

Windows File and Print Clustering

Contributed by David Noel-Davies

This white paper presents scenarios for creating highly available file and printer shares with the Cluster service feature in the Microsoft® Windows® 2000 operating system. Although this document by no means presents a comprehensive list of every possible scenario, it tries to provide enough examples so that administrators can implement Cluster service in the way that best serves their users' needs.

Introduction Print service availability is often considered critical in computer network deployments. Consequently, enabling administrators to create highly available print shares is one of the key goals of the Cluster service feature in the Microsoft® Windows® 2000 Advanced Server operating system. Because clustering is new to Windows, administrators must move away from machine-centric view of Windows. This document covers the basic information you need to know to implement Windows 2000 Cluster service correctly and describes the shared-nothing architecture of Cluster service, group structure, and virtual servers. This document also describes in detail how to create print spooler resources, and the Appendix A provides a detailed illustration of the cluster setup that is referred to throughout this document.

Basic Terms and Concepts This section covers the basic terms and concepts that are used throughout this document.

Naming Conventions If you're new to clustering in Windows, you may become confused by the different naming conventions that are used. During beta testing of the Cluster service, the product was code-named Wolfpack. The product was released with the name Microsoft Cluster service and is sometimes referred to as MSCS. In addition, some Knowledge Base articles refer to the technology as Cluster Server or Server Clustering. By any name, the technology is the same.

Note: Do not confuse This type of clustering should not be confused with Network Load Balancing or Component Load Balancing.

Shared-Nothing Architecture Cluster service uses a shared-nothing model of clustering. This model does not allow cluster members to access resources on other members of the cluster, as a shared-everything model would. It also avoids the overhead and inherent scalability limits that would be imposed by a shared-everything model. In terms of file and printer shares, this means that a file share is owned by only one node. To split file share load among multiple servers, the file shares themselves must be split into different virtual servers.

Resources A resource provides a service to a client, such as a file share, an Internet Protocol (IP) address, or a network name. Resources are the smallest management unit of Cluster service. Resources can depend on other resources and are organized into groups.

Dependencies A dependency is a relationship between resources similar to service dependencies. Dependencies additionally define the order that resources are brought online and offline and are typically represented as dependency trees. Figure 1 illustrates a standard dependency tree for a file share resource. This illustration shows that should an administrator attempt to take the disk resource offline, the file share resource also goes offline, followed by the disk resource. The network name resource does not go to an offline state, because there is no direct dependency on the disk for the network name. Conversely, if we then try to bring the file share resource online when the initial state of all the resources is offline, the disk and IP address are brought online, followed by the network name, and then the file share resource. Note that there is no explicit dependency between the file share resource and the IP address. This is because dependencies are said to be transitive—a definition that is not technically correct but provides a good description of the empirical behavior of a dependency.

Figure 1: Typical file share dependency tree Given the above information, it would appear that a dependency is strictly a one-way relationship. This is not the case, however—a point which is best illustrated here by the print spooler resource. When adding a printer to a virtual server, it is absolutely necessary to run the Cluster Application Wizard from the network name for which the print spooler has a dependency. Applications, network name resources, and generic services also exhibit some similar behavior.

Note: For more information on dependencies, see the following Microsoft Knowledge Base articles:

171791 Creating Dependencies in Microsoft Cluster Server

198893 Effects of Checking "Use Network Name for Computer Name" in MSCS

195462 WINS Registration and IP Address Behavior for MSCS 1.0

Groups A group is a logical collection of resources that must all run on the same node to function properly. For example, if the Information Store and Message Transfer Agent in Microsoft Exchange Server were to run on different servers, this would obviously generate some problems. This information also implies one very important fact: A group is the unit of failover in the cluster. Entire groups of resources move from one node to the other, not individual resources. The easiest way to build groups is around storage. Most applications and services require storage to function, as in the Exchange example. Because dependencies can be created only within the same group, every application or resource that uses the same disk must be in the same group as the disk resource. Groups are also generally organized as virtual servers, though more than one virtual server can exist in a group.

Failover and Failback Failover is the process by which a group moves from one server to another. Failover can occur for several reasons:

The administrator has moved a group to a new server.

A resource in the group has exceeded its failure threshold.

The group is configured for failback, and the original owner of the group has returned to service. The third reason provides the definition for failback—that is, the group fails back to a preferred owner if so configured. Preferred

owners are configured as a group property. Failback can be configured to occur either immediately or within a specific time period. Note: For more information on failover and failback, consult the following Knowledge Base articles:•

197047 Failover/Failback Policies on Microsoft Cluster Server•

171277 Microsoft Cluster Server Cluster Resource Failover TimeVirtual Server

Figure 2: Virtual server dependency treeA virtual server is a combination of two resources (for instance, an IP address resource and a network name resource) that work together to present a namespace to a client. Figure 2 illustrates the resources and dependency tree for a virtual server. Note that there is no reason that a group cannot contain multiple virtual servers, but this means that all virtual servers in the group can be owned by only one node at a time. Organizing groups into virtual servers provides finer granularity and better scalability, especially in scenarios where additional nodes are to be added later. Figure 3 illustrates the namespaces that are presented to a client by the cluster nodes. Note: Virtual servers are based on NetBIOS and thus have certain limitations.

Figure 3: Namespaces presented to clients by cluster nodes

See full-sized image.Creating a Virtual ServerYou can construct a virtual server in either of two ways. Although you can create a virtual server manually, the easiest method is to use the Cluster Application Wizard. The Cluster Application Wizard creates a virtual server with the option to configure a cluster server application. Creating a virtual server manually requires that an IP address and network name resource be configured in the appropriate group. For a detailed scenario that shows both procedures, see Appendix B.Printer SharesCluster service makes it possible to host multiple spooler resources. This allows the service to provide high availability for printer shares, along with the ability to distribute load among the cluster nodes. The print spooler resource is limited to one spooler resource per group. It has required dependencies for a network name resource (for consistent access to the spooler) and a physical disk (to store spooler files).After a spooler resource is created, printers must be added to it. In Windows 2000, the job is greatly simplified, because the spooler resource maintains information about clustered printer ports in the cluster configuration database. This eliminates the need to install the ports twice, once on each node. Drivers, however, must be installed on both nodes, because printer drivers are copied to the PRINT\$ share of the remote server.Clustered print spoolers support only standard port monitors and Line Printer Remote (LPR) printers. LPR ports do not support bidirectional printing. Currently, no other ports are supported.When a group containing a print spooler resource fails over to another node, the document that is currently being spooled to the printer is restarted from the other node after the failure. When you move a print spooler resource or take it offline, Cluster service waits until all documents are finished spooling or until the configured wait time has elapsed. Documents that are spooling from an application to a print spooler resource are discarded and must be re-spooled to the resource (or reprinted) if the group containing the resource fails over before the application has finished spooling.Printers hosted by a cluster node are added to Active Directory by the spooler service.Creating a Spooler ResourceThe print spooler resource has required dependencies on a network name and a storage class resource. Figure 2 shows the dependency tree for a spooler resource.

Figure 4: Dependency tree for a spooler resourceOn the resource parameters page of the New Resource Wizard, the only parameters provided are the location of the spool folder and the duration of the job completion timeout. The wizard, based on dependencies, provides these parameters automatically, so there is usually no need to modify them.

Figure 5: Print spooler parameters

See full-sized image.Adding PrintersCreating the virtual server and the spooler resource is only half of the job. The spooler resource is useless without printers. To add a printer, you must first note the owning node of the virtual servers, because the Add Printers Wizard copies the drivers to the owning node only. The driver must then be manually installed on the other node.The Add Printers Wizard must be started from the virtual server—specifically, the network name on which the print spooler depends. After the printer is successfully installed, you must add the driver to the other node. Although several different methods exist to add the drivers to the other node, the command line shown in Figure 6 below starts the Add Printers Wizard directly. You might find it useful to create a shortcut to this command line on a file and print cluster.

Figure 6: Add Printers Wizard command line

See full-sized image.For a detailed scenario, see Appendix C.Adding Additional DriversAdding additional drivers that are not Windows 2000 drivers is relatively uncomplicated. After the initial drivers are installed, you can install additional drivers through the printer user interface (UI). The procedure is as follows:

1.

Connect to the virtual server.

2.

Open the Printers folder.

3.

Right-click the printer to add drivers to, and then click Properties.

4.

Click the Sharing tab, and then click the Additional drivers option.

5.

Once the driver has been added, return to the Printers folder.

6.

Fail the group to the other node.

7.

Repeat steps 3 through 6. ScenarioReskit.com has multiple printer servers it would like to consolidate to one cluster, SEA-NA-CLUS-02. There are approximately 200 printers, and performance is of prime importance. Implementation To achieve better performance, the printers are split roughly evenly between two virtual servers, SEA-NA-PRINT-01 and SEA-NA-PRINT-02. Each virtual server is configured in its own resource group with a print spooler resource. Half the printers are installed on one virtual server with the remaining half allocated for the other virtual server. Preferred owners are configured for each group, and each is assigned to a particular group. Failback is configured to occur between 23 and 0 hours. Procedural Overview

1.

Create virtual servers SEA-NA-PRINT-01 and SEA-NA-PRINT-02, adding a print spooler in each virtual server.

2.

Add printers to the first node.

3.

Add printers to the second node.

4.

Configure failback policy for each group SEA-NA-PRINT-01 as follows:

1.

Right-click the group, and then click Properties.

2.

On the General tab, click Modify to add SEA-NA-CLN-01 as the preferred owners.

3.

In the Available Nodes box, select SEA-NA-CLN-01, and then click the right arrow button to move the node to the Preferred Owners box. Click OK.

4.

Click the Failback tab, and then click Allow failback.

5.

Click Failback between, and then enter 23 and 0 hours. Click OK.

5.

Repeat steps a through e for SEA-NA-PRINT-02 and SEA-NA-CLN-02, respectively.

6.

Move groups to their preferred owners. For a detailed scenario, see Appendix D. Conclusion Because of the importance of

print services to the end users on a computer network, Windows 2000 Advanced Server provides Cluster service, a powerful feature for improving the reliability and performance of print services. However, you must deploy Cluster service properly to provide these benefits to your users. Since clustering is new to Windows, you may be required to look at familiar tasks in new ways. For example, although clustered file shares behave exactly like normal file shares, creating and administering them is a bit different, because they run on multiple machines. Moreover, you must create clustered file shares using the Cluster Administrator in Administrative Tools, and they can involve more complicated administration of permissions. Similarly, the use of clustering for printer services requires that you apply dependencies for a network name resource and a physical disk. Related Links For more information about Windows 2000 Cluster service, see the following resources:

Step-by-Step Guide to Installing Cluster Service

Server Cluster Glossary

Windows 2000 Advanced Server, Clustering Microsoft ITG Infrastructure Services

Cluster Management and Administration

Windows Clustering Technologies See also the following Knowledge Base articles:

259267 Microsoft Cluster Service Installation Resources

258750 Recommended Private "Heartbeat" Configuration on Cluster Server

174812 Effects of Using Autodetect Setting on Cluster NIC

197047 Failover/Failback Policies on Microsoft Cluster Server

171277 Microsoft Cluster Server Cluster Resource Failover Time

171791 Creating Dependencies in Microsoft Cluster Server

198893 Effects of Checking "Use Network Name for Computer Name" in MSCS

195462 WINS Registration and IP Address Behavior for MSCS 1.0

278710 No Global Groups are Available Creating File-Share Permissions Appendix A: Cluster Setup

Figure 7: Cluster Setup

See full-sized image. Installation Notes

The heartbeat cards were set to a speed and configured according to information from the following articles in the Microsoft Knowledge Base: 258750 Recommended Private "Heartbeat" Configuration on Cluster Server
174812 Effects of Using Autodetect Setting on Cluster NIC

Warning: At no time should both systems be booted with access to the shared resources, unless at least one of them has Cluster service. Install the first node with the second node down, and then boot the second node and install Cluster service. FAILURE TO HEED THIS WARNING COULD RESULT IN FILE SYSTEM CORRUPTION THAT MAY BE UNDETECTABLE UNTIL A LATER DATE.

For additional cluster installation resources, see also the Knowledge Base article 259267 Microsoft Cluster Service Installation Resources as well as

<http://www.microsoft.com/windows2000/techinfo/planning/server/clustersteps.asp> Appendix B: Creating a Virtual Server Using the Cluster Application Wizard To start the Cluster Application Wizard, open the Cluster Administrator. On the File menu, click Configure Application.

Figure 8: Cluster Administrator window

See full-sized image. At the Welcome to the Cluster Application Wizard page, click Next.

Figure 9: Cluster Application Wizard

See full-sized image. On the Select or Create a Virtual Server page, you can configure an application for an existing virtual server or create a virtual server. Because this section is about creating virtual servers, Create a new virtual server is selected.

Figure 10: Selecting Create a new virtual server

See full-sized image. The next page is used to select a group for the virtual server or to create a new group. In general, you should use an existing cluster group, unless storage class resource is either not needed or will be created or moved to the new group. The storage class resource to be used by this virtual server is located in Disk Group 1.

Figure 11: Assigning a resource group

See full-sized image. Even though an existing group is selected, you'll want to give the group a new name and description to reflect its current function. In this particular instance, the group is named after the virtual server, but this is not necessary and may not even be desirable because a group can host more than one virtual server.

Figure 12: Naming the resource group

See full-sized image. After you configure the group name and description, the IP address and network name for the virtual server are provided.

Figure 13: Providing a network name and IP address

See full-sized image. After the virtual server properties are provided, the wizard lets you edit detailed information about the group and resources. This is not necessary at this time, so it is skipped.

Figure 14: Setting advanced properties

See full-sized image. At this point, the virtual server is configured. The wizard prompts you to create an application resource. In this example, only the virtual server itself is of interest, so select No, I'll create a cluster resource for my application later.

Figure 15: The Create Application Cluster Resource screen

See full-sized image. The wizard then provides a status page.

Figure 16: Cluster Application Wizard status page

See full-sized image. After you click Finish, the group is created in an offline state. To fully test the virtual server, right-click the group, and then click Bring Online.

Figure 17: Right clicking the group in the Cluster Administrator and clicking Bring Online

See full-sized image. If all is well, the group should come completely online. To check the virtual server, ping from a client by IP address. To test name resolution, ping by name. Using the Manual Method Rather than use the Cluster Application Wizard, you can create a virtual server manually simply by adding an IP address and a network name to an existing group. To begin, right-click the group in which the virtual server will exist, click New, and then click Resource. In this example, the Disk Group 2 was clicked. This starts the New Resource Wizard. The first resource that must be created is an IP address resource. This example shows the creation of the SEA-NA-FILE-02 virtual server. The first step is to identify the resource type. In the Name box, you can type a convenient label. The Description box lets you provide more detailed information.

Figure 18: Setting new resource parameters

See full-sized image. The next page in the wizard lets you specify possible owners. By default, all nodes in the cluster are possible owners. If a node is removed from the possible owners list of a resource, the group in which the resource exists cannot fail over to that node. In this example, the group is intended for failover, so the default is used.

Figure 19: Assigning resource owners

See full-sized image. The next step in correctly adding the resource is to configure resource dependencies. IP addresses have no dependencies, so you can leave this page as is.

Figure 20: Configuring resource dependencies

See full-sized image. The final step is to configure resource-specific properties. An IP address has the following parameters:

Address. The IP address that will be registered on the adapter.

Subnet mask. The subnet mask that will be used for this IP address. This field is populated automatically.

Network. The network on which the IP address will be registered.

Enable NetBIOS for this address. Clear this check box if NetBIOS is not needed for this virtual server. NetBIOS is required for browsing to function properly. There is a limit of 64 NetBT devices in Windows. If this limit is exceeded, additional IP address resources cannot use NetBIOS.

Figure 21: Enabling NetBIOS check box

See full-sized image. When you click Finish, the IP address is created and a confirmation message appears.

Figure 22: Resource confirmation dialog box

See full-sized image. Now you must configure a network name for the virtual server. The procedure is much the same, with the exception of the way you select the type of resource dependencies and resource-specific parameters. You start the New Resource Wizard from Disk Group 2 by right-clicking the group, clicking New, and then clicking Resource. The resource type is, of course, a network name.

Figure 23: Configuring a network name

See full-sized image. Leave Possible owners at the default setting.

Figure 24: The Possible Owners screen

See full-sized image. A dependency is configured for the previously created IP address. A dependency on an IP address is required for a network name.

Figure 25: Configuring resource dependencies

See full-sized image. The only resource-specific parameter for a network name is the name itself. This is the name of the virtual server on the network, and it must conform to standard computer naming rules.

Figure 26: Naming the virtual server

See full-sized image. Click Finish to create the network name. At this point, you can bring the resource group online in the same manner described in the Cluster Application Wizard example. To give the group a more descriptive name, right-

click the group, and then click Rename on the shortcut menu.

Figure 27: Right-clicking a group to rename the resource

See full-sized image. Appendix C: Creating a Print Spooler Resource with the Cluster Application Wizard In this example, the Cluster Application Wizard is used to create the virtual server SEA-NA-PRINT-01 and configure the print spooler resource. To start the Cluster Application Wizard, open the Cluster Administrator. On the File menu, click Configure Application.

Figure 28: Cluster Application Wizard

See full-sized image. In this example, the intent is to create a new virtual server and configure the application in one step. Select Create a new virtual server, and then click Next.

Figure 29: Selecting Create a new virtual server

See full-sized image. For this example, select an existing resource group with a storage class resource. This virtual server uses the storage class resource in Disk Group 3.

Figure 30: Selecting Use an existing resource group

See full-sized image. The group name and description are provided.

Figure 31: Naming the resource group

See full-sized image. No additional configuration items are needed for the Advanced Properties page. Click Next to create the virtual server.

Figure 32: Clicking Next to create the virtual server

See full-sized image. On the next screen, Yes, create a cluster resource for my application now is selected. Click Next.

Figure 33: Accepting the default selection

See full-sized image. The resource type is Print Spooler. Click Next.

Figure 34: Selecting Print Spooler as the resource type

See full-sized image. The name and description for the resource are provided. To add dependencies, click Advanced Properties.

Figure 35: Clicking Advanced Properties

See full-sized image. In the Advanced Resource Properties dialog box, click the Dependencies tab, and then click Modify.

Figure 36: Setting resource dependencies

See full-sized image. In the Modify Dependencies dialog box, move the appropriate dependencies from the Available Resources list to the Dependencies list.

Figure 37: Modifying dependencies

See full-sized image. After you have configured the dependencies, click OK to close all the dialog boxes until you see the screen below. The resource is created, and the wizard prompts you for resource-specific parameters. The wizard fills in the parameters based on dependencies. Usually, no additional modification is needed, so you can click Next.

Figure 38: The provided print spooler parameters generally do not need to be modified.

See full-sized image. The virtual server and spooler resource are now configured. The group can now be brought online. Adding a Printer In this example, the connection must be made to SEA-NA-PRINT-01. To start the connection, on the Start menu, click Run.

Figure 39: Connecting to the virtual server

See full-sized image. The Printers folder is selected.

Figure 40: Selecting the Printers folder

See full-sized image. From the Printers folder, start the Add Printer Wizard. From a remote connection, the only option is to add a printer to the remote print server.

Figure 41: Only Remote print server can be selected.

See full-sized image. Because this is a new printer, you must create a new port.

Figure 42: Selecting Create a new port

See full-sized image. The Add Standard TCP/IP Printer Port Wizard is started.

Figure 43: The AddStandardTCP/IPPrinterPort Wizard welcome screen

See full-sized image. The port name field is populated as soon as you type the printer name or IP address.

Figure 44: The Port Name field is populated automatically

See full-sized image. After you configure the port, the Add Printer Wizard prompts you for a driver. In this example, a Lexmark Optra is selected.

Figure 45: Selecting the correct printer driver

See full-sized image. You then name the printer.

Figure 46: Naming the printer

See full-sized image. The share name is populated automatically, but you can change it.

Figure 47: Setting printer sharing preferences

See full-sized image. You can then provide information about the printer's location and capabilities.

Figure 48: Adding printer location and comments

See full-sized image. When you click Next, the status screen shows you that the printer has been installed.

Figure 49: The Add Printer Wizard status screen for a completed installation

See full-sized image. Adding Drivers to the Other Node To start the connection, on the Start menu, click Run. The command shown in the Figure 50 starts the Add Printer Driver Wizard.

Figure 50: Starting the Add Printer Driver Wizard

See full-sized image. You can use this wizard to install the necessary print drivers for the other node.

Figure 51: Selecting the appropriate printer driver

See full-sized image. This step is necessary only for the first Windows 2000 printer driver. After you complete this step, you can install additional drivers for other platforms through the printer user interface.

Figure 52: Completing the Add Printer Driver Wizard

See full-sized image. Appendix D: Configuring Failback This section describes how to configure a group for failback to a preferred owner. The example is drawn from the scenario for creating normal file shares and configuring failback. To begin configuring failback, right-click the group, and then click Properties.

Figure 53: Right-clicking a group to open its Properties dialog box

See full-sized image. On the General tab in the group Properties dialog box, click Modify.

Figure 54: Modifying group properties

See full-sized image. In the Modify Preferred Owners dialog box, add the desired preferred owner.

Figure 55: Nodes available for addition to the preferred owners

See full-sized image. This example is drawn from the first scenario, so SEA-NA-CLN-01 is added.

Figure 56: Adding nodes to the preferred owners

See full-sized image. On the Failback tab, the default option of Prevent failback is changed to Allow failback, with Failback between configured for 23 and 0 hours.

Figure 57: Configuring failback settings

See full-sized image. When you click OK, the process is complete and the group is configured for failback to a preferred owner.