when specifying the driver name:

`DriverName.DriverSet.OrganizationalUnit.OrganizationName`

Ensure that there are no spaces (beginning or end) in the specified context, and do not use the `cn=DriverName.ou=OrganizationalUnitName.o=OrganizationName` format.

5 Specify the requested driver information for the driver in the other cluster.

Use the same format specified in [Step 4](#).

6 Click *Next*, then click *Finish*.

8.5 Enabling or Disabling the Synchronization of e-Mail Settings

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

8.6 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 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Identity Manager*, then click the *Identity Manager Overview* link.
- 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 the *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.

8.7 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 User objects and Cluster Resource objects in a manner that prevents Identity Manager synchronization loops. Identity Manager synchronization loops can cause excessive network traffic and slow server communication and performance.


For example, in a three-cluster business continuity cluster, an Identity Manager synchronization loop occurs when Cluster One is configured to synchronize with Cluster Two, Cluster Two is configured to synchronize with Cluster Three, and Cluster Three is configured to synchronize back to Cluster One. This is illustrated in [Figure 8-1](#) below.

Figure 8-1 Three-Cluster Identity Manager Synchronization Loop




A preferred method is to make Cluster One an Identity Manager synchronization master in which Cluster One synchronizes with Cluster Two, and Cluster Two and Cluster Three both synchronize with Cluster One. This is illustrated in [Figure 8-2](#) below.

Figure 8-2 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 8-3](#).

Figure 8-3 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 8-4](#).

Figure 8-4 Single-Tree, Four-Cluster Identity Manager Synchronization Scenario



8.8 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.

8.9 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. For information, see [Chapter 9, “Configuring BCC for Peer Clusters,”](#) on page 79.

Configuring BCC for Peer Clusters

9

After you have installed and configured Identity Manager, the Novell® 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. Make sure 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 8, “Configuring the Identity Manager Drivers for BCC,”](#) on page 65.

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

- ♦ [Section 9.1, “Enabling Clusters for Business Continuity,”](#) on page 79
- ♦ [Section 9.2, “Adding Peer Cluster Credentials,”](#) on page 80
- ♦ [Section 9.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 82
- ♦ [Section 9.4, “Adding Storage Management Configuration Information,”](#) on page 83
- ♦ [Section 9.5, “Configuring CIMOM Daemons to Bind to IP Addresses,”](#) on page 86
- ♦ [Section 9.6, “Enabling Linux POSIX File Systems to Run on Secondary Clusters,”](#) on page 86
- ♦ [Section 9.7, “Verifying BCC Administrator User Trustee Rights and Credentials,”](#) on page 87
- ♦ [Section 9.8, “Disabling BCC for a Peer Cluster,”](#) on page 88
- ♦ [Section 9.9, “What’s Next,”](#) on page 88

9.1 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.

- 1** Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed. This server must be in the same eDirectory™ tree as the cluster you are enabling for business continuity.
- 2** Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3** Ensure that the BCC-specific Identity Manager drivers are running:
 - 3a** In *Roles and Tasks*, click *Identity Manager*, then click the *Identity Manager Overview* link.
 - 3b** Search the eDirectory Container or tree for the BCC-specific Identity Manager drivers.

- 3c** For each driver, click the upper right corner of the driver icon to see if a driver is started or stopped.
- 3d** If the driver is stopped, start it by selecting *Start*.
- 4** In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
- 5** Specify a cluster name, or browse and select one.
- 6** Click the *Properties* button, then click the *Business Continuity* tab.
- 7** Select the *Enable Business Continuity Features* check box.
- 8** Repeat [Step 1](#) through [Step 7](#) for each cluster that you want to add to the business continuity cluster.
- 9** 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.
- 10** Continue with [Adding Peer Cluster Credentials](#).

9.2 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 (username and password) of the user who the selected cluster will use to authenticate to a selected peer cluster.

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

- ♦ [Section 9.2.1, “Using Console Commands to Add Credentials,” on page 80](#)
- ♦ [Section 9.2.2, “Using iManager to Add Credentials,” on page 81](#)

9.2.1 Using Console Commands to Add Credentials

To add peer cluster credentials, do the following for each node of every cluster in the business continuity cluster:

- 1** Open a terminal console on the cluster node where you want to add peer credentials, then log in as the `root` user.
- 2** At the terminal console prompt, enter

```
cluster connections
```
- 3** 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 75](#).
- 4** For each cluster in the list, enter the following command at the server console prompt, then enter the `bccadmin` username and password when prompted.

```
cluster credentials cluster_name
```

5 Repeat the following steps for every node in each 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 member of the `ncsgroup`.

```
id bccadmin
```

9.2.2 Using iManager to Add Credentials

You cannot use iManager on Linux to set eDirectory credentials for BCC. You must use iManager on NetWare[®] or Windows (the server must be in the same eDirectory tree), or use the Linux BCC command line interface from a console prompt to set credentials.

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

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

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

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

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

Rather than using the Admin user to administer your BCC, you should consider creating another user with sufficient rights to the appropriate contexts in your eDirectory tree to manage your BCC. For information, see [Section 4.3, “Configuring a BCC Administrator User and Group,” on page 46](#).

3 Repeat [Step 1](#) and [Step 2](#) for the other cluster that this cluster will migrate resources to.

4 Continue with [Adding Search-and-Replace Values to the Resource Replacement Script](#).

9.3 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 inbound resources that have IP addresses starting with “10.10.20.” 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 inbound resources that have IP addresses starting with “10.10.10.” 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 > Cluster Options*, select the Cluster object, click *Properties*, then select *Business Continuity*.
- 2 In the *Resource Replacement Script* section of the Business Continuity Cluster Properties page, click *New*.
- 3 Add the desired search-and-replace values.

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

For example, if you 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 addresses, and replaces them with the 192.168.1.1, 192.168.1.2, 192.168.1.3, and 192.168.1.4 addresses, respectively.

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

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

- ♦ [Regular Expressions \(http://www.opengroup.org/onlinepubs/007908799/xbd/re.html\)](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.

5 Click *Apply* to save your changes.

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

6 Verify that the change has been synchronized with the peer clusters by the Identity Vault.

7 Continue with [Section 9.4, “Adding Storage Management Configuration Information,” on page 83.](#)

9.4 Adding Storage Management Configuration Information

You can create BCC load and unload scripts for each BCC-enabled resource in each peer 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.

You can also add commands and Perl scripts to resource scripts to call other scripts. Any command that can be run at the Linux terminal console can be used. The scripts or commands you add are stored in eDirectory. If you add commands to call outside scripts, those scripts must exist 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 MasterIP 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.
 - ◆ 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.

To add storage management configuration information:

- 1** Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace `server_ip_address` with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2** Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3** In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
- 4** Specify a cluster name, or browse and select one.
- 5** Under *Cluster Objects*, select a cluster resource that is enabled for business continuity, then click *Details*.
Cluster resources that are enabled for business continuity have the BCC label on the resource type icon.
- 6** Click the *Business Continuity* tab, then click *BCC Scripts*.
- 7** Create BCC storage management load and unload scripts:
 - 7a** Under *BCC Load Scripts*, click *New* to bring up a page that lets you create a script to promote mirrored LUNs on a cluster.
You can also delete a script, edit a script by clicking *Details*, or change the order in which load scripts execute by clicking the *Move Up* and *Move Down* links.
 - 7b** 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.
CIMOM IP or DNS	If you selected the <i>CIM Client</i> check box on the previous page and you are not using a template, specify the IP address or DNS name for your storage system. This is the IP address or DNS name that is used for storage management.

Parameter	Description
Namespace	<p>If you selected the <i>CIM Client</i> check box on the previous page, accept the default namespace, or specify a different namespace for your storage system.</p> <p>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.</p>
Username and password	<p>If you selected the <i>CIM Client</i> check box on the previous page, accept the default namespace, or specify a different namespace for your storage system.</p>
Port	<p>If you selected the <i>CIM Client</i> check box on the previous page, accept the default port number or specify a different port number. This is the port number that CIMOM (your storage system manager) uses. Consult your storage system documentation to determine which port number you should use.</p>
Secure	<p>If you selected the <i>CIM Client</i> check box on the previous page, select the <i>Secure</i> check box if you want storage management communication to be secure (HTTPS). Deselect the <i>Secure</i> check box to use non-secure communications (HTTP) for storage management communications.</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 <i>New</i>, then provide the variable name and value in the fields provided. Click <i>OK</i> to save your entries. You can specify additional variables by clicking <i>New</i> again and providing variable names and values. You can also edit and delete existing script parameters by clicking the applicable link.</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 storage hardware. You can add a Perl script, or any commands that can be run on Linux.</p> <hr/> <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>
CIM enabled	<p>Select this check box if your storage system supports SMI-S and you did not select the <i>CIM Client</i> check box on the previous page. This causes the CIM-specific fields to become active on this page.</p>
Synchronous	<p>Select this check box to run scripts sequentially (that is, one at a time). Deselect this check box to allow multiple scripts to run concurrently. Most storage system vendors do not support running multiple scripts concurrently.</p>
Edit flags	<p>This is an advanced feature, and should not be used except under the direction of Novell Support.</p>

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

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

IMPORTANT: The CIMOM daemons on all nodes in the business continuity cluster should be configured to bind to all IP addresses on the server. For information, see [Section 9.5, “Configuring CIMOM Daemons to Bind to IP Addresses,”](#) on page 86.

9.5 Configuring CIMOM Daemons to Bind to IP Addresses

The CIMOM daemons on all nodes in the business continuity cluster must be configured to bind to all of the IP addresses on a server.

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

You can do this by editing the `/etc/openwbem/openwbem.conf` file. Set the `http_server.listen_addresses` parameter of the `openwbem.conf` file to `0.0.0.0`, which is the default.

Change the following section in the `openwbem.conf` file:

```
# http_server.listen_addresses option specifies the local addresses
# to listen on. The option is a space delimited list.
# Each item is either a hostname or an IP address.
# The value 0.0.0.0 means to listen on all local addresses.
# This is a multi-valued option. Whitespace is the separator.
# The default is 0.0.0.0
http_server.listen_addresses = 0.0.0.0
```

For more information about managing OpenWBEM, see the [OES 2: OpenWBEM Services Administration Guide](#) (http://www.novell.com/documentation/oes2/mgmt_openwbem_lx_nw/data/front.html).

9.6 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 EVMS CSM (Cluster Segment Manager) 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, and not for NSS pools.

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

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

- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
- 4 Specify a cluster name, or browse and select one.
- 5 Under *Cluster Objects*, select the business-continuity-enabled cluster resource that contains the Reiser or Ext3 file system, then click *Details*.
Cluster resources that are enabled for business continuity have the BCC label on the resource type icon.
- 6 Click the *Business Continuity* tab, then click *Storage Management*.
- 7 Under *BCC Load Scripts*, click *New* to bring up a wizard that lets you create a new script.
- 8 In the wizard, specify the following information:
 - ♦ **Name:** Convert CSM container.
 - ♦ **Description:** Converts the EVMS Cluster Segment Manager (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 7b](#) in “[Adding Storage Management Configuration Information](#)” on page 83 for more information and descriptions of the information in the value fields.
- 9 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 EVMS CSM (Cluster Segment Manager) container. You assigned this name to the EVMS CSM container when you created it. This value is case-sensitive and should exactly match the container name.
- 10 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 1.2 CD or ISO image.
- 11 Click *Apply* to save your changes.
- 12 Repeat [Step 1](#) through [Step 11](#) for the peer clusters in your BCC.
The information you provide in the steps should be unique for each cluster.

9.7 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 Administrative User](#)” on page 46.

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 47.

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 47](#).

9.8 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. Make sure to 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 11.6, “Disabling BCC for a Cluster Resource,” on page 102](#).

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.

9.9 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 9, “Configuring BCC for Peer Clusters,” on page 79](#).

Managing a Business Continuity Cluster

10

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

IMPORTANT: Identity Manager must be configured and running on one node in each peer cluster before any BCC enabled cluster resource changes are made. Make sure 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 8, “Configuring the Identity Manager Drivers for BCC,”](#) on page 65.

For information about using console commands to manage your business continuity cluster, see [Appendix A, “Console Commands for BCC,”](#) on page 127.

- ♦ [Section 10.1, “Migrating a Cluster Resource to a Peer Cluster,”](#) on page 89
- ♦ [Section 10.2, “Bringing a Downed Cluster Back in Service,”](#) on page 91
- ♦ [Section 10.3, “Changing Peer Cluster Credentials,”](#) on page 91
- ♦ [Section 10.4, “Viewing the Current Status of a Business Continuity Cluster,”](#) on page 92
- ♦ [Section 10.5, “Generating a Cluster Report,”](#) on page 93
- ♦ [Section 10.6, “Resolving Business Continuity Cluster Failures,”](#) on page 93

10.1 Migrating a Cluster Resource to a Peer Cluster

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

- ♦ [Section 10.1.1, “Understanding BCC Resource Migration,”](#) on page 89
- ♦ [Section 10.1.2, “Migrating Cluster Resources between Clusters,”](#) on page 90

10.1.1 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 a 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.

10.1.2 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 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace `server_ip_address` with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters*, then click the *BCC Manager* link.
- 4 Specify a cluster name, or browse and select the Cluster object of the cluster you want to manage the BCC from.
- 5 Select one or more cluster resources, then click *BCC Migrate*.

The cluster you chose in [Step 4](#) 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 from so that you understand the point of view that the status information is provided in the first table. 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*.

- 6 In the list of cluster, select the cluster where you want to migrate the selected resources, then click *OK*.

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

- 7 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.

10.2 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.

10.3 Changing Peer Cluster Credentials

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

IMPORTANT: Make sure the new administrator username meets the requirements specified in [Section 4.3, “Configuring a BCC Administrator User and Group,” on page 46](#).

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

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

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

- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters > BCC Manager*, then click the *Management* link.
- 4 Browse and select the Cluster object of the cluster.
- 5 Click *Connections* and select a peer cluster.
- 6 Edit the administrator username and password that the selected cluster will use to connect to the selected peer cluster, then click *OK*.

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

NOTE: If the business continuity cluster has clusters in multiple eDirectory™ trees, and you specify a username and password that is used by all peer clusters (that is, credentials for an administrator user who all peer clusters have in common), each eDirectory tree in the business continuity cluster must have the same username and password.

10.4 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 10.4.1, “Using iManager to View the Cluster Status,” on page 92](#)
- ♦ [Section 10.4.2, “Using Console Commands to View the Cluster Status,” on page 92](#)

10.4.1 Using iManager to View the Cluster Status

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters*, then click *BCC Manager*.
- 4 Browse and select the Cluster object of the cluster you want to manage.
- 5 Use the page to see if all cluster peer connections are up or if one or more peer connections are down. You can also see the status of the BCC resources in the business continuity cluster.

10.4.2 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
```

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

10.6 Resolving Business Continuity Cluster Failures

There are several failure types associated with a business continuity cluster that you should be aware of. Understanding the failure types and knowing how to respond to each can help you more quickly recover a cluster. Some of the failure types and responses differ depending on whether you have implemented 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, but your storage vendor documentation might use different terms such as *mask* and *unmask*.

- ♦ [Section 10.6.1, “Storage-Based Mirroring Failure Types and Responses,” on page 93](#)
- ♦ [Section 10.6.2, “Host-based Mirroring Failure Types and Responses,” on page 95](#)

10.6.1 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 94](#)
- ♦ [“Primary Cluster and Primary Storage System Both Fail” on page 94](#)
- ♦ [“Secondary Cluster Fails but Secondary Storage System Does Not” on page 94](#)
- ♦ [“Secondary Cluster and Secondary Storage System Both Fail” on page 94](#)
- ♦ [“Primary Storage System Fails and Causes the Primary Cluster to Fail” on page 94](#)
- ♦ [“Secondary Storage System Fails and Causes the Secondary Cluster to Fail” on page 94](#)
- ♦ [“Intersite Storage System Connectivity Is Lost” on page 95](#)
- ♦ [“Intersite LAN Connectivity Is Lost” on page 95](#)

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, you must 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. Make sure 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.

Intersite Storage System Connectivity Is Lost

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

Intersite 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 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 B, “Setting Up Auto-Failover,” on page 131](#).

10.6.2 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 95
- ♦ “Primary Cluster and Primary Storage System Both Fail” on page 95
- ♦ “Secondary Cluster Fails but Secondary Storage System Does Not” on page 95
- ♦ “Secondary Cluster and Secondary Storage System Both Fail” on page 96
- ♦ “Primary Storage System Fails and Causes the Primary Cluster to Fail” on page 96
- ♦ “Secondary Storage System Fails and Causes the Secondary Cluster to Fail” on page 96
- ♦ “Intersite Storage System Connectivity Is Lost” on page 96
- ♦ “Intersite LAN Connectivity is Lost” on page 96

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 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 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.

Intersite Storage System Connectivity Is Lost

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

Intersite 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 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 B, “Setting Up Auto-Failover,” on page 131](#).

Configuring BCC for Cluster Resources

11

After you have set up the Novell® Cluster Services™ clusters for a business continuity cluster by using Novell Business Continuity Clustering software, you are ready to configure the cluster resources for BCC. You can enable one or multiple cluster resources in each of the peer clusters for business continuity that if you want to be able to fail over between peer clusters. For each resource, you can specify the preferred peer clusters for failover.

- ♦ [Section 11.1, “Requirements for Cluster Resources,” on page 97](#)
- ♦ [Section 11.2, “BCC-Enabling Cluster Resources,” on page 98](#)
- ♦ [Section 11.3, “Configuring Search-and-Replace Values for an Individual Cluster Resource,” on page 99](#)
- ♦ [Section 11.4, “Assigning Preferred Peer Clusters for the Resource,” on page 100](#)
- ♦ [Section 11.5, “Assigning Preferred Nodes in Peer Clusters,” on page 101](#)
- ♦ [Section 11.6, “Disabling BCC for a Cluster Resource,” on page 102](#)
- ♦ [Section 11.7, “Changing the IP Address of a Cluster Resource,” on page 103](#)
- ♦ [Section 11.8, “Deleting or Unsharing a BCC-Enabled Shared NSS Pool Resource,” on page 103](#)

11.1 Requirements for Cluster Resources

- ♦ [Section 11.1.1, “LUNs for Cluster Pool Resources,” on page 97](#)
- ♦ [Section 11.1.2, “Volumes for Cluster Pool Resources,” on page 97](#)
- ♦ [Section 11.1.3, “Shared Disk Cluster Resources,” on page 98](#)
- ♦ [Section 11.1.4, “Cluster Resources for OES 2 Linux Services,” on page 98](#)

11.1.1 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.

11.1.2 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.

11.1.3 Shared Disk Cluster Resources

See [Table 11-1](#) for resources that explain how to create shared disk cluster resources on Novell Open Enterprise Server 2 Linux servers:

Table 11-1 Shared Disk Cluster Resources on OES 2 Linux Servers

Shared Disk Cluster Resources	Procedure
Dynamic Storage Technology shadow volume pairs	“Configuring DST Shadow Volumes with Novell Cluster Services for Linux” in the <i>OES 2 SP1: Dynamic Storage Technology Administration Guide</i>
Linux POSIX file systems	“Configuring Cluster Resources for Shared Linux POSIX Volumes” in the <i>OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide</i>
NCP™ (NetWare Core Protocol™) volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 2 SP1: NCP Server for Linux Administration Guide</i>
Novell Storage Services™ (NSS) pools and volumes	“Configuring Cluster Resources for Shared NSS Pools and Volumes” in the <i>OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide</i>

11.1.4 Cluster Resources for OES 2 Linux Services

For information about creating cluster resources for OES 2 Linux services, see the [OES 2 High-Availability documentation Web site \(http://www.novell.com/documentation/oes2/cluster-services.html#clust-config-resources\)](http://www.novell.com/documentation/oes2/cluster-services.html#clust-config-resources).

11.2 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 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
- 4 Specify a cluster name, or browse and select one.
- 5 Select the desired cluster resource from the list of Cluster objects, then click the *Details* link.

- 6 Select the *Business Continuity* tab.
- 7 Enable a cluster resource or cluster pool resource for business continuity by selecting the *Enable Business Continuity Features* check box.
- 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 by following the procedure in [Section 9.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 82, then continue with [Section 11.4, “Assigning Preferred Peer Clusters for the Resource,”](#) on page 100.
 - ♦ **Individual Resource Replacement Script:** (Not recommended) Create a replacement script that contains search-and-replace values only for an individual resource. Continue with [Section 11.3, “Configuring Search-and-Replace Values for an Individual Cluster Resource,”](#) on page 99.
- 9 Continue with [Section 11.4, “Assigning Preferred Peer Clusters for the Resource,”](#) on page 100.

11.3 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 replacement script for the entire cluster instead of for an individual resource. For information, see [Section 9.3, “Adding Search-and-Replace Values to the Resource Replacement Script,”](#) on page 82.

Before you create a replacement script for an individual resource, you should contact Novell Support.

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

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 82](#) 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 11.4, “Assigning Preferred Peer Clusters for the Resource,” on page 100](#).

11.4 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 (or remove a cluster and add it back in) to the business continuity cluster, 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 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 <i>Finish</i> .
Linux POSIX file system	Existing	Click <i>Apply</i> .
NSS pool	New	Click <i>Next</i> , then add the storage management configuration information. For information, see “Adding Storage Management Configuration Information” on page 83 .

Cluster Resource	New or Existing	Action
NSS pool	Existing	Add the storage management configuration information. For information, see “Adding Storage Management Configuration Information” on page 83.

- 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 11.5, “Assigning Preferred Nodes in Peer Clusters,”](#) on page 101.

11.5 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 by default set to all nodes in the cluster. 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.

IMPORTANT: Changes you make for a cluster resource configuration do not take effect while the resource is loaded or running on a server. You must offline and online the resource to apply the changes you make to its resource properties, policies, or scripts.

Before you begin, make sure that you have assigned the preferred peer clusters for the resource. For information, see [Section 11.4, “Assigning Preferred Peer Clusters for the Resource,”](#) on page 100.

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

- 1 In iManager, click *Clusters*, then click *Cluster Options*.
- 2 Browse to locate and select the Cluster object of the peer cluster you want to manage.
- 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 node assignment changes.

11.6 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.

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

To disable BCC for a cluster resource:

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

The primary peer cluster and the node where the resource is running cannot be moved from the *Assigned* list to the *Unassigned* list.
 - 7c** Click *Apply* to save node assignment changes.

- 8 On the *Details* page, click the *Business Continuity* tab, deselect the *Enable Business Continuity Features* check box, then click *Apply*.
- 9 Wait for Identity Manager to synchronize the changes.
This could take from 30 seconds to one minute, depending on your configuration.
- 10 Delete the Cluster Resource object on the clusters where you no longer want the resource to run.

11.7 Changing the IP Address of a Cluster Resource

You can change the IP address of a cluster resource by using the Cluster Properties page of the Clusters plug-in in iManager. Make sure to offline the cluster resource before attempting to change its IP address. When you online the cluster resource, the load script is updated with the new IP address.

11.8 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.
- ♦ Make sure to 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 iManager, select *Clusters > Cluster Options*, then select the Cluster object where the shared NSS pool resource is currently active.
- 2 If the resource is on a non-master node in the cluster, migrate it to the master node by selecting the resource, then clicking *Migrate* and specifying the master node.
- 3 Select the shared NSS pool resource, click *Properties*, then select *Business Continuity*.
- 4 Disable BCC for the pool cluster resource, then click *Apply*.

For information, see [Section 11.6, “Disabling BCC for a Cluster Resource,”](#) on page 102.

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

- 5 Select *Clusters > Cluster Options*, select the shared NSS pool, then select *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.

- 6 Do one of the following:

Task	Procedure
Keep the shared NSS pool to use only in the current cluster.	<ol style="list-style-type: none"> 1. In iManager, select <i>Clusters > Cluster Options</i>, select the shared NSS pool resource, select <i>Properties</i>, then select <i>Scripts</i>. 2. Modify the load script by removing any commands related to the BCC, then click <i>Apply</i>. 3. Modify the unload script by removing any commands related to the BCC, then click <i>Apply</i>. 4. Select <i>Clusters > Cluster Options</i>, select the shared NSS pool resource, then select <i>Online</i>.
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 <i>Storage > Pools</i>, then select the master node of the cluster. 2. Select the pool to view its <i>Pool Details</i>. 3. Mark the pool as <i>Not Shareable for Clustering</i>, then click <i>Apply</i>. 4. Select <i>Clusters > Cluster Options</i>, select the check box next to the Cluster Resource object you want to delete, then click <i>Delete</i>.
Unshare the pool, then delete it.	<ol style="list-style-type: none"> 1. In iManager, select <i>Storage > Pools</i>, then select the master node of the cluster. 2. Select the pool to view its <i>Pool Details</i>. 3. Mark the pool as <i>Not Shareable for Clustering</i>, then click <i>Apply</i>. 4. Select <i>Directory Administration > Delete Object</i>, select the Cluster Resource object for the shared NSS pool cluster resource, then click <i>OK</i>. 5. In iManager, select <i>Storage > Pools</i>, then select the master node of the cluster. 6. Select the pool to view the <i>Pool Details</i>. 7. Click <i>Delete</i>, then confirm the deletion.

Troubleshooting Business Continuity Clustering

12

This section contains the following topics to help you troubleshoot Novell® Business Continuity Clustering 1.2 for Linux.

- ◆ [Section 12.1, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted,” on page 106](#)
- ◆ [Section 12.2, “Identity Manager Plug-Ins Do Not Appear in iManager,” on page 106](#)
- ◆ [Section 12.3, “Cluster Connection States,” on page 106](#)
- ◆ [Section 12.4, “Driver Port Number Conflicts,” on page 107](#)
- ◆ [Section 12.5, “Excluded Users,” on page 108](#)
- ◆ [Section 12.6, “Security Equivalent User,” on page 108](#)
- ◆ [Section 12.7, “SSL Certificates,” on page 109](#)
- ◆ [Section 12.8, “Clusters Cannot Communicate,” on page 110](#)
- ◆ [Section 12.9, “BCC Startup Flags,” on page 110](#)
- ◆ [Section 12.10, “Identity Manager Drivers for Cluster Synchronization Do Not Start,” on page 111](#)
- ◆ [Section 12.11, “Identity Manager Drivers Do Not Synchronize Objects from One Cluster to Another,” on page 111](#)
- ◆ [Section 12.12, “Tracing Identity Manager Communications,” on page 112](#)
- ◆ [Section 12.13, “Peer Cluster Communication Is Not Working,” on page 113](#)
- ◆ [Section 12.14, “A Resource Does Not Migrate to a Peer Cluster,” on page 113](#)
- ◆ [Section 12.15, “Resource Cannot Be Brought Online,” on page 113](#)
- ◆ [Section 12.16, “Administration of Peer Clusters Is Not Functional,” on page 114](#)
- ◆ [Section 12.17, “Resource Script Search-and-Replace Functions Do Not Work,” on page 114](#)
- ◆ [Section 12.18, “Virtual NCP Server IP Addresses Won’t Change,” on page 114](#)
- ◆ [Section 12.19, “The IP Address, Virtual Server DN, or Pool Name Does Not Appear on the iManager Cluster Configuration Page,” on page 115](#)
- ◆ [Section 12.20, “Blank Error String iManager Error Appears While Bringing a Resource Online,” on page 115](#)
- ◆ [Section 12.21, “Mapping Drives in Login Scripts,” on page 115](#)
- ◆ [Section 12.22, “Mapping Drives to Home Directories by Using the %HOME_DIRECTORY Variable,” on page 116](#)
- ◆ [Section 12.23, “BCC Error Codes,” on page 117](#)
- ◆ [Section 12.24, “Clustered Pool Is Stuck in an eDirectory Synchronization State,” on page 118](#)

12.1 NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted

In some environments, a timing problem prevents the NDP user space application (`ndppapp`) from loading, and in turn, NSS cannot be loaded until the problem resolves itself. You can increase the UDEV event handling limits to 1024 to circumvent this problem. For instructions, see “[NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted \(Linux\)](#)” in the *OES 2 SP1: NSS File System Administration Guide*.

12.2 Identity Manager Plug-Ins Do Not Appear in iManager

After you properly install Identity Manager plug-ins in iManager, the plug-ins sometime 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, make sure 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 terminal console prompt as the `root` user:

```
rcndsd restart
```

12.3 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 12-1 BCC Connection States

BCC Connection State	Number	Description	Possible Actions
Normal	0	The connections between clusters are functioning normally.	None required.
Authenticating	1	BCC is in the process of authenticating to a peer cluster.	Wait until the authentication process is finished.
Invalid Credentials	2	You entered the wrong username or password for the selected peer cluster.	Enter the correct username and password that this cluster will use to connect to the selected peer cluster.

BCC Connection State	Number	Description	Possible Actions
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>etc/init.d/adminfs status</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 47.
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.

12.4 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 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Identity Manager*, then click the *Identity Manager Overview* link.
- 4 Select *Search Entire Tree*, then click *Search*.
- 5 Select the driver you want to check by clicking the red *Cluster Sync* icon or the blue *User Sync* icon.
- 6 Click the red or blue icon again, then click the *Identity Manager* tab (if it is not already selected).

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

For example, the *Authentication context* field might contain a value similar to 10.1.1.12:2003:2003. In this example, the first port number (2003) is the port number for the corresponding Identity Manager driver on the cluster 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* or *User Sync* icon (whichever you have chosen), then click *Restart driver*.
- 9 If you changed the port number in [Step 7](#) above, change the port numbers to be the same for the corresponding driver in the other cluster.

You can do this by repeating [Step 1](#) through [Step 8](#) for the Identity Manager driver on the other cluster.

12.5 Excluded Users

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

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

To see the excluded users list:

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Identity Manager*, then click the *Identity Manager Overview* link.
- 4 Select *Search Entire Tree*, then click *Search*.
- 5 Select the user synchronization driver you want to check by clicking the blue *User Sync* icon.
This is not necessary for the cluster synchronization driver.
- 6 Click the blue icon again, then click the *Identity Manager* tab if it is not already selected.
- 7 Click *Excluded Users*, then view, add, or remove users as desired.

12.6 Security Equivalent User

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

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

The IDM Driver object must have sufficient rights to 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 business-continuity-enabled resource are also placed in the landing zone.

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

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Identity Manager*, then click the *Identity Manager Overview* link.
- 4 Choose *Search Entire Tree*, then click *Search*.
- 5 Select the driver you want to check by clicking the red *Cluster Sync* icon or the blue *User Sync* icon.
- 6 Click the red or blue icon again, then click the *Identity Manager* tab if it is not already selected.
- 7 Click *Security Equals*, and view or add a security equivalent user as needed.
- 8 Repeat [Step 5](#) through [Step 7](#) for the other drivers in your BCC.

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

12.7 SSL Certificates

If SSL certificates are not present or have not been created, Identity Manager drivers might not start or function properly. See “[Creating SSL Certificates](#)” on [page 74](#) for more information on creating SSL certificates for BCC.

12.8 Clusters Cannot Communicate

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

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

12.9 BCC Startup Flags

There are three optional flags that can be used by BCC during startup.

Table 12-2 *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 `v`, `vt`, `t`, or `tv`.

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

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

12.10 Identity Manager Drivers for Cluster Synchronization Do Not Start

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

- ♦ A certificate has not been created. For information, see [“Creating SSL Certificates” on page 74](#).
- ♦ 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 8, “Configuring the Identity Manager Drivers for BCC,” on page 65](#) for more information.

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

- ♦ The driver has been disabled.

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

Selecting the Auto Start option is recommended.

- ♦ Unknown communications problems.

See [Section 12.12, “Tracing Identity Manager Communications,” on page 112](#).

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

```
rcnds 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 `iFolder Server` and you try to synchronize a resource named `iFolder_Server`, the synchronization fails. This is because the underscore character is mapped to a space. eDirectory returns an error that the entry already exists.

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

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

- ♦ The drivers are not communicating on the same port.

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

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

- ♦ See [Tracing Identity Manager Communications](#) below.

12.12 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 “[Adding Trace Levels in iManager](#)” (http://www.novell.com/documentation/idm36/idm_common_driver/data/b1rc6ea.html).

For information about using `ndstrace`, see “[Using ndstrace](#)” (<http://www.novell.com/documentation/edir88/edir88tshoot/data/bq0gvax.html>) in the *Novell eDirectory 8.8 Troubleshooting Guide* (<http://www.novell.com/documentation/edir88/edir88tshoot/data/front.html>).

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 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 Start your Internet browser and enter the URL for iManager.

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

1b Specify your username and password, specify the tree where you want to log in, then click *Login*.

1c Click the *View Objects* button at the top of the iManager page.

1d In tree view, browse to and right-click the desired DriverSet object, then select *Modify Object*.

1e Click the *General* tab, and in the list of valued attributes, click *DirXML-DriverTraceLevel*, then click *Edit*.

1f Ensure that the Trace Level is set to 4, then click *OK*.

1g Repeat [Step 1e](#) and [Step 1f](#) for the *DirXML-XSLTraceLevel* attribute, also setting the trace level to 4.

1h Repeat [Step 1d](#) through [Step 1g](#) for the other driver sets you want to trace.

2 At the 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.

12.13 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).

12.14 A Resource Does Not Migrate to a Peer Cluster

If you cannot migrate a resource from one cluster to a peer cluster, the problem might be caused by one of the following conditions:

- ♦ The resource has not been BCC-enabled.
- ♦ Remote clusters cannot communicate.

See [Section 12.13, “Peer Cluster Communication Is Not Working,”](#) on page 113.

- ♦ 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.

12.15 Resource Cannot Be Brought Online

If you cannot bring a BCC-enabled resource online, it is possible the resource might be set as secondary. If the NCS:BCC State attribute is equal to 1, the resource is set to secondary and cannot be brought online.

On the resource object, change the NCS:BCC State attribute to 0 to set the resource to the primary state. Also, increment the NCS:Revision attribute one number so that Novell Cluster Services recognizes that the resource properties have been updated. See [Step 1 on page 112](#) for an example of how to modify object attributes.

12.16 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:

- ♦ [“Novell eDirectory 8.8” on page 38](#)
- ♦ [“Configuring a BCC Administrator User and Group” on page 46](#)
- ♦ [“Verifying BCC Administrator User Trustee Rights and Credentials” on page 87](#)

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

12.18 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 IP Addresses, LDAP Server, or Administrator Credentials for a Cluster”](#) in the *OES 2 SP1: Novell Cluster Services 1.8.6 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, not in load and unload scripts.
This is the only way to change the IP address on the virtual NCP server object in eDirectory.

12.19 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 either missing or are invalid. The following list shows the required attributes. The format for the list is:

Object in the directory (Object Class Name)

Attribute Name -->The object in the directory the attribute points to

Attribute Name -->The object in the directory the attribute points to

Clustered Volume (Object Class “NCS:Volume Resource”)

NCS:NCP Server -->Virtual NCP Servers

NCS:Volumes-->All Volumes

Virtual NCP Server (Object Class “NCP Server”)

Resource-->Cluster Resource

NCS:NetWare Cluster-->Cluster Object

NCS:Volumes-->All Volumes

Volume (Object Class “Volume”)

nssfsPool-->Pool Object

Host Server-->Virtual NCP Server

Pool (Object Class “nssfsPool”)

Host Server-->Virtual NCP Server

12.20 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.

12.21 Mapping Drives in Login Scripts

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

- ♦ Using an FDN (such as map s:=BCCP_Cluster_HOMES.servers.:shared) to a cluster-enabled volume has been tested and does not work.

When the resource fails over to the secondary cluster, the DN does not resolve to a server/volume that is online. This causes the map command to fail.

- ♦ Using the *SLP Server Name/VOL:shared* syntax has been tested and works.

SLP Server Name is the name being advertised in SLP as specified in the resource load script. This method requires a client reboot.

- ♦ See [TID 10057730 \(http://support.novell.com/docs/Tids/Solutions/10057730.html\)](http://support.novell.com/docs/Tids/Solutions/10057730.html) for information on modifying the server cache Time To Live (TTL) value on the Novell Client™.

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

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

- ♦ Using the *%HOME_DIRECTORY* variable (such as `map u:=%HOME_DIRECTORY`) has been tested and does not work.

When you fail the resource over to the secondary cluster, the *%HOME_DIRECTORY* variable still resolves to the old volume object, and the `map` command fails.

- ♦ Using a temporary environment variable has been tested and does not work.

For example:

```
set FOO=%HOME_DIRECTORY
MAP u:=%FOO
```

- ♦ Using a false volume object along with ICE and LDIF has been tested and works.

- ♦ Create a new volume object that references the real volume object.

The Host Server attribute must point to the virtual NCP server in the primary cluster, and the Host Resource Name attribute must specify the name of the volume. This new volume object can be referred to as a volume reference.

- ♦ All User objects must be modified to have their Home Directory attribute reference the new volume object (volume reference).
- ♦ Use LDIF and ICE in the NSMI script (SAN Array Mapping Information) area.

This script modifies the new volume reference object and updates the Host Server attribute to point to the virtual NCP server in the secondary cluster.

NOTE: Using LDIF/ICE prevents you from using the NSMI script to control the SAN. If you want to use LDIF/ICE and the NSMI script, you must have two NCF files: one for the SAN, and one for LDIF/ICE. The NSMI script must then call each NCF file separately.

See [TID 10057730 \(http://support.novell.com/docs/Tids/Solutions/10057730.html\)](http://support.novell.com/docs/Tids/Solutions/10057730.html) for information on modifying the Server Cache Timeout value on the Novell Client.

A sample NSMI script is included below:

```
#!ICE -b -D LDAP -d cn=root,ou=servers,o=lab -w novell -S LDIF -f
#@ -s0 -w20
version: 1
dn: cn=HOMES_REF, ou=servers,o=lab
changetype: modify
replace: hostServer
hostServer:
cn=BCC_CLUSTER_HOMES_SERVER,ou=From_BCCP,ou=servers,o=lab
```

The first line in the sample script instructs NSMI to run the ICE utility.

- ♦ The `-b` parameter automatically closes the ICE window.

- ♦ The -d parameter is the administrator DN that is used to modify eDirectory.
- ♦ The -w parameter is the password.
- ♦ There must be a trailing space after the -f parameter.

The second line in the sample script includes NSMI-specific options.

- ♦ -s0 causes NSMI to not wait for a signal file.
- ♦ -w20 causes NSMI to wait 20 seconds before proceeding

Failure to add the wait causes the temporary LDIF file to be deleted before ICE can read it. This causes ICE to fail.

12.23 BCC Error Codes

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

Table 12-3 *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.

Error Code Number	Message
1020	CIM Client error.
1021	Error creating a system resource (mutex, semaphore, etc.).
1022	File IO error.
1023	No data.
1024	Not a member of the cluster.
1025	Invalid token in the script.
1026	Invalid or unknown cluster.
1027	The NSMI script is too long. It must be less than 64 KB.
1028	The cluster-enabled pool resource does not contain a volume.
1029	An operation has timed out.
1030	The specified resource is already busy.

The error code numbers are recorded in a log file that you can use to view status changes for BCC.

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.

12.24 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.

Security Considerations

13

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 13.1, “Security Features,” on page 119](#)
- ◆ [Section 13.2, “Security Configuration,” on page 120](#)
- ◆ [Section 13.3, “General Security Guidelines,” on page 124](#)
- ◆ [Section 13.4, “Security Information for Dependent Products,” on page 124](#)

13.1 Security Features

The following table contains a summary of the security features of Novell® Business Continuity Clustering 1.2:

Table 13-1 Business Continuity Clustering 1.2 Security Features

Feature	Yes/No	Details
Users are authenticated	Yes	Administrative users are authenticated via Novell eDirectory™. For information about configuring rights needed by BCC administrators, see Section 4.1.5, “Novell eDirectory 8.8,” on page 38 .
Users are authorized	Yes	Users are authorized via eDirectory trustees.
Access to configuration information is controlled	Yes	Access to the administrative interface is restricted to valid users who have write rights to the configuration files.
Roles are used to control access	Yes	Configurable through iManager.
Logging and/or security auditing is done	Yes	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 usernames and passwords are encrypted. Cluster credentials are stored encrypted in eDirectory.
Security is on by default	Yes	

13.2 Security Configuration

This section provides a summary of security-related configuration settings for Business Continuity Clustering 1.2.

- ♦ [Section 13.2.1, “BCC Configuration Settings,” on page 120](#)
- ♦ [Section 13.2.2, “Changing the NCS: BCC Settings Attributes in the BCC XML Configuration,” on page 121](#)
- ♦ [Section 13.2.3, “Disabling SSL for Inter-Cluster Communication,” on page 122](#)
- ♦ [Section 13.2.4, “Restricting the Network Address for Administration,” on page 124](#)

13.2.1 BCC Configuration Settings

[Table 13-2](#) lists the BCC configuration settings that are security-related or impact the security of BCC:

Table 13-2 BCC Security Configuration Settings

Configuration Setting	Possible Values	Default Value	Recommended Value for Best Security
Inter-cluster communications scheme	HTTP (port 5988) HTTPS (port 5989)	HTTPS	HTTPS
Identity Manager communications	Secure/Non-secure	This is the certificate in the Identity Manager driver setup. If you create the driver with an SSL certificate, it is secure. If not, it is in the clear.	Secure SSL certificates are mandatory for User Synchronization drivers to synchronize between trees.
BCC Administrator user	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 4.3, "Configuring a BCC Administrator User and Group," on page 46.
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

13.2.2 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 Novell Support. Doing so can have adverse affects on your cluster nodes and BCC.

The following XML for the NCS:BCC Settings attribute is saved on the local Cluster object in eDirectory. The BCC must be restarted for changes to these settings to take effect. These are advanced settings that are intentionally not exposed in the BCC plug-in for iManager.

```
<bccSettings>
  <adminGroupName>bccgroup</adminGroupName>
  <authorizationCacheTTL>300</authorizationCacheTTL>
  <cimConnectTimeout>15</cimConnectTimeout>
  <cimReceiveTimeout>30</cimReceiveTimeout>
  <cimSendTimeout>30</cimSendTimeout>
  <idlePriorityThreshold>3</idlePriorityThreshold>
  <initialNormalThreads>3</initialNormalThreads>
  <initialPriorityThreads>2</initialPriorityThreads>
  <ipcResponseTimeout>15</ipcResponseTimeout>
  <maximumPriorityThreads>20</maximumPriorityThreads>
  <minimumPriorityThreads>2</minimumPriorityThreads>
  <resourceOfflineTimeout>300</resourceOfflineTimeout>
  <resourceOnlineTimeout>300</resourceOnlineTimeout>
  <scanForNewDevicesDelay>5</scanForNewDevicesDelay>
</bccSettings>
```

On Linux, the above XML is written to `/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 13-3 provides additional information on each setting:

Table 13-3 BCC XML Settings

Setting	Description	Default Value
<adminGroupName>	The name of the LUM-enabled group that BCC uses on Linux.	bccgroup
<authorizationCacheTTL>	The number of seconds the authorization rights are cached in the BCC OpenWBEM provider.	300 seconds
<cimConnectTimeout>	BCC CIM client connect timeout in seconds.	15 seconds
<cimReceiveTimeout>	BCC CIM client receive timeout in seconds.	30 seconds
<cimSendTimeout>	BCC CIM client send timeout in seconds.	30 seconds
<idlePriorityThreshold>	The number of idle high priority threads before BCC starts killing priority threads.	3
<initialNormalThreads>	The number of normal threads created by BCC.	3
<initialPriorityThreads>	The number of high priority threads created by BCC at startup.	2
<ipcResponseTimeout>	The timeout in seconds that BCC waits for an IPC response.	15
<maximumPriorityThreads>	The maximum number of high priority threads BCC creates.	20
<minimumPriorityThreads>	The minimum number of high priority threads BCC keeps after killing idle high priority threads.	2
<resourceOfflineTimeout>	The number of seconds BCC waits for a resource to go offline during a BCC migrate.	300
<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

13.2.3 Disabling SSL for Inter-Cluster Communication

Disabling SSL for inter-cluster communication should only be done for debugging purposes, and should not be done in a production environment or for an extended period of time.

To turn off SSL for inter-cluster communication, or to specify a different communication port, you need to modify the Novell Cluster Services™ Cluster object that is stored in eDirectory by using an eDirectory management tool such as iManager or ConsoleOne®. See the *Novell iManager 2.7x Administration Guide* (<http://www.novell.com/documentation/imanager27/>) for information on using iManager.

Disabling SSL communication to a specific peer cluster requires changing the BCC management address to the peer cluster. The address is contained in the NCS:BCC Peers attribute that is stored on the NCS Cluster object.

For example, a default NCS:BCC Peers attribute could appear similar to the following example where https:// is assumed and is never specified explicitly:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES_TREE</tree>
  <address>10.1.1.10</address>
</peer>
```

To disable SSL for inter-cluster communication, you would change the <address> attribute to specify http:// with the IP address, as shown in the following example:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES_TREE</tree>
  <address>http://10.1.1.10</address>
</peer>
```

The BCC management address of `chicago_cluster` now specifies non-secure HTTP communication.

The BCC management port can also be changed by modifying the NCS:BCC Peers attribute values. The default ports for secure and non-secure inter-cluster communication are 5989 and 5988 respectively.

For example, if you want to change the secure port on which OpenWBEM listens from port 5989 to port 1234, you would change the <address> attribute value in the above examples to:

```
<peer>
  <cluster>chicago_cluster</cluster>
  <tree>DIGITALAIRLINES_TREE</tree>
  <address>10.1.1.10:1234</address>
</peer>
```

The attribute now specifies that inter-cluster communication uses HTTPS over port number 1234.

The NCS:BCC Peers attribute has a value for each peer cluster in the BCC. Attribute values are synchronized among peer clusters by the BCC-specific Identity Manager driver, so a change to an attribute value on one cluster causes that attribute value to be synchronized to each peer cluster in the BCC.

The changes do not take effect until either a reboot of each cluster node, or a restart of the Business Continuity Clustering software on each cluster node.

[Table 13-4](#) provides an example of possible combinations of scheme and port specifier for the <address> tag for values of the NCS:BCC Peers attribute:

Table 13-4 Example of Scheme and Port Specifier Values for the NCS:BCC Peers Attribute

Value	Protocol Used	Port Used
10.1.1.10	HTTPS	5989

Value	Protocol Used	Port Used
10.1.1.10:1234	HTTPS	1234
http://10.1.1.10	HTTP	5988
http://10.1.1.10:1234	HTTP	1234

13.2.4 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.

13.3 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.

13.4 Security Information for Dependent Products

[Table 13-5](#) provides links to security-related information for other products that impact the security of BCC:

Table 13-5 Security Information for Other Products

Product Name	Links to Security Information
eDirectory	Security for eDirectory is provided by NCI. See the NICI 2.7x Administration Guide (http://www.novell.com/documentation/nici27x/nici_admin_guide/data/a20gkue.html).
Identity Manager	“Security: Best Practices” (http://www.novell.com/documentation/idm36/idm_security/data/b1bsw73.html) in the <i>Identity Manager 3.6x Security Guide</i> .
iSCSI	<p>For information about securing iSCSI target devices on NetWare 6.5 SP8, see “Configuring Access Control to iSCSI Targets” in the <i>NW 6.5 SP8: iSCSI 1.1.3 Administration Guide</i>.</p> <p>For information about securing iSCSI targets and initiators for Linux iSCSI solutions, see “Mass Storage over IP Networks—iSCSI” (http://www.novell.com/documentation/sles10/sles_admin/data/cha_inst_system_iscsi.html) in the <i>SUSE® Linux Enterprise Server 10 Installation and Administration Guide</i>.</p>
Linux User Management (LUM)	See <i>OES 2 SP1: Novell Linux User Management Technology Guide</i> .
Novell Cluster Services for Linux	In the <i>OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide</i> , see <ul style="list-style-type: none">◆ “Configuration Requirements”◆ “Assigning Install Rights for Container Administrators”
Novell Storage Services for Linux	In the <i>OES 2 SP1: NSS File System Administration Guide</i> , see: <ul style="list-style-type: none">◆ “Access Control for NSS on Linux”◆ “File Access for Users”◆ “Securing Access to NSS Volumes, Directories, and Files”◆ “Security Considerations”
OpenWBEM	OpenWBEM should be configured on each node to allow only the necessary users. OpenWBEM by default allows all users. For information, see “Changing the Authentication Configuration” in the <i>OES 2 SP1: OpenWBEM Services Administration Guide</i> .

Console Commands for BCC

A

Novell® Business Continuity Clustering (BCC) provides server console commands to help you perform certain BCC management tasks. Some of the commands can also be used to manage Novell Cluster Services™ clusters.

IMPORTANT: For Novell Cluster Services console commands, see “[Console Commands for Novell Cluster Services](#)” in the *OES 2 SP1: Novell Cluster Services 1.8.6 for Linux Administration Guide*.

- ♦ [Section A.1, “Using Console Commands,” on page 127](#)
- ♦ [Section A.2, “Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh,” on page 130](#)

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 Novell 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 [cluster]</code>	Lets you change the administrator username 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 [<i>resource</i>]	<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> <hr/> <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, make sure to remove the secondary peer clusters from its <i>Assigned</i> list. For information, see Section 11.6, “Disabling BCC for a Cluster Resource,” on page 102.</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> <hr/> <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 [<i>resource</i>]	<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 [<i>source/resource</i>] [<i>destination/nodename</i>]	<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> <hr/> <p>IMPORTANT: Make sure the appropriate resource is offline in all peer clusters in the BCC before executing this command.</p> <hr/> <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 Novell 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 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

There is a framework for customized device scanning that Novell Cluster Services provides 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 currently in a cluster.

By default, the script is empty so that administrators can add the Linux shell commands they need to refresh their systems. Any changes that are made to the script are not overwritten when Novell Cluster Services is upgraded.

- 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.

Setting Up Auto-Failover

B

Auto-failover is available beginning in Novell® Business Continuity Clustering 1.1. To set up the auto-failover feature, you must enable it, then configure the auto-failover settings.

WARNING: Auto-failover is disabled by default and is not recommended. It should only be enabled after a thorough examination and review of your network and geographic site infrastructure. You should seriously consider the adverse conditions that might occur as a result of enabling this feature.

These conditions may include but are not limited to:

- ♦ Data loss at one or more geographic sites
- ♦ Data corruption at one or more geographic sites
- ♦ Data divergence at one or more geographic sites

For example, if there is a loss of communication between two clusters and auto-failover has been enabled and configured, each cluster will assert ownership of BCC-enabled cluster resources. These resources then automatically load on both clusters.

When communication between cluster has been restored, some of the data on each cluster is different. This is called data divergence. Also, the mirroring or synchronization process either fails, or attempts to overwrite any changed data on one cluster. This causes either data loss or data corruption.

-
- ♦ [Section B.1, “Enabling Auto-Failover,” on page 131](#)
 - ♦ [Section B.2, “Creating an Auto-Failover Policy,” on page 132](#)
 - ♦ [Section B.3, “Refining the Auto-Failover Policy,” on page 132](#)
 - ♦ [Section B.4, “Adding or Editing Monitor Configurations,” on page 134](#)

B.1 Enabling Auto-Failover

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

- 1 Start your Internet browser and enter the URL for iManager.
The URL is `http://server_ip_address/nps/iManager.html`. Replace *server_ip_address* with the IP address or DNS name of the server that has iManager and the Identity Manager preconfigured templates for iManager installed.
- 2 Specify your username and password, specify the tree where you want to log in, then click *Login*.
- 3 In *Roles and Tasks*, click *Clusters*, then click the *Cluster Options* link.
- 4 Specify a cluster name, or browse and select the Cluster object.
- 5 Click the *Properties* button, then click the *Business Continuity* tab.
- 6 On the Business Continuity page, click the *Main* link.
- 7 Click the *AutoFailover* button.

- 8 Click the *Auto-Failover* link just under the tabs.
- 9 Select the *Enable Automatic Failover of Business Continuity Cluster Resources* check box, then click *Apply*.
- 10 Continue with [Section B.2, “Creating an Auto-Failover Policy,” on page 132](#) to create a failover policy.

Auto-failover is not completely enabled until you create an auto-failover policy.

B.2 Creating an Auto-Failover Policy

By default, no auto-failover policy exists for BCC. You must create an auto-failover policy for each cluster in your BCC where you want auto-failover enabled. This is required to automatically fail over resources from one cluster to another.

- 1 In iManager, under *Cluster Membership Monitoring Settings*, select a cluster and click the *Edit* link.
- 2 Under *Membership Threshold*, select the *Enable* check box, select either *Percent Fail* or *Nodes Fail*, and specify either the percentage of failed nodes or the number of failed nodes.

The node failure number or percentage you specify must be met for the selected cluster before resources automatically fail over to another cluster.

IMPORTANT: Do not use a membership condition of total node failure (either 100 percent or the total number of nodes); the condition cannot be satisfied because the cluster will not be up to report this state.

If a cluster has been totally downed, you must bring up the master node in the downed cluster, then run the `cluster resetresources` command on that node before you begin manually migrating the BCC-enabled resources to peer clusters.

- 3 Under *Communication Timeout*, select the *Enable* check box and specify the number of minutes that must elapse without any communication between clusters before resources automatically fail over to another cluster.
- 4 Click *OK* to finish editing the policy.
- 5 Click the *Apply* button to save your settings.

B.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 if the connection to the peer cluster is up. Adding this condition changes the behavior of the rule. With this rule, a graceful automatic failover of resources can happen when the connection to the peer cluster is up.

You can also specify or change the criteria for percent or number of nodes that are used to determine if an automatic failover can occur.

IMPORTANT: Do not use a membership condition of total node failure (either 100 percent or the total number of nodes); the condition cannot be satisfied because the cluster will not be up to report this state.

If a cluster has been totally downed, you must bring up the master node in the downed cluster, then run the `cluster resetresources` command on that node before you begin manually migrating the BCC-enabled resources to peer clusters.

You should create a separate rule with a connection down condition. Adding a connection down condition to an existing rule with a condition that tests cluster membership is not recommended. It is highly unlikely that cluster membership information for a specific cluster will be reported to peer clusters when the connection to that specific cluster is down.

For example, a rule might contain only one condition that tests whether a connection to a specific cluster has been down for five or more minutes. Failover occurs when peer clusters agree that the connection to the cluster specified in the rule has been down for five or more minutes. If the peer clusters do not agree about the connection being down (that is, one cluster has a valid connection to the specified cluster), failover does not occur. More complex rules can be constructed that contain multiple conditions.

If previously configured, the fields under *Failover Policy Configuration* should already contain information on the policies that were created in the *Cluster Membership Monitoring Settings* section of the page.

- 1 Under *Failover Policy Configuration*, select a policy and click *Edit* to further refine a rule. Click *Delete* to remove the rule, or click *New* to create a new rule that you can add the additional failover conditions to.
- 2 Select the cluster that you want the rule to apply to, or select *All* to apply the policy to all clusters.
- 3 Under *Conditions*, choose the type of condition and the appropriate values. To add multiple conditions to the rule, click the *Add* button below the condition.

You can use the default setting of *Monitor* if you don't want to apply the cluster up or cluster down criteria to this policy. You can also specify or change the percent or number of nodes criteria that are used to determine if an auto failover can occur.

- 4 Click *Apply* to save your settings.

B.4 Adding or Editing Monitor Configurations

Clicking the *Advanced* button also displays an additional section on this page called *Health Monitor Configuration*. Monitors are an important part of the automatic failover feature, and are separate processes that perform a specialized task to analyze the health of a specific cluster or all clusters in the BCC. These monitors report an indication of health to BCC. BCC in turn uses the reported information to analyze the failover policy to determine if resources should be migrated from a specific cluster. Business Continuity Clustering 1.1 ships with two monitors (nodecnt and 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 Business Continuity Clustering 1.1. Although default values have already been set, you can customize some of the monitor settings for the cluster membership monitors. If you have created your own custom monitor, you can click *New* to add configuration settings to your monitor.

- 1 In iManager, under *Monitor Name* in the *Health Monitor Configuration* section, select a monitor and click *Edit*.
- 2 Under *Clusters*, select the cluster or clusters that you want this monitor to apply to.
- 3 Specify the maximum health indication that the monitor will report.

This value is used when creating a failover policy to validate the rules. This is the maximum value that can be used for the threshold type when you create a failover policy. For example, if you specified percent fail membership monitoring, the maximum values would be 100 (for 100 percent) for the nodepnt monitor. If you specified nodes fail membership monitoring, the maximum value for the nodecnt monitor is the maximum number of nodes permitted in a cluster, which is 32. If you created your own custom monitor, the values could be different.

For the nodepnt and nodecnt monitors, the *Maximum Health Indication* value is for information only, and should not be changed.

- 4 Under *Short Polling Interval*, specify the number of seconds the monitor will wait each time it contacts the cluster or clusters to get health information.
- 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 *Long Polling Interval* is not used with the default nodepnt and nodecnt monitors. This value might be used for some custom monitors.

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.

Configuring Host-Based File System Mirroring for NSS Pools

C

Several methods and scenarios exist for mirroring data between geographically separate sites. Each method has its own strengths and weaknesses. For a Novell® Business Continuity Clustering system, you need to choose either host-based mirroring or storage-based mirroring (also called array-based mirroring) and whether you want the mirroring to be synchronous or asynchronous.

Figure C-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 2 SP1 Linux. 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 C.1, “Creating and Mirroring NSS Pools on Shared Storage,” on page 136](#)
- ♦ [Section C.2, “Creating NSS Volumes,” on page 138](#)
- ♦ [Section C.3, “Novell Cluster Services Configuration and Setup,” on page 138](#)
- ♦ [Section C.4, “Checking NSS Volume Mirror Status,” on page 138](#)

C.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 2 SP1: NSS File System Administration Guide*.
- ◆ Novell CIFS for Linux and Novell AFP for Linux are available in OES 2 SP1 Linux and later. If you plan to mark CIFS or AFP as an advertising protocol for the NSS pool resource, make sure 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 138](#).

- 6 Specify the virtual server name, IP address, and advertising protocols.

When you cluster-enable a pool, a virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name is `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

The *NCP* option is selected by default. NCP™ (NetWare Core Protocol™) is used by NSS to enforce the Novell trustee model of file access for all users of the data, and also as an advertising protocol for NCP clients. Selecting *NCP* causes commands to be added to the pool resource load and unload scripts to activate the NCP protocol on the cluster.

IMPORTANT: Using the *CIFS* check box or the *AFP* check box requires that Novell CIFS for Linux or Novell AFP for Linux is already installed and configured on the nodes in the cluster and nodes in the peer clusters where you want to be able to fail over this pool cluster resource.

Select the *CIFS* check box to add it as an advertising protocol. Specify a CIFS share name for users to access the resource. Selecting *CIFS* causes commands to be added to the pool resource load and unload scripts to activate the Novell CIFS for Linux protocol on the cluster.

Select the *AFP* check box to add it as an advertising protocol. Selecting *AFP* causes commands to be added to the pool resource load and unload scripts to activate the Novell AFP for Linux protocol on the cluster.

- 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 the F3 key.
- 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 2 SP1: NSS File System Administration Guide*.
- 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](#) to [Step 12](#) for each additional shared pool that you want to mirror.
- 14 Continue with [“Creating NSS Volumes” on page 138](#).

C.2 Creating NSS Volumes

After you create the pool and its mirror, you must create an NSS volume on each of the pool resources so that it works properly as a cluster resource. To create an NSS volume on a shared pool, follow the instructions in “[Creating NSS Volumes on a Shared Pool](#)” in the *OES 2 SP1: Novell Cluster Services 1.8.6 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.

C.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 2 SP1: Novell Cluster Services 1.8.6 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.

C.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 137](#).
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 and 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 `cluster scan for new devices` command on a server in each of the peer clusters.

Configuration Worksheet for the BCC Drivers for Identity Manager

D

Use this worksheet to gather the information you need to configure the Novell® Business Continuity Clustering 1.2 drivers for Identity Manager. Repeat this process for each connection between clusters. For information about how to set up the drivers, see [Chapter 8, “Configuring the Identity Manager Drivers for BCC,”](#) on page 65.

- ♦ [Section D.1, “Cluster1 to Cluster2,”](#) on page 141
- ♦ [Section D.2, “Cluster2 to Cluster1,”](#) on page 143

D.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
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
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.	

Driver Set Information	Value
Create a new partition on this driver set	Deselect this option.

Driver Instance for the Connection from Source to Destination

Driver Set Information	Value
------------------------	-------

Pre-configured driver template to import (pick one):

`BCCClusterResourceSynchronization.xml`

`UserObjectSynchronization.xml`

TIP: The User Object Synchronization driver is needed only when creating a driver instance for communications between clusters in different trees. In this case, create one driver instance for resource synchronization and one driver instance for user object synchronization.

Driver name for this driver instance

Example: cluster1_to_cluster2 BCCCR Sync

Example: cluster1_to_cluster2 BCCUO Sync

Default: BCC Cluster Sync

SSL certificate name to use for this driver instance

Example: cluster1_to_cluster2 BCCCR Sync

Example: cluster1_to_cluster2 BCCUO 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)

Example: 2001 (default for user object 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.

Driver Set Information	Value
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.	

D.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
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
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.	

Driver Set Information	Value
Create a new partition on this driver set	Deselect this option.

Driver Instance for the Connection from Source to Destination

Driver Set Information	Value
------------------------	-------

Preconfigured driver template to import (pick the same one that you selected for the Cluster1-to-Cluster2 synchronization):

`BCCClusterResourceSynchronization.xml`

`UserObjectSynchronization.xml`

TIP: The User Object Synchronization driver is needed only when creating a driver instance for communications between clusters in different trees. In this case, create one driver instance for resource synchronization and one driver instance for user object synchronization.

Driver name for this driver instance

Example: cluster2_to_cluster1 BCCCR Sync

Example: cluster2_to_cluster1 BCCUO Sync

Default: BCC Cluster Sync

SSL Certificate name for this driver instance

Example: cluster2_to_cluster1 BCCCR Sync

Example: cluster2_to_cluster1 BCCUO 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

Unique IP port to use for this driver instance (same as you used for the Cluster1-to-Cluster2 driver instance)

Example: 2002 (default for resource sync)

Example: 2001 (default for user object sync)

TIP: Both clusters must use the same port for this connection.

Driver Set Information	Value
Distinguished name of the source cluster for this driver instance	Example: cluster2.clusters.siteB.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.siteB.example
TIP: Browse to select the context.	

Using Dynamic DNS with BCC 1.2

E

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 Novell® Business Continuity Cluster (BCC) 1.2 for Novell Open Enterprise Server 2 SP1 Linux, you can configure BCC to act as a dynamic DNS (Domain Name System) client.

- ♦ [Section E.1, “Requirements and Assumptions,” on page 147](#)
- ♦ [Section E.2, “Configuring the DNS Server for Dynamic DNS,” on page 148](#)
- ♦ [Section E.3, “Configuring the Cluster Resources for Dynamic DNS,” on page 154](#)
- ♦ [Section E.4, “Testing the Dynamic DNS Solution,” on page 159](#)

E.1 Requirements and Assumptions

- ♦ [Section E.1.1, “Software Requirements,” on page 147](#)
- ♦ [Section E.1.2, “DNS Server,” on page 147](#)
- ♦ [Section E.1.3, “TSIG Keys,” on page 148](#)
- ♦ [Section E.1.4, “DNS Record Time-to-Live Values,” on page 148](#)

E.1.1 Software Requirements

Install and configure your business continuity cluster using the following components:

- ♦ Novell Business Continuity Clustering 1.2 for Linux
- ♦ Novell Cluster Services™ 1.8.6 for Linux
- ♦ Novell Open Enterprise Server 2 SP1 for Linux

To integrate dynamic DNS in your BCC solution, use the `bcc_dyn_dns.pl` NSMI script that is included in the BCC 1.2 ISO image in the `<media>/nsmi_scripts/linux` directory.

E.1.2 DNS Server

The examples used in this BCC 1.2 for Linux 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 2 SP1: Novell DNS/DHCP Administration Guide for Linux*.

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 a bit tricky. 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 148](#)
- [Section E.2.2, “Configuring the DNS Server with the Public Key,” on page 149](#)
- [Section E.2.3, “Configuring the DNS Server Zones,” on page 150](#)
- [Section E.2.4, “Testing the DNS Server,” on page 150](#)

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
```

where `60303` represents a randomly generated number created by the utility.

- 3 Store these files in a secure location, then continue with [Section E.2.2, “Configuring the DNS Server with the Public Key,”](#) on page 149.

The DNS administrator uses these keys to configure your master DNS server.

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 148. 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 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 150.

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 make sure 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 150.

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 Make sure 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 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
wkstn1.clusters.site1.company.com

;; ADDITIONAL SECTION:
wkstn1.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
wkstn1.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 to your home directory on the Linux client computer (such as `/home/yourhomedir/`) the public and private keys created in [Section E.2.1, “Creating the TSIG Keys for DNS Server Authentication,”](#) on page 148.

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.sitel.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
wkstnl.clusters.sitel.company.com.

;; ADDITIONAL SECTION:
wkstnl.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
wkstnl.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 make sure 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
wkstnl.clusters.sitel.company.com.
root.wkstnl.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 154.

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 155](#)
- ♦ [Section E.3.2, “Public and Private Keys,” on page 157](#)
- ♦ [Section E.3.3, “Testing the Perl Wrapper Script,” on page 158](#)

E.3.1 Modifying the BCC Load Script

For your convenience, the `bcc_dyn_dns.pl` NSMI script is included in the BCC 1.2 ISO image in the `<media>/nsmi_scripts/linux` directory. This is the script you need as the wrapper for the `nsupdate` utility. You can add the script to the BCC load script without modifications.

NOTE: The `bcc_dyn_dns.pl` NSMI script cannot be executed directly by a Perl interpreter because the interpreter does not understand the variables used in lines 62–66 of the script. If you need to test the script outside of the BCC environment, substitute values for the variables. Put the variables back in before using the script in BCC.

- 1 Open iManager, then log in as the BCC Administrator user.
- 2 In *Roles and Tasks*, select *Clusters > Cluster Options*.
- 3 Select the Cluster object of the secondary peer cluster in the business continuity cluster environment for the cluster resource. (That is, choose a peer cluster where the resource is not currently assigned and running.)
- 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.

Clusters > Cluster Options > Resource Properties

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	
Name	Value
No items	

CIM Enabled Synchronous Edit Flags Flags:

OK Cancel Apply

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, Username, and Password controls are all disabled.
Name	Specify the name of the script that is shown in the BCC log files (such as in <code>/var/log/messages</code>). 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.
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. This field is not needed for the dynamic DNS script.
Secure	This field is used by CIM or SMI-S enabled scripts. This field is not needed for the dynamic DNS script.
Username	This field is used by CIM or SMI-S enabled scripts. This field is not needed for the dynamic DNS script.
Password	This field is used by CIM or SMI-S enabled scripts. This field 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 Novell Support, so leave it disabled (deselected).

7 Specify the Script Parameters.

Click *New* to insert a new editable row into the Script Parameters table. The left field is the parameter name and the right field is the parameter value. When you are done adding the name/value pair, click *OK* to save the parameter.

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.
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.site1.company.com.+157+60303.private</code>	Specify the location of the private key file. Remember that the public key must be in this same directory.

- 8 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 `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 148](#). The `nsupdate` utility really 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.** Although this is simple and relatively quick to do initially, it is a maintenance nightmare. 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.** 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.

- 1 Create a 10 MB shared volume with the Ext3 file system mounted at `/mnt/bcc-master`.
For example, name the resource `bcc-master`. Make sure it is configured with the *Resource Follows Master* enabled.
For information about creating a shared Linux POSIX file system, see “[Configuring Cluster Resources for Shared Linux POSIX Volumes](#)” in the *OES 2 SP1: Novell Cluster Services 1.8.6 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 Make sure 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 manually 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 – 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 modify your DNS server. You can test the results as outlined in [Section E.2.4, “Testing the DNS Server,”](#) on page 150.

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 onlined 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 step in [Section E.2.4, “Testing the DNS Server,” on page 150](#).

- 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 onlined 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 step in [Section E.2.4, “Testing the DNS Server,” on page 150](#).

Using Virtual IP Addresses with BCC 1.2

F

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 Novell® Business Continuity Cluster (BCC) 1.2 for Novell Open Enterprise Server 2 SP1 Linux, 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 161](#)
- ♦ [Section F.2, “Virtual IP Address Benefits,” on page 162](#)
- ♦ [Section F.3, “Planning a Virtual IP Network Implementation,” on page 163](#)
- ♦ [Section F.4, “Configuring a Virtual Router with OSPF,” on page 165](#)
- ♦ [Section F.5, “Configuring Virtual IP Addresses,” on page 166](#)

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 161](#)
- ♦ [Section F.1.2, “Host Mask,” on page 161](#)
- ♦ [Section F.1.3, “Internal Router,” on page 161](#)

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. For information about configuring and using Quagga, see the [Quagga Routing Software Suite Documentation Web site \(http://www.quagga.net/docs.php\)](http://www.quagga.net/docs.php).

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 162](#)
- ♦ [Section F.2.2, “Unlimited Mobility,” on page 163](#)
- ♦ [Section F.2.3, “Automatic Name Resolution,” on page 163](#)

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 if a destination is reachable or not. However, with secondary IP addresses, you know only if 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 side effects that directly follow from the highly reachable nature of virtual IP addresses:

- ♦ A multihomed 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 164](#)
- ♦ [Section F.3.2, “LAN Routers,” on page 164](#)
- ♦ [Section F.3.3, “Internal Router,” on page 164](#)

- ♦ [Section F.3.4, “IP Addresses for BCC-Enabled Cluster Resources,” on page 164](#)
- ♦ [Section F.3.5, “Host Mask,” on page 165](#)

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 164](#). 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. Make sure 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 Software Suite package.
- 3** Edit the configuration files for the Quagga Routing Software Suite.

For information, see the [Quagga Routing Software Suite Documentation Web site \(http://www.quagga.net/docs.php\)](http://www.quagga.net/docs.php).

- 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 163](#), 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 165](#).

- 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 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/ncsfncs
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 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/ncsfncs
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 previous steps for the remaining cluster resources.

Removing Business Continuity Clustering Core Software



If you need to uninstall Novell® Business Continuity Clustering, you can do so by using the Business Continuity Clustering installation program.

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

```
yast2 novell-bcc
```
- 3 Deselect the *Install Core Business Continuity Clustering Software and Configure Core Software* option, select *Remove Installed Software*, then click *Next*.

Documentation Updates



This section contains information about documentation content changes made to the *Novell® Business Continuity Clustering 1.2 Administration Guide for Linux* since the initial 1.2 release. If you are an existing user, review the change entries to readily identify modified content. If you are a new user, simply read the guide in its current state.

This document was updated on the following dates:

- ♦ [Section H.1, “February 18, 2010,” on page 171](#)
- ♦ [Section H.2, “February 7, 2010 \(BCC 1.2.0 Patch for OES 2 SP1 Linux\),” on page 171](#)
- ♦ [Section H.3, “January 7, 2010,” on page 172](#)
- ♦ [Section H.4, “September 9, 2009,” on page 174](#)
- ♦ [Section H.5, “August 14, 2009,” on page 174](#)
- ♦ [Section H.6, “May 11, 2009,” on page 175](#)
- ♦ [Section H.7, “April 28, 2009,” on page 176](#)

H.1 February 18, 2010

Updates were made to the following section. The changes are explained below.

- ♦ [Section H.1.1, “Updating \(Patching\) BCC 1.2.0 on OES 2 SP1 Linux,” on page 171](#)

H.1.1 Updating (Patching) BCC 1.2.0 on OES 2 SP1 Linux

Location	Change
Step 5 on page 56	Restart Tomcat by entering the following command at a terminal console prompt:
Step 5 on page 57	<code>rcnovell-tomcat5 restart</code>

H.2 February 7, 2010 (BCC 1.2.0 Patch for OES 2 SP1 Linux)

Updates were made to the following sections. The changes are explained below.

- ♦ [Section H.2.1, “Updating \(Patching\) BCC 1.2 on OES 2 SP1 Linux,” on page 171](#)
- ♦ [Section H.2.2, “What’s New for BCC 1.2,” on page 172](#)

H.2.1 Updating (Patching) BCC 1.2 on OES 2 SP1 Linux

This section is new.

H.2.2 What's New for BCC 1.2

Location	Change
Section 2.1, "BCC 1.2.0 Patch (January 2010)," on page 23	This section is new. A patch for BCC 1.2 for OES 2 SP1 Linux is available in the OES 2 SP1 patch channel.

H.3 January 7, 2010

Updates were made to the following sections. The changes are explained below.

- ◆ [Section H.3.1, "Configuring BCC for Peer Clusters," on page 172](#)
- ◆ [Section H.3.2, "Console Commands for BCC," on page 172](#)
- ◆ [Section H.3.3, "Installing Business Continuity Clustering," on page 173](#)
- ◆ [Section H.3.4, "Managing a Business Continuity Cluster," on page 173](#)
- ◆ [Section H.3.5, "Removing Business Continuity Clustering Core Software," on page 173](#)
- ◆ [Section H.3.6, "Troubleshooting Business Continuity Clustering," on page 173](#)
- ◆ [Section H.3.7, "Upgrading Identity Manager Nodes to Identity Manager 3.6.1," on page 173](#)
- ◆ [Section H.3.8, "Using Virtual IP Addresses with BCC 1.2," on page 174](#)

H.3.1 Configuring BCC for Peer Clusters

Location	Change
Chapter 9, "Configuring BCC for Peer Clusters," on page 79	IMPORTANT: Identity Manager must be configured and running on one node in each peer cluster before configuring clusters for business continuity. Make sure 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 8, "Configuring the Identity Manager Drivers for BCC," on page 65 .
Section 9.4, "Adding Storage Management Configuration Information," on page 83	Added guidelines for adding Perl scripts to the load script and unload script of a BCC-enabled cluster resource.

H.3.2 Console Commands for BCC

Location	Change
Section A.2, "Setting Up Linux Scan Commands in /opt/novell/ncs/bin/device_scan.sh," on page 130	This section is new.

H.3.3 Installing Business Continuity Clustering

Location	Change
“Identity Manager Engine and eDirectory Driver” on page 41	Each Identity Manager 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.

H.3.4 Managing a Business Continuity Cluster

Location	Change
Chapter 10, “Managing a Business Continuity Cluster,” on page 89	IMPORTANT: Identity Manager must be configured and running on one node in each peer cluster before any BCC enabled cluster resource changes are made. Make sure 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 8, “Configuring the Identity Manager Drivers for BCC,” on page 65.
Section 10.1.1, “Understanding BCC Resource Migration,” on page 89	The scan for devices follows the unload script and load script runs.

H.3.5 Removing Business Continuity Clustering Core Software

This section is new.

H.3.6 Troubleshooting Business Continuity Clustering

Location	Change
Section 10.16 Slow Failover	This section was removed. It applied only to BCC 1.1 for NetWare.

H.3.7 Upgrading Identity Manager Nodes to Identity Manager 3.6.1

Location	Change
Section 6.1, “Upgrading to 32-Bit Identity Manager 3.6.1,” on page 59	It is not necessary to re-create the BCC drivers if you perform an in-place upgrade of 32-bit Identity Manager 3.6 to 32-bit Identity Manager 3.6.1.
Section 6.2, “Upgrading to 64-Bit Identity Manager 3.6.1,” on page 59	Upgrading to 64-bit Identity Manager 3.6.1 requires a rebuild of the 64-bit cluster node and a re-creation of the BCC drivers.

H.3.8 Using Virtual IP Addresses with BCC 1.2

Location	Change
Appendix F, “Using Virtual IP Addresses with BCC 1.2,” on page 161	This section is new.

H.4 September 9, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section H.4.1, “Configuring BCC for Cluster Resources,” on page 174](#)
- ◆ [Section H.4.2, “Converting BCC Clusters from NetWare to Linux,” on page 174](#)

H.4.1 Configuring BCC for Cluster Resources

Location	Change
Section 11.1.4, “Cluster Resources for OES 2 Linux Services,” on page 98	Fixed a broken link.

H.4.2 Converting BCC Clusters from NetWare to Linux

Location	Change
Section 7.2, “Prerequisites for Converting from NetWare to Linux,” on page 61	Fixed broken links.

H.5 August 14, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section H.5.1, “Console Commands for BCC,” on page 175](#)
- ◆ [Section H.5.2, “Installing Business Continuity Clustering,” on page 175](#)
- ◆ [Section H.5.3, “Upgrading Identity Manager Nodes to Identity Manager 3.6.1,” on page 175](#)
- ◆ [Section H.5.4, “What’s New for BCC 1.2,” on page 175](#)

H.5.1 Console Commands for BCC

Location	Change
Appendix A, "Console Commands for BCC," on page 127	Updated to add a link to Novell Cluster Services console commands.

H.5.2 Installing Business Continuity Clustering

Location	Change
Section 4.1.7, "Identity Manager 3.6 Bundle Edition," on page 40	Identity Manager 3.6.1 supports the 64-bit OES 2 SP2 operating system.
Section 4.1.16, "Web Browser," on page 45	This section is new.

H.5.3 Upgrading Identity Manager Nodes to Identity Manager 3.6.1

This section is new.

H.5.4 What's New for BCC 1.2

Location	Change
Section 2.2, "Identity Manager 3.6.1 Support (June 2009)," on page 24	This section is new.
	Identity Manager 3.6.1 now supports installation on 64-bit OES 2 SP1 Linux operating systems.

H.6 May 11, 2009

Updates were made to the following sections. The changes are explained below.

- ♦ [Section H.6.1, "Converting BCC Clusters from NetWare to Linux," on page 175](#)

H.6.1 Converting BCC Clusters from NetWare to Linux

Location	Change
Chapter 7, "Converting BCC Clusters from NetWare to Linux," on page 61	Updated for clarity.

Location	Change
Chapter 7, “Converting BCC Clusters from NetWare to Linux,” on page 61	Updated for clarity.

H.7 April 28, 2009

Updates were made to the following sections. The changes are explained below.

- ◆ [Section H.7.1, “Converting BCC Clusters from NetWare to Linux,” on page 176](#)

H.7.1 Converting BCC Clusters from NetWare to Linux

Location	Change
Chapter 7, “Converting BCC Clusters from NetWare to Linux,” on page 61	Updated for clarity.
Chapter 7, “Converting BCC Clusters from NetWare to Linux,” on page 61	Updated for clarity.