Novell
Cluster Services™

1.7

RESOURCE CONFIGURATION GUIDE

www.novell.com

Novell®

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.

You may not use, export, or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside.

Copyright ©1999- 2005 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. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at http://www.novell.com/company/legal/patents/ and one or more additional patents or pending patent applications in the U.S. and in other countries.

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

www.novell.com

Novell Cluster Services 1.7 Resource Configuration Guide
March 31, 2005

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

ConsoleOne is a registered trademark of Novell, Inc. in the United States and other countries.

eDirectory is a trademark of Novell, Inc.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Client is a trademark of Novell, Inc.

Novell Cluster Services is a trademark of Novell, Inc.

Novell Storage Services is a trademark of Novell, Inc.

Third-Party Materials

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

Contents

About This Guide

Before running any application on servers in a Novell® Cluster Services™ cluster, the application must be installed and configured properly to take advantage of shared storage and to achieve high availability. This document provides information to help you set up and configure some of the applications commonly used in a NetWare® 6.5 environment. See the Novell Cluster Services Overview and Installation documentation for cluster installation and configuration, and see your application documentation for specific application installation and setup instructions.

Documentation Updates

For the most recent version of the *Novell Cluster Services 1.7 Resource Configuration Guide*, see the latest NetWare 6.5 documentation (http://www.novell.com/documentation/lg/nw65/index.html).

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as UNIX*, should use forward slashes as required by your software.

User Comments

We want to hear your comments and suggestions about this manual and the other documentation included with NetWare 6.5. To contact us, use the User Comments feature at the bottom of any page in the online documentation.

1 Configuring NSS Mirroring with Novell Cluster Services

NSS mirroring is a checkpoint based synchronous mirroring solution. Data blocks are written synhronously to multiple storage devices. It is an alternative to SAN array based synchronous replication options.

Hardware configuration and placement for NSS mirroring can vary depending on geography and fault tolerance requirements. For example, the following figure depicts a simple hardware configuration in which one side of the mirrored NSS volume is located in a separate building from the rest of the cluster hardware. Should a disaster occur in one building, data would still be safe on the mirrored NSS volume in the other building.

Figure 1   Single Cluster with Mirrored NSS Volumes in Separate Buildings
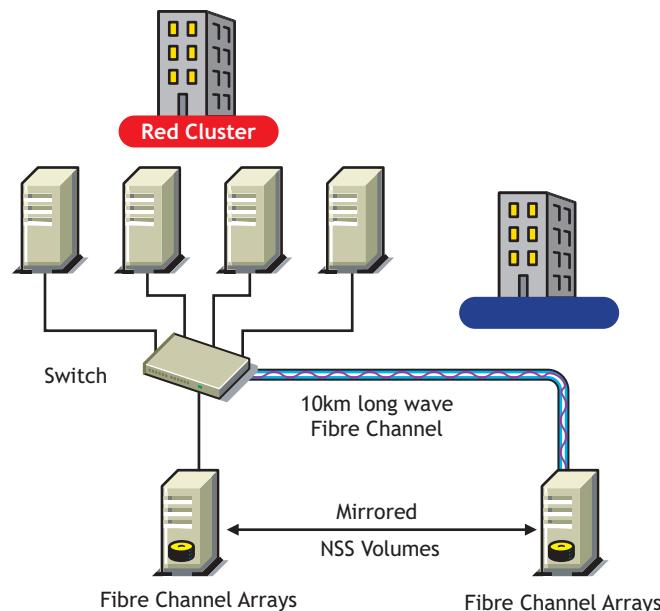
The following figure depicts a more complex hardware configuration in which two clusters are placed in separate buildings, but each has one side of a mirrored NSS volume in the building with the other cluster. Should a disaster occur in one building, data would still be safe and immediately accessible on the mirrored NSS volume in the other building.

Figure 2   Two Clusters with Mirrored NSS Volumes in Separate Buildings
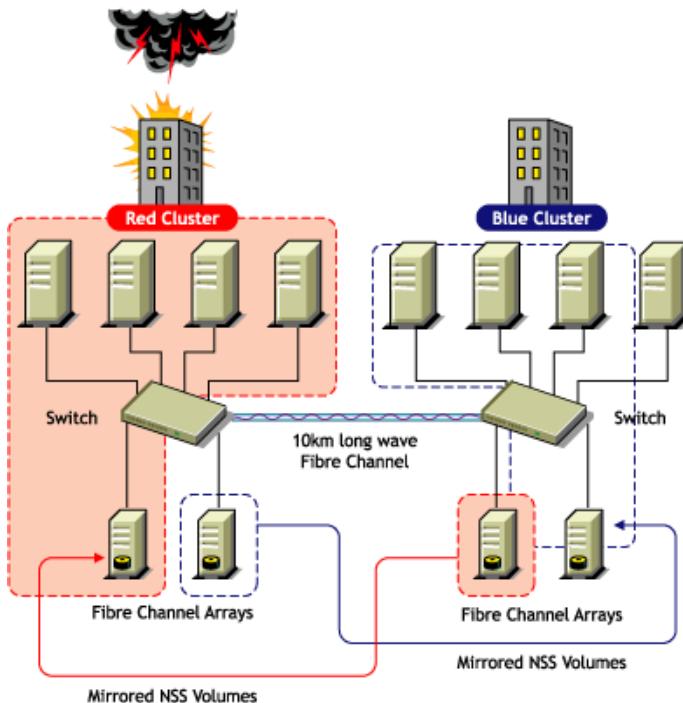
Configure NSS Mirroring

NSS partitions must be mirrored after they are created. If you have an existing partition that you want to mirror, you can either create another partition of equal size on another device to mirror the first partition to, or let the mirroring software automatically create another partition of equal size on another device.

Novell® Cluster Services™ should be installed and running prior to creating and mirroring partitions on shared storage.

When you create a Novell Cluster Services system that utilizes shared storage space (a Storage Area Network or SAN), it is important to remember that all servers attached 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 noncluster servers.

Create and Mirror NSS Partitions on Shared Storage

Prior to creating and mirroring NSS partitions on shared storage, ensure that you have

- NetWare® 6.5 installed on all servers that will be part of the cluster

- All servers in the cluster connected to a shared storage system

- One or more drive arrays configured on the shared storage system

- At least 15 MB of free space on the shared storage system for a special cluster partition

To create and mirror NSS partitions complete the following:

1 Start NSSMU by entering NSSMU at the server console of a cluster server.

2 Select Partitions from the NSSMU main menu.

3 Press the Insert key and select the device on your shared storage system where you want to create a partition.

If a device is marked as sharable for clustering, all partitions on that device will automatically be sharable.

Device names are not changeable and might be labeled something like 0x2 or 0x1.

If Cluster Services was previously installed and shared disk partitions were already created, the Partitions List will include this information.

4 Select NSS as the partition type, then specify the parition size and, if desired, an NSS pool name and label.

If you specify a pool name, a pool by that name will automatically be created on the partition. If no pool name is specified, you will have to create a pool on the partition later.

4a If you chose to create a pool, choose whether you want the pool to be activated and cluster enabled when it is created.

The Activate on Creation feature 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 will have to manually activate it later before it can be used.

The Cluster Enable on Creation feature 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 . If you want to cluster enable the pool at a later date, see the Novell Cluster Services 1.7 Administration Guide for more information.

4b Specify the Virtual Server Name, IP Address, Advertising Protocols and, if necessary, the CIFS Server Name.

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

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

You can select one or all of the advertising protocols. NCP is the protocol used by Novell clients, CIFS is the protocol used by Microsoft clients, and AFP is the protocol used by Macintosh clients. Selecting any of the protocols will cause lines to be added to the pool resource load and unload scripts to activate the selected protocols on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to all your clients.

If you select CIFS as one of the protocols, a CIFS Server Name is also required. This is the server name CIFS clients see when they browse the network. A default server name is listed, but you can change the server name by editing the text in the field.

4c Select Create to create and cluster enable the pool.

5 Select the partition you want to mirror (this may be the partition you just created) and press the F3 key.

6 Select the device with free space or the partition you want to mirror to, and then select YES to mirror the partition.

To ensure disaster recovery, the device you select to mirror should be in another storage array.

Create an NSS Pool and Volumes

Once an NSS partition has been created and mirrored, if you have not already done so, you must create an NSS pool and volume on that partition. To do this, follow the instructions in ""Create NSS Pools" " in the "Installation and Setup" section of the Novell Cluster Services Overview and Installation documentation.

Novell Cluster Services Configuration and Setup

After configuring NSS mirroring and creating a pool and volume on the mirrored NSS partition, if you did not cluster enable the NSS pool on the mirrored partition when you created it, do so by following the instructions in the "Installation and Setup" section of the Novell Cluster Services Overview and Installation documentation.

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.

Check NSS Volume Mirror Status

Once 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 **volumes** at the server console.

2 Check the mirror status of the mirrored partition by entering **mirror status** at the server console of the server where the NSS pool on the mirrored partition is active.

After entering mirror status, you should see a message indicating that mirror status is 100 percent or a message indicating that the mirrored object is fully synchronized.

3 Migrate the pool to another server in the cluster and again check to ensure the volumes on the pool are mounted by entering **volumes** at the server console.

4 Check the mirror status of the partition again by entering **mirror status** at the server console.

Additional Novell Cluster Services Configuration Information

For additional information on setting up and configuring Novell Cluster Services, see the *Novell Cluster Services 1.7 Administration Guide*.

2 Configuring Web Search and Novell Cluster Services

The NetWare Web Search Server Synchronization feature provides both a fault-tolerant and load-balanced environment. Therefore, it is unnecessary to set up Novell® Cluster Services™ for use with Web Search. However, if you want to include failover capabilities, then you might consider using it.

Setting Up Web Search with Novell Cluster Services

Installing and configuring NetWare™ Web Search Server for use with Novell Cluster Services requires three primary steps:

1. Setting up your clustered environment, which includes the Apache Web server and the Tomcat servlet container (both of these are typically installed and configured by default).

2. Installing Web Search Server to the same shared volume from each of your clustered servers.

3. Creating virtual search servers using the DNS name and IP address of your cluster server.

Step One: Setting Up Your Clustered Environment

Before NetWare Web Search Server will run in a clustered environment, you must install and configure the following components on a server in your cluster.

❑ Novell Cluster Services

    See ""Installation and Setup"".

❑ Apache Web Server

    See ""Apache with Novell Cluster Services"".

❑ Tomcat Application Server

    See ""Configuring Tomcat and Novell Cluster Services"".

Step Two: Installing Web Search Server to a Shared Volume

Install Web Search Server (using the NetWare 6.5 Installation CD) to the same shared volume in your clustered environment. For information about setting up a shared volume in a clustered environment, see ""Cluster Enable Pools and Volumes"" in the *Novell Cluster Services 1.7 Administration Guide*.

Installing Web Search Server to each server in your cluster allows the installation program to correctly register Web Search with each of your server's Web and application servers.

Installing it to the same shared volume centralizes all of the required templates, indexes, and configuration settings for use by all servers in your cluster. Each time you create a new virtual

search server, default settings and templates are drawn from the shared volume. Future upgrades or patches to the Web Search Server need to be installed only to the shared volume and can be done from any server in your cluster.

Step Three: Creating Virtual Search Servers

You can now begin creating virtual search servers from any of the Web Search Servers in your cluster.

When prompted to provide a name and alias for your new virtual search server, enter the DNS name and IP address (alias) of your cluster server. The Path field can be left blank because it is already known by Web Search Server.

If You Are Not Using a Shared Volume

We recommend that you use a shared volume. However, if for some reason your clustered environment is not utilizing a shared volume, you can install Web Search Server to each of your servers and then use Web Search Synchronization to keep the configuration settings and indexes up-to-date across all of your servers.

If you do this, each time you upgrade or apply a patch to the Web Search Server, you will need to repeat the upgrade installation or patch on each of the servers in your cluster.

3 Configuring Tomcat and Novell Cluster Services

Tomcat is a JSP-supported servlet engine that, by default, is installed with NetWare® 6. Tomcat runs on a Web server and compiles JSP (Java* Server Page) files whose output is HTML code when displayed in a browser.

If your Web server is running Tomcat in a clustered environment, some addtional configuration is required.

This documentation assumes that you already have Novell® Cluster Services™ installed and running on your network.

Configuring Tomcat

1 Locate the class files that are required to run your Web site, and then copy these files to the same directory on each Web server in the cluster.

This directory can be anywhere that you choose. For example:

SYS:*Your_Web_Site_Name*

2 Choose any Web server in the cluster and edit the sys:\tomcat\33\bin\tomcat33.ncf file by adding a new line to the -classpath parameter. This new line should point to the directory where you copied your class files in Step 1.

For example:

envset TOMCAT CLASSPATH=$TOMCAT CLASSPATH; SYS:*Your_ Web_Site_Name*

3 Copy the edited .ncf file (from Step 2) to the sys:\tomcat\33\bin directory of every Web server in the cluster.

4 Create an NWAPPS-*any_name*.XML file in the sys:\tomcat\33\conf directory of any Web server in the cluster, and then put the path to your Web site in this file.

TIP: This path sets the Tomcat context. Tomcat will service only those requests that originate at or below this context.

For example, if you put all of your Web and JSP files in the same directory, which you called SYS:*Your_Web_Site_Name*, you would enter a slash for the Context path and `SYS:\Your_Web_Site_Name` for the docBase.

The context path is the path that you need in order to access your Web files from a browser and the docBase is the path that you need in order to access the Web files on your hard drive.

To see an example of what this .xml file should contain, go to the CONF directory of your NetWare 6 server, and then open the apps-example.xml file. If you need more explanation of how to use the elements in the example .xml file, see the FAQ document located in the sys:\tomcat\33\doc directory.

5 Copy the .xml file (from Step 4) to the sys:\tomcat\33\conf directory of every Web server in the cluster.

6 At the server console of every Web server in the cluster, enter `tomcat33 -stop` to stop the Tomcat service.

7 Wait approximately 15 seconds for the cleanup operations to finish.

8 At the server console of every Web server in the cluster, enter `tomcat33` to restart the Tomcat service.

Tomcat is now configured to run with Novell Cluster Services.

4 Apache with Novell Cluster Services

There are obvious advantages to clustering your Apache Web servers. Not only will your customers appreciate the high availability of your Web content, but they will also no longer bear the brunt of a Web server failing or going down for maintenance. With Novell ®Cluster Services™, if your Web server fails, you can use any of the servers in the cluster to host your Web site, which results in virtually zero down-time for your customers.

Before Apache and Novell Cluster Services can work together, there is additional configuration that needs to take place when you are installing and configuring both Novell Cluster Services *and* Apache.

Apache Configuration

When you install NetWare® 6 or 6.5, Apache is installed by default. There is one exception to this rule: if you select the Pre-Migration server option while installing Netware 6, Apache is *not* installed by default.

Novell Cluster Services Configuration and Setup

1 Create an NSS pool or use an existing pool.

For more information on creating pools, see the *Novell Cluster Services Overview and Installation Guide* (http://www.novell.com/documentation/lg/ncs6p/index.html).

2 Create an NSS volume on a pool.

If a pool already exists, make sure there is enough space on the pool for the NSS volume that you plan to make.

For more information on creating NSS volumes, see the *Novell Cluster Services Overview and Installation Guide* (http://www.novell.com/documentation/lg/ncs6p/index.html).

3 Cluster enable the pool you created in Step 2.

Cluster enabling the pool automatically custer enables the NSS volume that you created on the pool.

When you cluster enable a pool, this allows you to access the cluster enabled volumes through a NetWare client. This step involves assigning an IP address. Remember this IP address because you will need it later in the configuration.

For more information, see the *Novell Cluster Services Overview and Installation Guide* (http://www.novell.com/documentation/lg/ncs6p/index.html) > Installation and Setup > Setting Up Cluster Services.

4 Put your Web content on the cluster-enabled volumes that you enabled in in Step 3.

5 Choose a server in the cluster and edit the httpd.conf file, located in the sys:\apache2\conf (or, if you are running NetWare 6, sys:\apache\conf) directory.

Make the following two modifications to the httpd.conf file:

- Set the value of the DocumentRoot directive to the directory where you placed your Web content in Step 4. (The default value of this parameter is sys:\Apache2\htdocs, or on NetWare 6, sys:\apache\htdocs.)

  For example, if you placed your Web content in sys:\Apache2\htdocs, the DocumentRoot directive would be:

  DocumentRoot "sys:\Apache2\htdocs"

- Uncomment the Listen parameter, and then add the IP address that you assigned to your cluster-enabled volumes in Step 3 to that parameter.

6 Copy the edited httpd. conf file to the sys:\apache\conf directory on every server in the cluster.

5 Configuring DHCP with Novell Cluster Services

Benefits of Configuring DHCP for High Availability

A NetWare® 6.5 DHCP server automatically assigns IP addresses and other configuration information to clients upon request or when the clients are restarted. If for some reason the NetWare 6.5 DHCP server is not accessible, clients lose their ability to connect to the network because they cannot obtain an IP address. Configuring DHCP with Novell® Cluster Services™ ensures the IP address range required by users to connect to the network is highly available. This is possible because DHCP configuration information and IP address ranges are stored in Novell eDirectory™ and the DHCP server is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services.

DHCP Installation and Configuration

Novell Cluster Services (NCS) 1.7 must be installed before configuring DHCP. NCS 1.7 provides a DHCP resource template, which facilitates configuring DHCP in a cluster environment. NetWare 6.5 includes NCS 1.7 and licenses for two cluster nodes.

Before running DHCP with Novell Cluster Services, DHCP must be configured properly on one server in the cluster. To see DHCP documentation for instructions on installing and configuring DHCP with NetWare 6.5, go to the *Novell DNS/DHCP Services for NetWare Administration Guide for NetWare 6.5*.

DHCP server software is included with NetWare 6.5. You must select the Customized NetWare Server installation option and check the Novell DNS/DHCP Services component check box or the DNS/DHCP Pre-Configured Server option to install DHCP during the NetWare 6.5 installation. You cannot install DHCP with the Basic NetWare File Server installation option. You will need to configure DHCP after your NetWare 6.5 server is installed. And, as stated earlier, you should configure DHCP on only *one* server in the cluster.

Because all DHCP configuration information and IP address ranges are stored in eDirectory, no shared storage is required to run DHCP with Novell Cluster Services.

Any server in the cluster that will run DHCP must have at least a Read/Write eDirectory replica. This allows Novell Cluster Services to modify DHCP-specific eDirectory objects regardless of which server in the cluster holds the objects.

Novell Cluster Services Configuration and Setup

After DHCP is properly configured on one server in your cluster, you must create and configure a DHCP resource in Novell Cluster Services. This includes configuring DHCP load and unload scripts, setting DHCP Failover and Failback modes, and assigning the DHCP resource to specific servers in your cluster.

Creating a DHCP Cluster Resource

Novell Cluster Services includes a DHCP resource template, which greatly simplifies the process for creating a DHCP cluster resource. Much of the DHCP cluster resource configuration is perfomed automatically by the DHCP resource template.

To create a cluster resource:

1 In ConsoleOne®, browse to and select the Cluster object.

2 Click File > New > Cluster > Cluster Resource.

3 Enter a name for the DHCP cluster resource.

4 Enter the DHCP template name in the Inherit from Template field, or browse and select it from the list.

5 Check the Define Additional Properties check box, click Create, and then continue with

The DHCP resource template configures the DHCP resource by automatically creating DHCP load and unload scripts, setting failover and failback modes, and assigning DHCP as a resource to all nodes in the cluster.

Editing DHCP Load and Unload Scripts

The DHCP resource template automatically creates a load and unload script to start and stop the DHCP service on servers in your cluster. The load script contains commands that you must customize for your specific DHCP configuration.

To edit the DHCP load script:

1 In ConsoleOne, right-click the DHCP Resource object and then click Properties.

If you just created the DHCP resource object and checked the Define Additional Properties check box, the DHCP Resource object should already be displayed.

2 Click the Scripts tab on the property page and select Cluster Resource Load Script.

The load script will appear similar to the following example:

```
CLUSTER DHCP CN=SERVER.O=ORG.T=TREE
dhcpsrvr --servaddr=A.B.C.D
```

The CLUSTER DHCP command is necessary to pass DHCP-specific configuration information to another server in the cluster if the server where DHCP is running fails.

The DHCPSRVR command and command switches are used to start the DHCP server (dhcpsrvr.nlm) on a NetWare server.

3 Edit the CLUSTER DHCP command to add the context of the Server object for the server in your cluster that has DHCP configured.

For example, if the server context is CN=PROD1.O=SALES.T=ACME_TREE, the line in the load script would appear as follows:

```
CLUSTER DHCP CN=PROD1.O=SALES.T=ACME_TREE
```

The server context specified in the load script does not change even though the load script will run on another server in the cluster during a failover.

4 Edit the DHCPSRVR command to assign an IP address to the DHCP cluster resource.

For example, if the IP address for the DHCP cluster resource is 111.152.153.10, the line in the load script would appear as follows:

```
dhcpsrvr --servaddr=111.152.153.10
```

The IP address for the DHCP cluster resource allows clients to reconnect to that address regardless of which server is hosting it.

5 Click the Unload Script tab on the property page to view the unload script.

The UNLOAD DHCPSRVR command is by default the only command in the unload script.

Setting DHCP Start, Failover, and Failback Modes

The DHCP resource template sets the DHCP resource Start Mode and Failover Mode to AUTO, and the Failback Mode to DISABLE.

If the DHCP resource Start Mode is set to AUTO, the DHCP server automatically loads on a designated server when the cluster is first brought up. If the DHCP Start Mode is set to MANUAL, you can manually start DHCP on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the DHCP Failover Mode is set to AUTO, the DHCP server automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the DHCP Failover Mode is set to Manual, you can intervene after a failure occurs and before the DHCP server is started on another node.

If the DHCP Failback Mode is set to DISABLE, the DHCP server will continue running on the node it has failed to. If the DHCP Failback Mode is set to AUTO, the DHCP server automatically moves back to its preferred node when the preferred node is brought back online. Set the DHCP Failback Mode to MANUAL to prevent the DHCP server from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

To view or change DHCP Start, Failover, and Failback modes, do the following:

1 In ConsoleOne, right click the DHCP Resource object and select Properties.

2 Click the Policies tab on the property page.

3 View or change the Start, Failover, or Failback modes for DHCP.

View or Edit DHCP Resource Server Assignments

The DHCP resource template automatically assigns the DHCP resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list. To view or edit DHCP resource node assignments or change the server failover order:

1 In ConsoleOne, right-click the DHCP Resource object and select Properties.

2 Click the Nodes tab on the property page.

3 View the current DHCP resource server assignments and, if necessary, click the right or left arrow button to assign or unassign servers to the DHCP resource.

You can also click the up or down arrow button to change the failover order of the servers assigned to DHCP.

Additional Novell Cluster Services Configuration Information

If you have DHCP clients that are located on a separate LAN segment or IP subnet from the servers in your cluster, some additional configuration is necessary. This additional configuration consists of adding and deleting a secondary IP address on the cluster server where DHCP is running and setting up your router to forward DHCP requests.

To add a secondary IP address:

1 Edit the DHCP cluster resource load script and add the following line:

`ADD SECONDARY IPADDRESS A.B.C.D`

Replace A.B.C.D with a unique IP address. Your router will forward DHCP requests from other LAN segments to this address.

2 Edit the DHCP cluster resource unload script and add the following line:

`DEL SECONDARY IPADDRESS A.B.C.D`

Replace A.B.C.D with the same IP address you added in the load script.

To configure a router to forward DHCP requests, see your router documentation. If you are using a NetWare server as a router, ensure the following line is added to the autoexec.ncf file of the server acting as a router:

`BOOTPFWD SERVER=A.B.C.D`

Replace A.B.C.D with the secondary IP address you specified in the load and unload scripts above.

For additional information on setting up and configuring Novell Cluster Services, See the Novell Cluster Services 1.7 Administration Guide.

6 Configuring iPrint with Novell Cluster Services

This section will help you do the following

- Understand the benefits of high availability printing
- "Creating a Shared Disk Partition for Printing" on page 23
- Create and configure a broker and a print manager
- Configure the print cluster volume

Benefits of High Availability Printing

Running iPrint on a NetWare® cluster provides the following benefits:

- Centralized management
    - By configuring the Print Management Database to reside on a shared disk cluster volume, you no longer need to replicate information to achieve high availability.
    - The cluster console GUI allows you to monitor and control the location of cluster volumes with the Broker and Print Manager services using them from a single management workstation.
    - Because cluster volumes are tightly consistent and highly available, distributed print management tasks are simplified.
- Automatic restart of print services
    - If a server fails, the cluster volume containing the Print Management Database and spool area automatically remounts on a surviving server in the cluster. The Broker and Print Manager NLM™ software will also automatically restart without user intervention.
    - On restart, the Broker and Print Manager will discover the Print Management Database on a cluster volume mounted on their server. Because the cluster volume is the same regardless of which server it is mounted on, no print-related information is lost or out of date.

Creating a Shared Disk Partition for Printing

After you have installed, set up, and configured Novell Cluster Services, you can begin configuring services to run on a cluster. The first task is to create a shared disk partition for printing. While you can use Novell iManager, NSSMU, or Novell Remote Manager to create the shared disk partition, the steps below use the Novell Remote Manager method. For information on using Novell iManager or NSSMU, see "Setting Up Novell Cluster Services" in the *Novell Cluster Services Administration Guide for NetWare*.

1 In Novell Remote Manager, click Partitions under the Clustering heading.

A screen appears displaying a list of devices that are currently accessible to servers in the cluster. For each device, the list displays the partitions, NSS pools, volumes, and free space on that device.

2 Locate the drive that corresponds to your shared storage system and click Create on the free space where you want to create the partition for print.

3 Click Create a New Pool and Volume.

4 Enter the following information, and then click Create.

Pool Size: Enter the size in megabytes you want for your print services. While a 2,000 MB size should be sufficient for an average print system of 300 - 500 printer agents with moderate printing, the size is dependent on the number of printers and the printing load.

Pool Name: Enter a name for the pool. Because only one pool exists in each partition, the pool size you specify will be the size of the partition that gets created. The volume that is created will be allowed to grow to the pool size.

Volume Name: Enter a volume name. The volume name can be the same as the pool name.

Volume Attributes: Select the volume attributes you want enabled. Print does not require any specific attributes.

Cluster Enable Pool: Ensure this is selected.

5 On the Cluster Pool Enabling Information page, accept the default settings, enter the Pool IP Address, and click Create.

The Pool IP Address is the address of the pool for printing that you are creating.

After clicking Create, the Cluster Status is displayed and the newly created print volume is shown.

Creating and Configuring a Broker and Print Manager

Before running iPrint in a cluster environment, you must create a Broker and Print Manager on the cluster-enabled volume. You should dedicate a volume for printing.

To create a Broker and Print Manager, use iManager.

To prepare the print shared disk partition for a print broker:

1 Using Windows* Explorer, browse and select the Print volume in the eDirectory™ tree.

This creates a client connection to the server currently hosting the volume.

2 From a server's local volume, copy the sys:\ndps directory and subdirectories to the root directory of the shared cluster volume created for printing.

Ensure that \ndps is the top -level directory at the root of the Print volume.

To create a Broker using iManager:

1 From iManager, click iPrint Management > Create Broker.

2 Fill in the fields.

For the RMS Volume, select the shared cluster volume created for printing.

NOTE: You cannot create a Broker object until you have copied the resource directory as explained in Step 2 in the above instructions for preparing the print volume for a print broker.

For more information, see "Creating a Broker" in the *OES iPrint Administration Guide for NetWare*.

To create a Print Manager:

1 From iManager, click iPrint Management > Create Print Manager.

2 Fill in the fields.

For the Database Volume select the shared cluster volume created for printing.

For more information, see"Creating a Print Manager" in the *OES iPrint Administration Guide for NetWare*.

IMPORTANT: While you can have more than one Print Manager running in a cluster, only one Print Manager can run on a cluster resource node. For example, if you have a four node cluster with serverA, serverB, serverC, and serverD, you can have two print managers: one running on serverA and one on serverB and then assign them to failover to different servers: serverC and serverD. If you assign the print managers to failover to the same cluster resource node, the Print Managers will collide and not failover.

Configuring the Print Cluster Pool

After you have created a Broker and Print Manager, you must complete the following configuration tasks:

Configuring the Print Pool Load and Unload Scripts

Load and unload scripts are required by Novell Cluster Services to start and stop iPrint and mount the print volume on servers in the cluster. Use the default timeout value on the Scripts page.

To configure the print volume load and unload scripts:

1 In iManager, click Clusters > Cluster Options.

2 Browse to and select the Cluster you want to manage.

3 Select the check box of the Print Cluster object you want to modify, and then click the Properties link under Cluster Objects heading > Scripts.

4 Click the Load Script tab on the property page and add the following lines to the end of the load script.

```
load broker .broker_name.broker_context /allowdup

load ndpsm .manager_name.manager_context /dbvolume=nocheck /
dnsname=print_service_dns
```

Replace *broker_name* and *broker_context* with the Broker object's name and context; replace the *manager_name* and *manager_context* with the Print Manager's name and context; replace *secondary_cluster_address* with the IP address assigned to the cluster-enabled volume.

The /allowdup command line switch instructs the Broker to avoid its check for duplicate Brokers at start-up. Because Novell Cluster Services ensures that only one Broker NLM runs and references the same eDirectory Broker configuration object at a time, this option allows the Broker to start faster without human intervention by avoiding a communication timeout.

The /dbvolume=nocheck command line switch instructs the Print Manager to avoid checking whether its database volume is local or remote. This switch allows the Print Manager to simply open the cluster volume identified in its eDirectory configuration object.

The /dnsname command line switch assigns a DNS name for the print service to the Print Manager. When the print service fails over to another node, the DNS name for the print service also moves to the new node.

NOTE: In prior releases the /ipaddress option was recommended and is still a valid option; however, if the IP address was ever changed, installed printers could no longer print. We recommend you use the / dnsname option.

Example Load Script:

```
nss /poolactivate=IPP
mount IPP VOLID=254
CLUSTER CVSBIND ADD IPP_CLUST_IPP_SERVER 192.0.127.75
NUDP ADD IPP_CLUST_IPP_SERVER 192.0.127.75
add secondary ipaddress 192.0.127.75

load broker .broker1.print.novell /allowd
unload ndpsm .psm1.print.novell /dbvolume=nocheck
/dnsname=printing.novell.com
```

5 Click Unload Script and add the following lines to the *beginning* of the unload script to dismount the print volume from a server. In addition, add the following lines to the

```
UNLOAD NDPSM
```

```
UNLOAD BROKER
```

Example Unload Script:

```
unload ndpsm
unload broker

del secondary ipaddress 192.0.127.75
CLUSTER CVSBIND DEL IPP_CLUST_IPP_SERVER 192.0.127.75
NUDP DEL IPP_CLUST_IPP_SERVER 192.0.127.75
nss /pooldeactivate=IPP /overridetype=question
```

Setting the Print Volume's Host Server Assignments

Novell Cluster Services by default assigns the print volume resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list. To view or edit the print volume resource node assignments or change the server failover order, do the following:

1 In iManager, click Clusters > Cluster Options the print volume resource object and select Properties.

2 Browse to and select the Cluster you want to manage.

3 Select the check box of the Print Cluster object you want to modify, and then click the Properties link under Cluster Objects heading.

4 Enable the Ignore Quorum by selecting the checkbox because you want print to be launched immediately on any server in the Assigned nodes list as soon as any server in the list in brought online.

5 View the current Preferred Nodes assignments, and if necessary, click the right or left arrow button to assign or unassign servers to the print volume resource.

You can click the up or down arrow button to change the failover order of the servers assigned to the print volume.

IMPORTANT: In order for the Broker to load and enable the Resource Management Service properly you must ensure that the IP address of the print volume is added to the \etc\hosts file of each server in the cluster that will host the print volume.

Use a text editor to open the hosts file and if necessary add the IP address you assigned to the print volume resource when you cluster-enabled it. For example, if the print cluster volume is named cluster_printvol_server and the IP address for that cluster volume is 192.0.34.70, you would add the following line:

```
192.0.34.70 cluster_printvol_server
```

Setting Start, Failover, and Failback modes for Printing

Novell Cluster Services by default sets the print cluster volume Start Mode and Failover Mode to Auto, and the Failback Mode to Disable. If the print volume Start Mode is set to Auto, iPrint automatically starts on a server when the cluster is first brought up. If the print volume Start Mode is set to Manual, you must manually start print on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the print volume Failover Mode is set to Auto, iPrint automatically starts on the next server in the Assigned Nodes list in the event of a hardware or software failure. If the print volume Failover Mode is set to Manual, you can intervene after a failure occurs and before iPrint is moved to another node.

If the print volume Failback Mode is set to Disable, iPrint will not fail back to its most preferred node when the most preferred node rejoins the cluster. If the print volume Failback Mode is set to Auto, iPrint will automatically fail back to its most preferred node when the most preferred node rejoins the cluster. Generally, you want a failback to a preferred node to occur after business hours. You can set the print volume Failback Mode to Manual to prevent iPrint from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

IMPORTANT: We recommend you use the default settings for Start, Failover and Failback modes.

To view or change the print volume Start, Failover, and Failback modes:

1 In iManager, click Clusters > Cluster Options the print volume resource object and select Properties.

2 Browse to and select the Cluster you want to manage.

3 Select the check box of the Print Cluster object you want to modify, and then click the Properties link under Cluster Objects heading.

The Failover and Failback modes are displayed.

Using iPrint in a Clustered Environment

If you are using the default iPrint configurations or are embedding URLs in your iPrint Web pages, you need to be aware of the following:

◆ Links to the default iPrint Printer List Web page or other custom pages in \ippdocs should use the DNS name of the print service. This address moves with the volume to other servers in the cluster during a failover. If you use the cluster's master IP address or the server's primary DNS name or IP Address, after a failover the Printer List might display blank or give an error.

- If you are using custom map files that are saved in sys:\apache2\htdocs\ippdocs\images\maps, you must copy the maps to each file server in the cluster that is hosting iPrint. This will make the maps available if a failover occurs.

- When adding drivers to Resource Management Service, make sure that the files in the ndps\resdir directory are not flagged as read-only. This is not an issue when copying the \resdir to sys:; however, if you extract the files from .zip files on a CD, the directories might be flagged read-only.

7 Configuring Novell exteNd Application Server and MySQL with Novell Cluster Services

Novell® exteNd™ Application Server requires MySQL*, so this document includes instructions for configuring both applications with Novell Cluster Services™.

Benefits of Configuring the exteNd Application Server and MySQL for High Availability

Novell exteNd Application Server provides the foundation for building and deploying cross-platform, high-performance, standards-based applications. MySQL is an open source, structured query language (SQL) database. If the NetWare® 6.5 server running the exteNd Application Server or MySQL goes down, the services and databases provided by exteNd Application Server and MySQL will not be accessible unless they are in a clustered environment. Configuring exteNd Application Server and MySQL with Novell Cluster Services ensures that they are scalable and highly available.

Novell Cluster Services Configuration and Setup

NetWare 6.5 and Novell Cluster Services 1.7 must be installed before configuring exteNd Application Server and MySQL to run with Novell Cluster Services. NetWare 6.5 includes Novell Cluster Services 1.7 and licenses for two cluster nodes.

After Novell Cluster Services is installed, you must create and cluster-enable a pool/volume for the MySQL and exteNd Application Server databases. Creating a cluster-enabled pool/volume resource in Novell Cluster Services also requires you to configure volume resource load and unload scripts, set volume resource Failover and Failback modes, and assign the volume resource to specific servers in your cluster.

Creating and Cluster-Enabling a Pool/Volume Resource

You must already have an NSS partition created on a device that is sharable for clustering (shared disk system or SAN) prior to creating and cluster-enabling a pool/volume. For more information on creating cluster partitions on shared storage, see "Create Shared Disk Partitions" in the *Novell Cluster Services 1.7 Administration Guide*.

Only one NSS pool can be created on a partition. Storage pools can be cluster-enabled at the same time they are created or they can be cluster-enabled at a later time after they are created.

To create and cluster-enable a pool at the same time:

1 From the NSSMU main menu, select Pools, press Insert, then type a name for the new pool you want to create.

2 Select the partition on your shared storage where you want the pool created.

3 Choose to activate and cluster-enable the pool when it is created.

The Activate on Creation feature 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 must manually activate it later before it can be used.

The Cluster Enable on Creation feature is also enabled by default. To cluster-enable the pool at the same time it is created, accept the default entry (Yes).

4 Specify the Virtual Server Name, IP Address, Advertising Protocols and, if necessary, the CIFS Server Name.

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

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

You can select one or all of the advertising protocols. NCP™ is the protocol used by Novell clients, CIFS is the protocol used by Microsoft* clients, and AFP is the protocol used by Macintosh* clients. Selecting any of the protocols will cause lines to be added to the pool resource load and unload scripts to activate the selected protocols on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to all your clients.

If you select CIFS as one of the protocols, a CIFS Server Name is also required. This is the server name CIFS clients see when they browse the network. A default server name is listed, but you can change the server name by editing the text in the field.

5 Select Create to create and cluster-enable the pool.

If you have already created a pool and want to cluster-enable it, see "Cluster Enable Pools and Volumes" in the *Novell Cluster Services 1.7 Administration Guide*.

Editing Cluster Volume Load and Unload Scripts

When you cluster-enable a pool or volume, commands are automatically added to the pool load and unload scripts to start and stop the pool resource on servers in your cluster. In addition to the commands to start and stop the pool resource, you must add commands to start and stop the Application Server and MySQL.

To edit the pool load script:

1 In ConsoleOne®, right-click the Pool Resource object, then click Properties.

If you just created or cluster-enabled the cluster pool and checked the Define Additional Properties check box, the cluster pool object should already be displayed.

2 Click the Scripts tab on the property page, then select Cluster Resource Load Script.

The load script will appear similar to the following example:

```
nss /poolactivate=POOL1
mount VOL1 VOLID=254
CLUSTER CVSBIND ADD CLUSTER1_POOL1_SERVER 111.152.101.190
NUDP ADD CLUSTER1_POOL1_SERVER 111.152.101.190
add secondary ipaddress 111.152.101.190
```

3 Add the following lines to the end of the load script .

```
mysqld_safe bind_ipaddress=a.b.c.d --autoclose
load silverserver.nlm
```

Replace *a.b.c.d* with the IP address you assigned to the cluster-enabled volume.

For example, if the IP address for the cluster volume resource is 111.152.101.10, the line in the load script would appear as follows:

```
mysqld_safe bind_ipaddress=111.152.101.10 --autoclose
```

If you are not configuring exteNd Application Server, the load silverserver.nlm line should not be added.

4 Click the Unload Script tab on the property page and add the following lines to the beginning of the unload script.

```
unload silverserver.nlm
unload mysqld
```

If you are not configuring exteNd Application Server, don't add the unload silverserver.nlm line.

Setting Cluster Volume Start, Failover, and Failback Modes

When you cluster enable a pool/volume, Novell Cluster Services sets the pool resource Start Mode and Failover Mode to AUTO, and the Failback Mode to DISABLE.

If the pool resource Start Mode is set to AUTO, the cluster pool automatically starts on a designated server when the cluster is first brought up. If the cluster pool/volume Start Mode is set to MANUAL, you can manually start the cluster pool/volume on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the cluster pool Failover Mode is set to AUTO, the cluster pool/volume automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the cluster pool/volume Failover Mode is set to Manual, you can intervene after a failure occurs and before the pool/volume is started on another node.

If the cluster pool/volume Failback Mode is set to DISABLE, the cluster pool/volume will continue running on the node it has failed to. If the cluster pool/volume Failback Mode is set to AUTO, the cluster pool/voume automatically moves back to its preferred node when the preferred node is brought back online. Set the cluster pool/volume Failback Mode to MANUAL to prevent the pool/volume from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

To view or change cluster pool/volume Start, Failover, and Failback modes:

1 In ConsoleOne, right-click the cluster pool resource object, then select Properties.

2 Click the Policies tab on the property page.

3 View or change the Start, Failover, or Failback modes for the cluster-enabled pool/volume.

Viewing or Editing Cluster Volume Resource Server Assignments

When you cluster-enable a pool/volume, the volume resource is automatically assigned to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list.

To view or edit the cluster pool/volume resource node assignments or change the server failover order:

1 In ConsoleOne, right-click the cluster pool resource object, then select Properties.

2 Click the Nodes tab on the property page.

3 View the current cluster pool resource server assignments and, if necessary, click the right or left arrow button to assign or unassign servers to the cluster pool/volume resource.

You can also click the up or down arrow button to change the failover order of the servers assigned to cluster pool/volume.

MySQL and exteNd Application Server Configuration and Setup

Before configuring exteNd Application Server and MySQL with Novell Cluster Services, both exteNd Application Server and MySQLmust be installed and configured properly on all servers in the cluster where you intend to have them run. For instructions on installing and configuring exteNd Application Server with NetWare 6.5, see the exteNd Application Server documentation (http://www.novell.com/documentation/lg/extendas50/index.html). For instructions on installing and configuring MySQL with NetWare 6.5, see the *MySQL for NetWare Administration Guide for NetWare 6.5*.

exteNd Application Server and MySQL software is included with NetWare 6.5. You should select the Customized NetWare Server installation option and check the MySQL and exteNd Application Server component check boxes to install both products during the NetWare 6.5 installation. You can also use the exteNd J2EE* Web Application Server Pre-Configured Server option to install both MySQL and the exteNd Application Server during the NetWare 6.5 installation.You cannot install MySQL or exteNd Application Server with the Basic NetWare File Server installation option.

In order for exteNd Application Server and MySQL to function properly in a cluster, the database data must reside on a cluster-enabled volume (virtual server). This means that a shared storage system (SAN) is required.

Edit Autoexec.ncf Files

When you install MySQL and exteNd Application Server, commands are added to the autoxec.ncf file on the server to automatically start both programs when the NetWare server starts. Because Novell Cluster Services starts and stops the programs, these commands must be removed or commented out from each server in the cluster where they are installed.

Edit the autoexec.ncf file on each server in the cluster where MySQL and exteNd Application Server are installed and comment out or remove the following lines:

```
mysqld_safe  --autoclose
```

```
load silverserver.nlm
```

If you are only configuring MySQL, and do not have exteNd Application Server installed, the load silverserver.nlm line will not be present.

Configure MySQL on the Cluster-Enabled Volume

In order for MySQL to take advantage of the benefits provided by Novell Cluster Services, you must make some configuration changes to MySQL. These configuration changes include editing

a MySQL configuration file to change the MySQL data directory to point to the cluster-enabled volume and installing the MySQL database on the shared volume.

These instructions assume no database has been created. If a MySQL database currently exists, the database must be moved to the cluster-enabled volume, rather than creating it as described in .

1 On all cluster nodes where MySQL is installed, use a text editor to open the sys:\etc\my.cnf file and change the MySQL data directory to point to the directory you will create on the cluster-enabled volume.

For example, the current line should read:

datadir=sys:/mysql/data

If you named the cluster-enabled volume vol1, you would change the line to read:

datadir=vol1:/mysql/data

2 On the server in the cluster that has the cluster-enabled volume mounted, create a mysql\data directory at the root of the cluster volume.

3 On the server in the cluster that has the cluster-enabled volume mounted, enter **mysqld_install_db** at the server console.

This installs the MySQL database on the shared volume.

4 Check the mysql\data directory on the cluster-enabled volume to ensure that the MySQL database was created.

Configure exteNd Application Server on the Cluster-Enabled Volume

Some exteNd Application Server configuration changes are also required in order for it to function properly in a cluster. These configuration changes include changing the jdbc URL to cause the MySQL SilverMaster to point to the cluster-enabled volume and creating the SilverMaster database on the cluster-enabled volume.

These instructions assume no database has been created. If SilverMaster databases currently exist, the databases must be moved to the cluster-enabled volume, rather than creating them as described in .

1 On all cluster nodes where exteNd Application Server is installed, use a text editor to open the sys:/exteNd/AppServer/Resources/httpd.props file and change the jdbc URL to point to the directory you will create on the cluster-enabled volume.

For example, if the DNS name for the server is server1.acme.com, the current line should read:

http-server.Jdbc.URL=jdbc:mysql://server1.acme.com/SilverMaster50

If the cluster-enabled volume (virtual server) name is cluster1_pool1_server, you would change the line to read:

http-server.Jdbc.URL=jdbc:mysql://cluster1_pool1_server/SilverMaster50

2 On the server in the cluster that has the cluster-enabled volume mounted, enter **silvermasterinit** at the server console.

This creates the SilverMaster database on the cluster-enabled volume.

If you have secured the SilverMaster database, you must supply a username and password when running SilverMasterInit. To add a username and password you would type **silvermasterinit -u *username* -p *password***.

3 Check the mysql\data directory on the cluster-enabled volume to ensure that the SilverMaster50 database was created.

Starting MySQL and exteNd Application Server in the Cluster

To start MySQL and exteNd Application Server in the cluster you must unload MySQL and App Server from all servers in the cluster where they are installed, and then bring the cluster-enabled volume resource offline and online again.

1 At the server console of each server in the cluster where MySQL and exteNd Application Server are installed and running, type **unload silverserver** and then **unload mysqld.**

If you are configuring only MySQL, it is not necessary to unload silverserver.

2 At the server console of the machine that has the cluster-enabled volume mounted, enter **cluster offline *pool1_server.***

Replace pool1_server with the name of the cluster-enable pool.

3 At the same server console, enter **cluster online *pool1_server***.

Again, replace pool1_server with the name of the cluster-enable pool.

When you bring the cluster enabled volume resource back online MySQL and exteNd Application Server will start on a specified node in the cluster.

Additional Novell Cluster Services Configuration Information

For additional information on setting up and configuring Novell Cluster Services, see the *Novell Cluster Services 1.7 Administration Guide*.

8 Configuring DNS with Novell Cluster Services

A NetWare® 6.5 primary DNS server automatically resolves IP addresses with domain names. If for some reason the NetWare 6.5 DNS server is not accessible, clients will lose their ability to access certain network resources and applications. Configuring DNS with Novell® Cluster Services™ can help avoid this loss of access.

Benefits of Configuring DNS for High Availability

Configuring DNS with Novell Cluster Services ensures the IP address/domain name resolution required by users is always highly available. This is possible because DNS configuration information is stored in eDirectory™ and the DNS server is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services.

Requirements

The following requirements must be met prior to configuring DNS with Novell Cluster Services.

❑ NetWare 6.5 installed and running on all servers that will be configured with DNS and Novell Cluster Services.

❑ A cluster set up and configured with Novell Cluster Services installed and running on each node in the cluster where DNS will be installed.

❑ All cluster nodes that will have DNS installed and configured contain a replica of the eDirectory partition where the DNS objects are located.

Configuring DNS

To properly configure DNS with Novell Cluster Services, the DNS server must be configured with a cluster-enabled pool (virtual server) on a shared disk system (SAN). This makes it possible for the DNS server configuration to move to different servers in the cluster with the DNS server.

DNS can be installed by choosing the DNS/DHCP component that is part of the NW 6.5 custom installation option or by choosing the DNS/DHCP pattern installation option. DNS must be installed on all cluster nodes where it will run. See the NetWare 6.5 Overview and Installation Guide for more information on NetWare 6.5 custom and pattern installation options.

DNS can also be installed after the NetWare 6.5 server installation. See "Installing DNS" on page 37.

Creating a Cluster-Enabled Pool (Virtual Server)

Storage pools can be created and cluster enabled at the same time the partitions that contain them are created.

To create an NSS partition and create and cluster enable a pool on that partition at the same time, complete the following steps:

1 Start Internet Explorer 5 or later 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 a NetWare 6.5 server in the cluster or the IP address for Apache-based services.

2 Enter your username and password.

3 In the left column, locate Storage, then click the Pools link.

4 Enter a cluster server name or browse and select one, then click the New link.

5 Specify the new pool name, then click Next.

6 Select the check box next to the device where you want to create the pool, then specify the size of the pool.

The pool should be at least 100 MB.

Because only one pool exists in each partition, the pool size you specify will be the size of the partition that is created.

7 Choose *not* to activate and cluster-enable the pool when it is created, then click Next.

The Activate On Creation check box is used to determine if the pool you are creating is to be activated as soon as it is created. The Activate On Creation check box is selected by default. Ensure that you deselect the check box. You will manually activate the pool later in the setup process.

8 Specify the Virtual Server Name, Pool IP Address, Advertising Protocols and, if necessary, the CIFS Server Name, then click Finish.

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

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-

enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

You can select one or all of the advertising protocols. NCP™ is the protocol used by Novell clients, CIFS is the protocol used by Microsoft* clients, and AFP is the protocol used by Macintosh* clients. Selecting any of the protocols causes lines to be added to the pool resource load and unload scripts to activate the selected protocols on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to all your clients.

If you select the CIFS check box, the Server Name field becomes active. The CIFS server name is the server name CIFS clients see when they browse the network. A default server name is listed, but you can change the server name by editing the text in the field.

You can also create and cluster enable pools after a partition has been created. For more information, see "Setting Up Novell Cluster Services" in the Novell Cluster Services 1.7 Administration Guide.

Creating a Cluster Volume on the Cluster-Enabled Pool

1 In the left column of the main iManager page, locate Storage, then click the Volumes link.

2 Enter a cluster server name or browse and select one, then click the New link.

3 Specify the new volume name, then click Next.

4 Select the check box next to the cluster pool where you want to create the volume and either specify the size of the volume (Volume Quota) or select the check box to allow the volume to grow to the size of the pool, then click Next.

The volume quota is the maximum possible size of the volume. If you have more than one volume per pool, you should specify a quota for each volume rather than allowing multiple volumes to grow to the pool size.

5 Review and change volume attributes as necessary.

The Flush Files Immediately feature helps ensure the integrity of volume data. Enabling the Flush Files Immediately feature improves file system reliability but hampers performance. You should consider this option only if necessary.

6 Choose to *not* have the volume activated and mounted when it is created, then click Finish.

Installing DNS

If you did not install DNS on all cluster nodes where it will run during the NetWare 6.5 server installation, you must install it before configuring DNS on cluster nodes. Complete the following steps to install DNS.

1 On a cluster server where DNS will run, enter **DNIPINST** at the server console.

2 Log in to the eDirectory tree as user Admin or equivalent.

A message should appear indicating that the NDS® schema extensions were successfully added. If this message is not displayed, there could be eDirectory issues that need to be resolved before proceeding.

3 If this is the first server in the cluster where DNIPINST is run, continue with Step 4 on page 37. If you have already run DNIPINST on another server in the cluster, continue with Step 5 on page 38.

4 Enter the eDirectory contexts for the Locator object, Group object, and RootSrvrinfo object.

These can all be in the same context. The objects will be created in the contexts you specify and are the location where the DNS/DHCP information is stored in eDirectory.

5 Enter the eDirectory context for the Locator object.

This is the same context you specified in Step 3 above.

6 Repeat the above steps on each cluster server where DNS will run.

Configuring DNS

Configuring DNS consists of

- Setting up the DNS/DHCP scope and Administration scope
- Setting up the DNS Server object
- Setting up the DNS zone
- Creating an IN-ADDR ARPA zone
- Creating a DNS Resource Record

Setting Up the DNS/DHCP Scope and Administration Scope

Setting up a DNS/DHCP scope and an Administration scope is not required, but doing so can optimize performance for DNS administration. You can still manage DNS without setting up a DNS/DHCP scope and an Administration scope, but warning messages will be displayed in iManager. Setting the DNS/DHCP scope prevents the warnings from appearing; management can be optimized further by also setting the administration scope.

Setting up the scopes for management optimization is only valid for the current session.

1 Start Internet Explorer 5 or later 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 a NetWare 6.5 server in the cluster or the IP address for Apache-based services.

2 Enter your username and password.

3 In the left column, click DNS, then click DNS/DHCP Scope Settings.

4 Enter the context for the DNS/DHCP Locator object.

This is the eDirectory context of the container object that will provide the administrative scope for the current session.

5 Enter the Administrative Scope for the current session.

This is the context for the admin user that is currently authenticated to the eDirectory tree.

If you specify only the eDirectory context of the DNS/DHCP Locator object and not the administrative scope of the current session, you can proceed with administrative tasks without inviting a warning message. Performance however, is further optimized if you also define the administrative scope.

Setting Up the DNS Server Object

1 In the left column of the main iManager page, click DNS, then click DNS Server Management.

2 Select the Create Server operation, then click OK.

3 In the NCP Server Name field, enter the name of the virtual server (cluster-enabled pool) you previously created.

See "Creating a Cluster-Enabled Pool (Virtual Server)" on page 36.

4 Specify a unique host name for the DNS server and a domain name for the server in the applicable fields, then click Create.

5 Click OK and view the message to ensure that the DNS server was successfully created.

Setting Up the DNS Zone

A DNS Zone object is an eDirectory container object that holds all the data for a single DNS zone. Multiple DNS domains can be represented within eDirectory by using separate, independent DNS Zone objects.

The DNS Zone object contains data that correlates to a DNS Start of Authority (SOA) resource record (RR), a member list of all eDirectory DNS servers that serve the zone, and Dynamic DNS (DDNS) server information.

1 In the left column of the main iManager page, click DNS, then click Zone Management.

2 Select Create Zone, then click OK.

3 Ensure that Create New Zone is selected as the zone type, then specify the eDirectory context for the zone and the domain name for the Zone object.

The eDirectory conext is the location where the Zone object will be stored in eDirectory.

4 Select Primary as the zone type.

A primary zone object is associated with a primary name server and contains the most up-to-date information about the zone and all the hosts within it.

5 Select the DNS virtual server object as the Authoritative Zone Server.

6 Click Create and ensure the DNS zone was successfully created.

Creating an IN-ADDR ARPA zone

ARPA zones are not required for DNS operation, but are required for reverse DNS lookup. For this reason, we recommend creating ARPA zones for all DNS zones.

An IN-ADDR.ARPA zone enables a client to request a hostname by providing an IP address.

1 In the left column of the main iManager page, click DNS, then click Zone Management.

2 Select Create Zone, then click OK.

3 Select IN-ADDR ARPA as the zone type, then Select Primary as the zone type.

4 Specify the eDirectory context for the Zone object.

The eDirectory context is the location where the Zone object will be stored in eDirectory.

5 Enter the IP address for the network.

This is the network IP address, not the IP address of the host server.

6 Select the DNS virtual server object as the Authoritative Zone Server.

7 Click Create and view the message to ensure the ARPA zone was successfully created.

The DNS Resource Record Set (RR set) object is contained within a DNS Zone object. An RR Set object represents an individual domain name within a DNS zone. Each RR Set object has one or more resource records beneath it containing additional information about the domain.

The most common resource records are Address (A) records, which map a domain name to an IP address, and Pointer (PTR) records, which map an IP address to a domain name within an IN-ADDR.ARPA zone.

1 In the left column of the main iManager page, click DNS, then click Resource Record Management.

2 Select Create Resource Record, then click OK.

3 From the drop-down menu, select the domain name where the resource record is to be created, then click Create.

4 In the Owner name field, enter the DNS name for the resource within the specified domain.

5 Select the type of record needed.

In most instances, this will be an A record.

6 Enter the IP address for the resource, then click Create and view the message to ensure the resource was successfully created.

An ARPA record will automatically be added for this DNS record if you created an ARPA zone.

Configuring the DNS Cluster Resource

The DNS cluster resource is the cluster-enabled pool (virtual server) you created earlier. Before bringing the DNS cluster resource online, you must edit the resource load and unload scripts, set failover and failback modes for the resource, and assign the resource to the desired cluster servers.

Edit Load and Unload Scripts

Before DNS will function correctly as a cluster resource, you must edit the cluster-enabled pool load and unload scripts to start and stop the DNS service.

1 In the left column of the main iManager page, locate Cluster Administration, then click the Configure link.

2 Enter the cluster name or browse and select it, select the check box next to the resource whose load script you want to edit, then click the Properties link.

3 Click the Scripts tab, then click the Load Script link.

4 Add the **named -v volume_name** command to the end of the load script, then click Apply.

For example, if the volume you created in "Creating a Cluster Volume on the Cluster-Enabled Pool" on page 37 is named DNS, the load script would appear similar to the following example.

```
nss /poolactivate=DNSPOOL

mount DNS VOLID=254

CLUSTER CVSBIND ADD DNS_vServer 10.1.1.52
```

```
NUDP ADD DNS_vServer 10.1.1.52

add secondary ipaddress 10.1.1.52

named -v DNS
```

5 Click the Scripts tab, then click the Unload Script link.

6 Add the **unload named** command to the beginning of the unload script, then click Apply.

The unload script should now appear similar to the following example.

```
unload named

del secondary ipaddress 10.1.1.52

CLUSTER CVSBIND DEL DNS_vServer 10.1.1.52

NUDP DEL DNS_vServer 10.1.1.52

nss /pooldeactivate=DNSPOOL /overridetype=question
```

The **named** command (NAMED.NLM) specified in the load script provides the DNS server functionality. NAMED.NLM reads configuration information from the NCP server object. Adding the -v switch and volume name of the cluster-enabled volume (virtual server) to the command enables NAMED.NLM to determine which NCP server to read it's configuration from.

Setting DNS Start, Failover, and Failback Modes

When you created a cluster-enabled pool (virtual server) for DNS, Novell Cluster Services automatically sets the pool resource Start Mode and Failover Mode to AUTO, and the Failback Mode to DISABLE.

If the DNS pool resource Start Mode is set to AUTO, the pool resource automatically loads on a designated server when the cluster is first brought up. If the pool resource Start Mode is set to MANUAL, you can manually start the pool resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the pool resource Failover Mode is set to AUTO, the pool resource automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the pool resource Failover Mode is set to Manual, you can intervene after a failure occurs and before the pool resource is started on another node.

If the pool resource Failback Mode is set to DISABLE, the pool resource will continue running on the node it has failed to. If the pool resource Failback Mode is set to AUTO, the pool resource automatically moves back to its preferred node when the preferred node is brought back online. Set the pool resource Failback Mode to MANUAL to prevent it from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

To view or change the DNS pool resource Start, Failover, and Failback modes:

1 In the left column of the main iManager page, locate Cluster Administration, then click the Configure link.

2 Enter the cluster name or browse and select it, check the box next to the DNS pool resource, then click the Properties link.

3 Click the General tab.

4 View or change the Start, Failover, and Failback modes for the DNS pool resource.

View or Edit DNS Resource Server Assignments

When you created a cluster-enabled pool (virtual server) for DNS, Novell Cluster Services automatically assigned the DNS pool resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list. To view or edit DNS resource node assignments or change the server failover order:

1 In the left column of the main iManager page, locate Cluster Administration, then click the Configure link.

2 Enter the cluster name or browse and select it, select the check box next to the DNS pool resource, then click the Properties link.

3 Click the General tab.

4 View the current DNS pool resource server assignments and, if necessary, click the right or left arrow button to assign or unassign servers to the DNS pool resource.

Additional Information

For more information on configuring Novell Cluster Services, see the *Novell Cluster Services 1.7 Administration Guide*. For more information on configuring DNS, see the *Novell DNS/DHCP Services for NetWare Administration Guide for NetWare 6.5*.

# 9 Configuring NetStorage with Novell Cluster Services

Benefits of Configuring NetStorage for High Availability

NetStorage provides secure Internet-based access to files and folders on a NetWare® server on your network using either a browser or Microsoft* Web Folders (Microsoft's implementation of WebDAV). If the NetWare® 6.5 server running NetStorage goes down, Internet-based access to files and folders will not be possible. Configuring NetStorage with Novell Cluster Services ensures Internet-based access to files and folders can continue even if a server running NetStorage goes down.

NetStorage Installation and Configuration

Novell NetStorage software is included with NetWare 6.5 and can be installed during the NetWare 6.5 installation. You must select the Customized NetWare Server installation option and check the Novell NetStorage component check box to install NetStorage during the NetWare 6.5 installation. Because NetStorage requires iManager and Apache, you must also select the Apache Web server and iManager check boxes. See NetWare 6.5 NetStorage Administration Guide for more information on installing and configuring NetStorage.

Because most networks require NetStorage to be installed on only one server, it is generally only necessary to install NetStorage on two servers in the cluster. Users gain access to files and folders by connecting to one NetStorage server. The other NetStorage server acts as a backup in case the first NetStorage server fails. The following figure shows how a typical NetStorage cluster configuration might look.
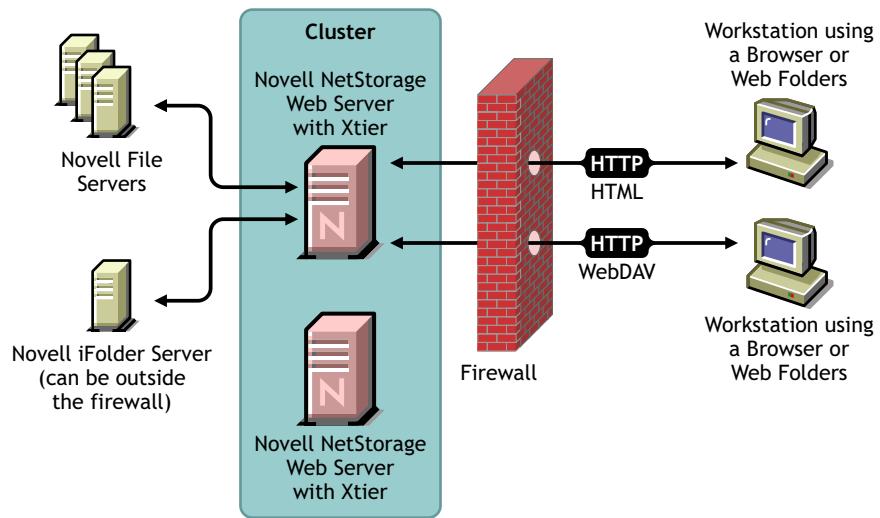
When you install NetStorage with NetWare 6.5, you are asked to specify the DNS name or IP address of the Primary eDirectory Server. This is the IP address or registered DNS name of a server in your eDirectory tree that has the master replica or a read/write replica of eDirectory. If you have iFolder installed on your network, you are also asked to specify the DNS Name or IP Address of for iFolder. For NetStorage to function properly in a cluster, you must specify same DNS name or IP address of the primary eDirectory Server on all servers in the cluster that have NetStorage installed. The same is true for iFolder. If you have iFolder installed, you must also specify the same iFolder DNS Name or IP Address on all servers in the cluster.

Any NetStorage-specific configuration changes made to one NetStorage server must be made on all other NetStorage servers in the cluster. In order for NetStorage to function properly with Novell Cluster Services, all NetStorage servers in the cluster must be identically configured.

Novell Cluster Services Configuration and Setup

Novell Cluster Services must be installed and running on the servers that have NetStorage installed before performing the remainder of the NetStorage and Novell Cluster Services configuration.

Because NetStorage is installed and configured identically on the cluster servers where it will run and because that configuration information is stored on each server, no shared storage is required to run NetStorage with Novell Cluster Services.

Creating a NetStorage Cluster Resource

A cluster resource containing a secondary IP address is necessary to ensure users continue to have access to NetStorage should a server running NetStorage go down.

Novell Cluster Services includes a Generic IP Service resource template, which simplifies the process for creating a NetStorage cluster resource.

1 In ConsoleOne®, browse to and select the Cluster object.

2 Click File > New > Cluster > Cluster Resource.

3 Enter a name for the NetStorage cluster resource.

4 Enter the Generic IP Service template name in the Inherit from Template field, or browse and select it from the list.

5 Check the Define Additional Properties check box, click Create, and then continue with .

The Generic IP Service template helps configure the NetStorage cluster resource by automatically creating load and unload scripts, setting failover and failback modes, and assigning the resource to all nodes in the cluster.

You should only assign the resource to those nodes in the cluster that have NetStorage installed and running.

Editing NetStorage Load and Unload Scripts

The Generic IP Service resource template automatically creates a load and unload script to start and stop the NetStorage resource on servers in your cluster. The load script contains commands that you must customize for your specific NetStorage configuration.

To edit the NetStorage load script:

1 In ConsoleOne, right-click the NetStorage Resource object and then click Properties.

If you just created the NetStorage resource object and checked the Define Additional Properties check box, the NetStorage Resource object should already be displayed.

2 Click the Scripts tab on the property page and select Cluster Resource Load Script.

The load script includes instructions and commands, all of which should currently be commented out with the # symbol.

3 Uncomment the following line by removing the # symbol.

```
#add secondary ipaddress A.B.C.D
```

4 Replace A.B.C.D with the IP address you want to assign to the NetStorage resource.

For example, if the IP address for the NetStorage cluster resource is 111.152.153.10, the line in the load script would now appear as follows:

```
add secondary ipaddress 111.152.153.10
```

Clients will now use the NetStorage cluster resource IP address as part of the URL to access NetStorage instead of the NetStorage server's IP address or DNS name.

5 Click Unload Script on the Scripts tab on the property page to edit the unload script.

All instructions and commands in the unload script should currently be commented out with the # symbol, like they are in the load script.

6 Uncomment the following line by removing the # symbol.

```
#del secondary ipaddress A.B.C.D
```

7 Replace A.B.C.D with the IP address you want to assigned to the NetStorage resource in Step 4 on page 45.

For example, if the IP address for the NetStorage cluster resource is 111.152.153.10, the line in the unload script would now appear as follows:

```
del secondary ipaddress 111.152.153.10
```

Setting the NetStorage Resource Start, Failover, and Failback Modes

The Generic IP Service resource template sets the NetStorage resource Start Mode and Failover Mode to AUTO, and the Failback Mode to DISABLE.

If the NetStorage resource Start Mode is set to AUTO, the resource automatically starts on a designated server when the cluster is first brought up (secondary IP address specified in the resource is automatically added to the server). If the NetStorage resource Start Mode is set to MANUAL, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the NetStorage resource Failover Mode is set to AUTO, the NetStorage resource automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the NetStorage resource Failover Mode is set to Manual, you can intervene after a failure occurs and before the resource is started on another node.

If the NetStorage resource Failback Mode is set to DISABLE, the resource will continue running on the node it has failed to. If the NetStorage resource Failback Mode is set to AUTO, the resource automatically moves back to its preferred node when the preferred node is brought back online. Set the NetStorage resource Failback Mode to MANUAL to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

To view or change the NetStorage resource Start, Failover, and Failback modes, do the following:

1 In ConsoleOne, right click the NetStorage Resource object and select Properties.

2 Click the Policies tab on the property page.

3 View or change the Start, Failover, or Failback modes for the NetStorage resource.

View or Edit NetStorage Resource Server Assignments

The Generic IP Service resource template automatically assigns the NetStorage resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list. You should only assign the NetStorage resource to those servers in the cluster that have NetStorage installed. To view or edit the NetStorage resource node assignments or change the server failover order:

1 In ConsoleOne, right-click the NetStorage Resource object and select Properties.

2 Click the Nodes tab on the property page.

3 View the current NetStorage resource server assignments and, if necessary, click the right or left arrow button to assign or unassign servers to the resource.

You can also click the up or down arrow button to change the failover order of the servers assigned to NetStorage.

Accessing NetStorage after Cluster Configuration

To access NetStorage after configuring it with Novell Cluster Services you should inform your clients to use the NetStorage cluster resource IP address as part of the URL to access NetStorage instead of the NetStorage server's IP address or DNS name. For more information on accessing NetStorage, see "Using NetStorage "in the NetWare 6.5 NetStorage Administration Guide.