

Parallels Virtuozzo Containers 4.6 for Windows

Deploying Microsoft Clusters

*Parallels Holdings, Ltd.
c/o Parallels International GmbH.
Parallels International GmbH
Vordergasse 49
CH8200 Schaffhausen
Switzerland
Tel: + 41 526320 411
Fax: + 41 52672 2010
www.parallels.com*

Copyright © 1999-2015 Parallels Holdings, Ltd. and its affiliates. All rights reserved.

This product is protected by United States and international copyright laws. The product's underlying technology, patents, and trademarks are listed at <http://www.parallels.com/trademarks>.

Microsoft, Windows, Windows Server, Windows NT, Windows Vista, and MS-DOS are registered trademarks of Microsoft Corporation.

Apple, Mac, the Mac logo, Mac OS, iPad, iPhone, iPod touch, FaceTime HD camera and iSight are trademarks of Apple Inc., registered in the US and other countries.

Linux is a registered trademark of Linus Torvalds.

All other marks and names mentioned herein may be trademarks of their respective owners.

Contents

Microsoft Cluster Types	5
Failover Clusters	5
NLB Clusters	5
Deploying Parallels Containers Failover Clusters	6
Designing a Cluster Configuration	7
Active/Passive Failover Clusters	7
Active/Active Failover Clusters	9
Setting Up Hardware for the Cluster	11
Configuring Physical Servers	12
Setting Up a Domain.....	12
Creating a Domain.....	13
Adding Nodes to the Domain	20
Connecting Servers to a Shared Storage	22
Setting Up a Cluster	23
Creating a Cluster Service Account.....	23
Deploying a Failover Cluster in Windows Server 2008 and Windows Server 2008 R2	25
Deploying a Failover Cluster in Windows Server 2003	38
Installing Parallels Containers on Cluster Nodes.....	60
Configuring Automatic Failback	65
Adding Nodes to the Failover Cluster	66
Testing the Failover Cluster	67
Troubleshooting Failover Clusters.....	69
Creating Clusters From Containers	72
Setting Up a SAN-Based Cluster.....	72
Setting Up a 'Loopback File'-Based Cluster	74
Setting Up an NLB Cluster	78
Index	82

Microsoft Cluster Types

The Windows Server operating system comes with two main built-in clustering technologies:

- *Failover cluster*
- *Microsoft Network Load Balancing (NLB)*

Failover Clusters

Failover clusters (known as *server clusters* in Windows Server 2003 and earlier Windows Server versions) allow you to provide high availability for your mission-critical applications (including the Parallels Containers software itself) through the failover of resources. In Parallels Containers, you can deploy failover clusters in one of the following ways:

- Create a failover cluster of two or more Hardware Nodes that will host the Parallels Containers software including all its mission-critical services and Containers. If one Hardware Node fails, another Node will take its responsibilities.
- Make Containers participate as full members in the cluster, like any other standalone server running the Windows Server operating system, and increase the availability of your mission-critical applications.

The process of deploying failover clusters in Parallels Containers-based systems is described in *Deploying Parallels Containers Failover Clusters* (p. 6) and *Creating Clusters From Containers* (p. 72).

Note: Throughout this guide, we use the term *failover cluster* to refer to the clustering technology in Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003 though in Windows Server 2003 this technology is called *server cluster*. When there are differences between these two clustering models, we describe them in separate sections.

NLB Clusters

The *NLB* technology allows you to provide scalability and high availability for TCP- and UDP-based services and applications (e.g., for Web servers and Terminal Services) by load-balancing incoming IP traffic across several servers in the cluster. For detailed information on deploying NLB clusters, see *Setting Up an NLB Cluster* (p. 78).

Deploying Parallels Containers Failover Clusters

In Parallels Containers 4.6, you can use the Windows failover cluster feature to create failover clusters consisting of two or more Hardware Nodes. Including Hardware Nodes in a failover cluster increases the availability of the Parallels Containers software and Containers. If one Hardware Node in the cluster fails, another Node will take its responsibilities.

The process of deploying a Parallels Containers failover cluster includes the following steps:

- 1 Choosing the cluster model that best suits your needs.
- 2 Setting up hardware for the cluster.
- 3 Configuring all physical servers to participate in the cluster (hereinafter, we will also refer to them as *cluster nodes* or *nodes*).
- 4 Setting up an Active Directory domain and adding the cluster nodes to it.

Note: You can also add the cluster nodes to an existing Active Directory domain, if you have one.

- 5 Setting up a shared storage and connecting it to the cluster nodes.
- 6 Creating a failover cluster and adding all your nodes to the cluster.
- 7 Setting up the failover cluster.

All these steps are described in the following subsections in detail.

In This Chapter

Designing a Cluster Configuration.....	7
Setting Up Hardware for the Cluster.....	11
Configuring Physical Servers.....	12
Setting Up a Domain	12
Connecting Servers to a Shared Storage	22
Setting Up a Cluster	23
Installing Parallels Containers on Cluster Nodes	60
Troubleshooting Failover Clusters	69

Designing a Cluster Configuration

Before you start deploying a Parallels Containers failover cluster, you should decide on the cluster configuration that will meet your demands best of all. You can create Parallels Containers failover clusters of one of the following types:

- *Active/passive clusters.* An active/passive cluster includes both active and standby nodes. In this type of clusters, standby nodes are only used if one of the active nodes fails.
- *Active/active clusters.* An active/active cluster consists of active nodes only, each running the Parallels Containers software and hosting a number of Containers. In the event of a failover, all Containers running on the problem node are failed over to one of the healthy active nodes.

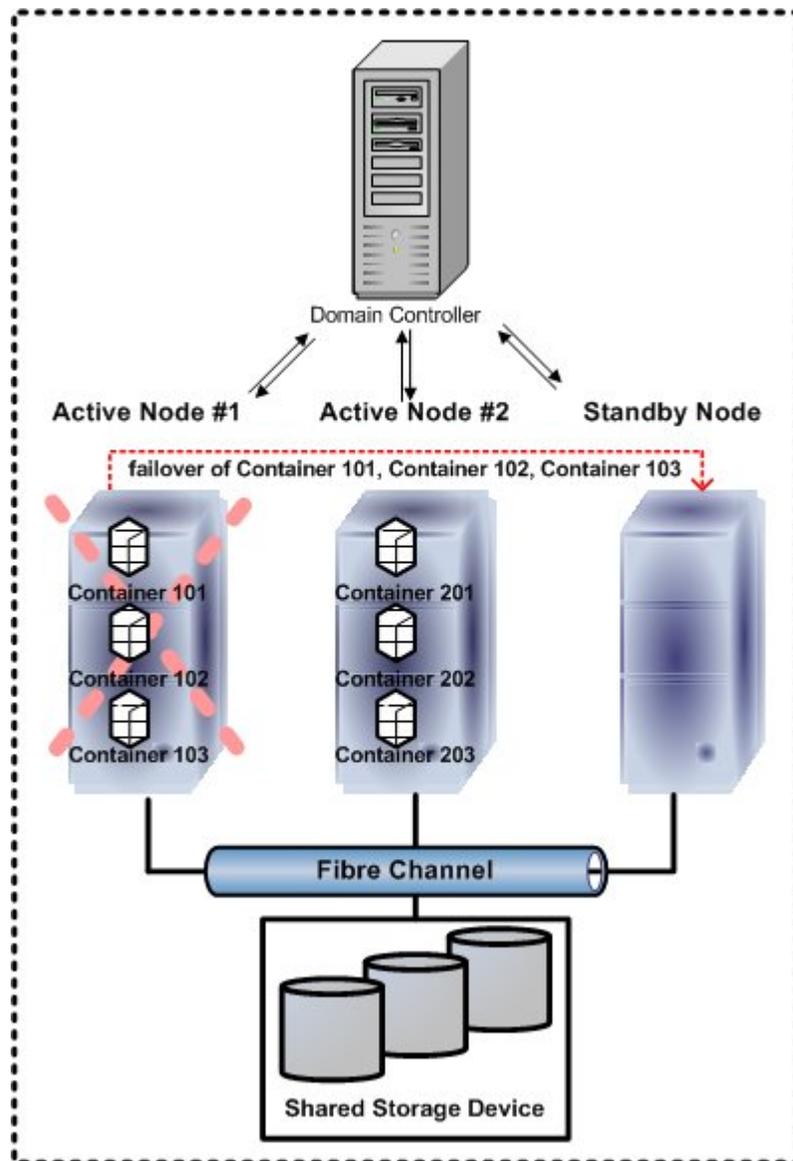
The following subsections discuss both types of failover clusters in detail.

Active/Passive Failover Clusters

Active/passive clusters include both active and standby nodes. An example of an active/passive cluster is a four-node cluster where three nodes are active, running the Parallels Containers software and hosting a number of Containers, and one node is in the standby state, waiting for an active node failure and ready to take the responsibility of its resources. An active/passive cluster can also consist of five active cluster nodes and two standby nodes where each of the active nodes is able to fail over to any of the two standby nodes.

The following picture shows an example configuration of an active/passive failover cluster:

Note: We will use the given configuration throughout this chapter to demonstrate the Parallels Containers failover cluster creation.



In this clustering scenario:

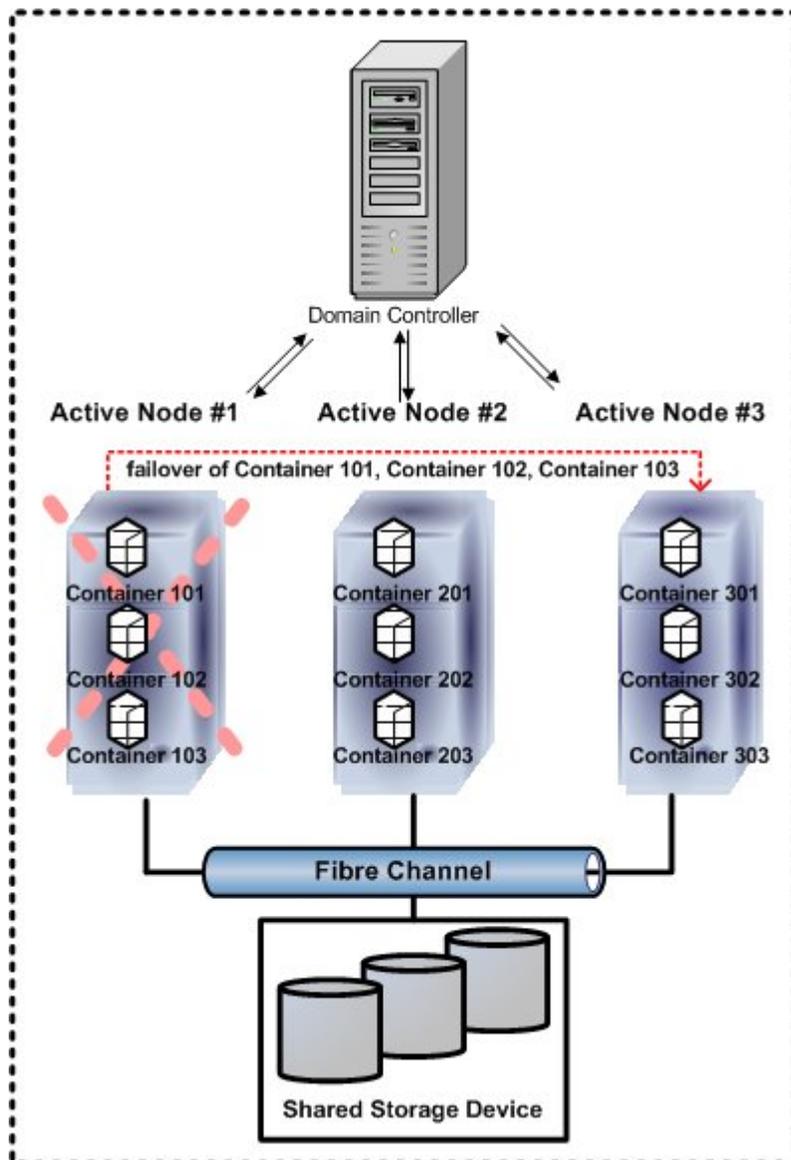
- The cluster comprises two active nodes (*Active Node #1* and *Active Node #2*) hosting a number of Containers and one passive node (*Standby Node*). All nodes in the cluster have one and the same version of the Windows Server operating system and the Parallels Containers software installed. If one of the active nodes in the cluster fails or is taken offline as part of planned maintenance, the services and Containers from this node are failed over to the standby node.
- All the nodes in the cluster are connected via Fibre Channel optical cables to a reliable and fault tolerant shared SCSI storage device. It is used to store all the cluster configuration data and Parallels Containers resources and allows active nodes in the cluster to access these resources.
- All the cluster nodes are located in the same Active Directory (AD) domain. A dedicated server is set up in this domain to act as the domain controller (shown as *Domain Controller* in the picture).

Note: In our example, we will create an Active Directory domain from scratch. However, you can make use of an existing AD domain and include all your cluster nodes in this domain.

Active/Active Failover Clusters

Parallels Containers allows you to create active/active failover clusters where all Hardware Nodes are acting as active nodes. The concept of active/active clusters is based on the ability to simultaneously run more than one group with Parallels Containers resources (i.e. with Parallels Containers-related services and Containers) on one and the same Hardware Node. In such clusters, each node is capable of taking on the additional workload, the failed over services and Containers, from a problem node along with running their own version of Parallels Containers and hosting their own Containers.

The following picture shows you an example of an active/active cluster configuration:



This clustering configuration is similar to that described in the previous section for an active/passive cluster. The cluster comprises three nodes: *Active Node #1*, *Active Node #2*, *Active Node #3*. All nodes are running one and the same version of the Windows Server operating system and the Parallels Containers software, located in the same AD domain, and connected via Fibre Channel optical cables to a shared SCSI storage device.

However, unlike the active/passive configuration, all nodes in this cluster are active and host a number of running Containers. If any of the nodes in the cluster fails or is taken offline as part of planned maintenance, the Parallels Containers-related services and Containers from this node are automatically failed over to one of the two remaining active nodes. If, say, *Active Node #1* hosting Container 101, 102, and 103 fails, all these Containers will be moved to and started on either *Active Node #2* or *Active Node #3*. The decision on where to move the resources from the problem node is made by the clustering software.

Setting Active/Active Failover Cluster to Work

The process of deploying an active/active Parallels Containers failover cluster slightly differs depending on whether you are creating your failover cluster from scratch or transforming any of your existing active/passive clusters:

- If you are creating a new Parallels Containers failover cluster, follow the instructions given in the following subsections to set up the cluster.
- If you already have a Parallels Containers active/passive cluster and want to transform it into an active/active one, do the following:
 1. Reinstall Parallels Containers on each passive node in the cluster by deleting the current software version and installing it again using the instructions in the **Installing Parallels Containers on Cluster Nodes** section (p. 60).

Note: When reinstalling the Parallels Containers software, you must select the **Active node** radio button for each node that will participate in the active/active cluster.

2. Install all available Parallels Containers updates. Refer to the *Parallels Containers 4.6 User's Guide* to learn how you can do it.

Important! We highly recommend that you install the same set of Parallels Containers and Windows Server updates on all nodes in the Parallels Containers failover cluster.

Active/Active Failover Cluster Restrictions

Before starting to deploy an active/active Parallels Containers failover cluster, you should have a clear understanding of all limitations imposed on your virtual infrastructure by this kind of cluster:

- Parallels Containers failover clusters do not currently have a mechanism to correctly process Containers with identical IDs. So, you should always see to it that all Containers residing on all nodes in the cluster have different IDs.
- Containers failed over from a problem node to a healthy one make use of the OS and application templates installed on the healthy node. So, to provide an error-free operation of your Containers, make sure that all nodes in the cluster have the same set of OS and application templates installed.

- Deploying an active/active cluster requires the proper hardware configuration of all cluster nodes. As the resources from a problem node may be failed over to any of the remaining active nodes, all nodes in the active/active cluster must have enough resources (CPU, main memory, etc.) to take on the additional workload from a failed node and ensure adequate performance in the case of a failover.

If any of the nodes does not have sufficient resources to support this additional workload (e.g. the node's own services and Containers consume the majority of system resources), it is recommended to deploy an active/passive failover cluster to avoid degradation in your overall cluster performance.

Setting Up Hardware for the Cluster

Once you know which cluster model to deploy, you can start setting up hardware for your cluster. This procedure includes the following basic operations:

- 1 Deciding on the hardware to be included in the cluster. The amount and type of hardware may vary depending on the purpose and availability requirements of your cluster. However, the following types of hardware are always present in any cluster:
 - *Cluster nodes* represented by two or more physical servers capable of running the Windows Server operating system.
 - *Shared storage device* usually represented by a common disk array on a SAN and used to store all Parallels Containers-related data (Container data, templates, backups, and so on).
 - *Switches* (Fibre Channel or SCSI) providing client access to the cluster and enabling the communication between each cluster node and the shared cluster storage.
- 2 Uniting all hardware components into one subnet and ensuring that each hardware component can access all the other components in the subnet.

For more information about installing and configuring cluster hardware, follow these links:

- for Windows Server 2003:
<http://technet2.microsoft.com/windowsserver/en/library/32c40202-1043-4211-8dba-bb57356f46811033.msp?mfr=true>
- for Windows Server 2008 and Windows Server 2008 R2:
[http://technet.microsoft.com/en-us/library/dd197454\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd197454(WS.10).aspx)

To deploy our example cluster configuration that will include two active nodes and one passive node, we use the following hardware:

- A shared SCSI storage device.
- A switch.
- Fibre optic cables.
- Four physical servers capable of running the Windows Server operating system. Three servers will act as cluster nodes and one server will be used as a domain controller.

All the physical servers in our example meet the following hardware requirements:

- Pentium IV processor

- 2048 MB of memory
- 80 GB of disk space
- two network interface cards (NICs) for each cluster server and one network card for the domain controller

Configuring Physical Servers

After you have prepared the hardware for your failover cluster, you need to configure three physical servers for participating in the cluster and one physical server for converting into a domain controller. To do this, perform the following operations:

- 1 Install the Windows Server operating system on each server (mandatory). The process of installing Windows Server is described in detail in the documentation shipped with this operating system.

Note: All cluster nodes must run the same version of the operating system, for example, the 64-bit version of Windows Server 2008 (Enterprise Edition).

- 2 Install the failover clustering feature on the physical server (mandatory). This operation must be performed only on physical servers running Windows Server 2008 and Windows Server 2008 R2. For more information, see <http://technet.microsoft.com/en-us/library/cc772178.aspx>.

- 3 Configure the network adapters on all physical servers (mandatory):

For each cluster node:

- Assign a static IP address to the first network adapter. This adapter will be used to connect the server to the private network providing communication among all nodes in the cluster.
- Assign a static IP address to the second network adapter. This adapter will be used to connect the server to the public network (usually, your local area network). In comparison to the IP address assigned to the first network adapter, this IP address must belong to a different subnet.

For the domain controller:

- Assign a static IP address to the network adapter. This adapter will be used to connect the domain controller to other servers in the Active Directory domain.

- 4 Validate your physical server configuration (optional though highly recommended). This operation can be performed only on physical servers running Windows Server 2008 or Windows Server 2008 R2. For more information, see [http://technet.microsoft.com/en-us/library/cc732035\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc732035(WS.10).aspx).

Setting Up a Domain

In the next step, you need to set up an Active Directory domain and make your servers members of this domain. This procedure includes the following steps:

- 1 Deploying an Active Directory domain.

- 2 Creating a domain administrator account.
- 3 Adding the cluster nodes to the domain.

All these steps are described in the following subsections in detail.

Notes:

1. If you already have an existing Active Directory domain, you can add your cluster nodes to this domain. In this case you can skip Steps 1 and 2 and proceed directly to Step 3.
 2. Throughout this section, we assume that you are setting up a domain on a server running Windows Server 2008 R2. The steps you need to perform to set up a domain on a server running Windows Server 2003 or Windows Server 2008 may slightly differ. For information on creating and managing domains under these operating systems, refer to their documentation.
-

Creating a Domain

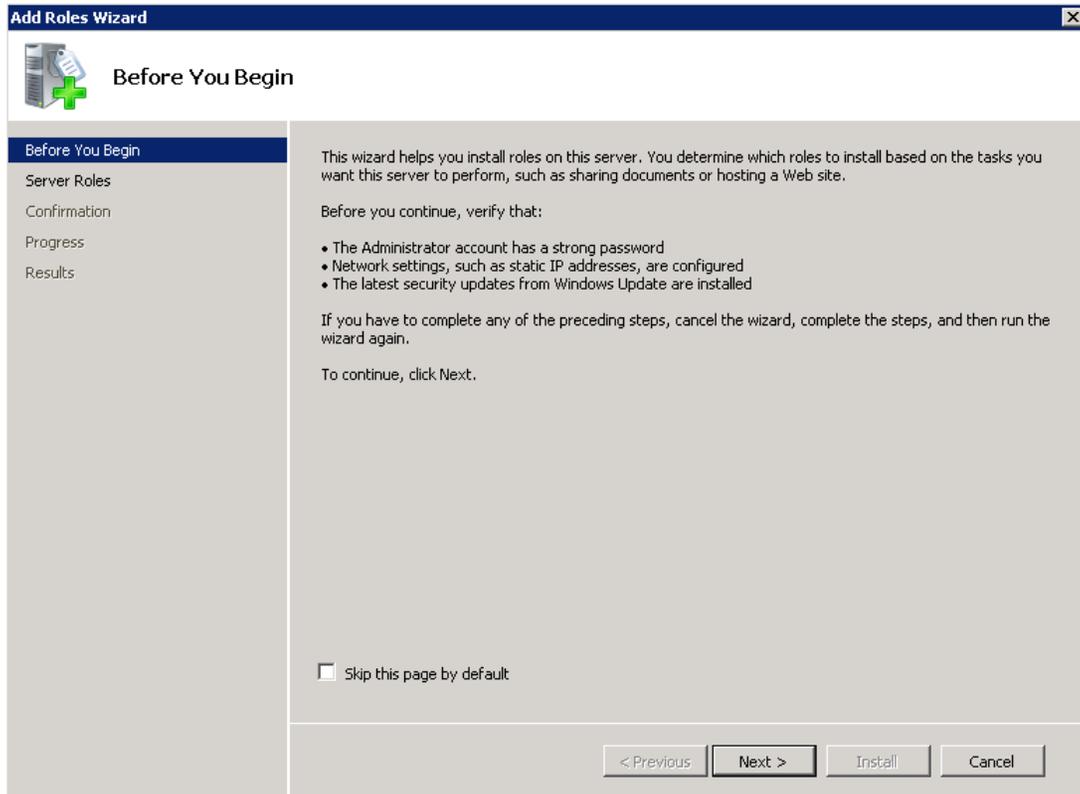
First, you have to set up an Active Directory domain. To do this, you need a dedicated physical server capable of running the Windows Server operating system. This server will act as a domain controller and have the Active Directory service installed. The process of setting up an Active directory domain includes two steps:

- 1 Installing the Active Directory Domain Controller role.
- 2 Creating the Active Directory domain.

Installing the Active Directory Controller role

In the first step, you need to install the Active Directory Domain Controller role on the physical server that is to act as the domain controller. By default, this role is not installed on servers during the Windows Server 2008 R2 installation. To install the role:

- 1 Open the Server Manager console.
- 2 Click the **Roles** item in the left pane of the console.
- 3 In the **Roles** pane, click **Add Roles** to launch the **Add Roles** wizard.

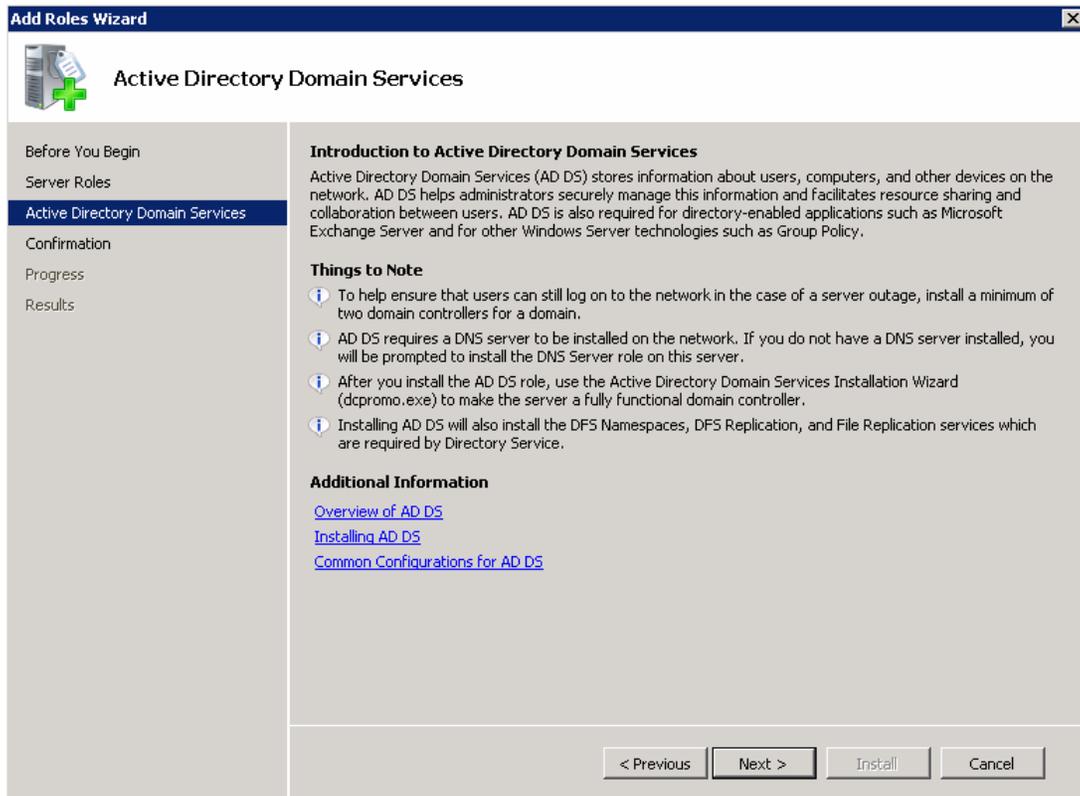


4 Click Next.

5 Select the check box next to **Active Directory Domain Services**. The wizard may ask you to install a number of additional features required for this role. For example:



Click Add Required Features, and then click Next.

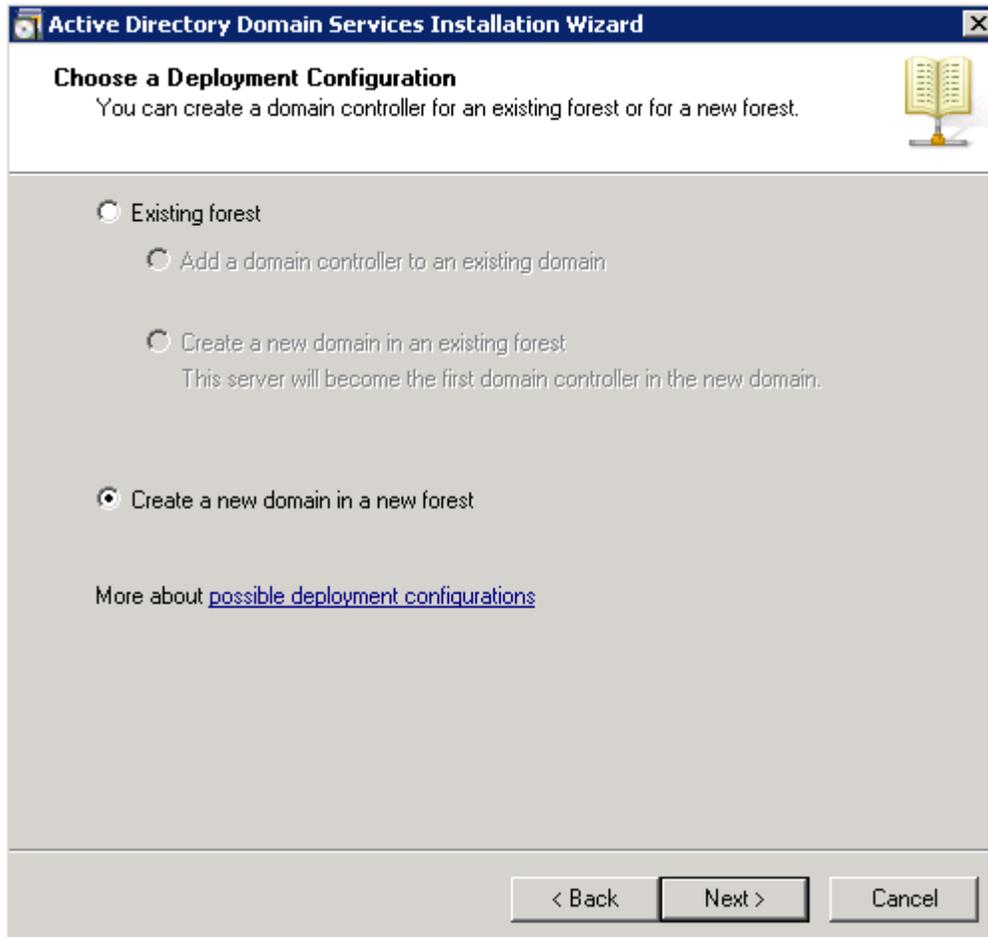


- 6 Read the information, and click **Next**.
- 7 In the **Confirm Installation Selections** window, review the roles and features that will be installed on the server, and click **Install**.
- 8 Once the installation is complete, click **Close**.

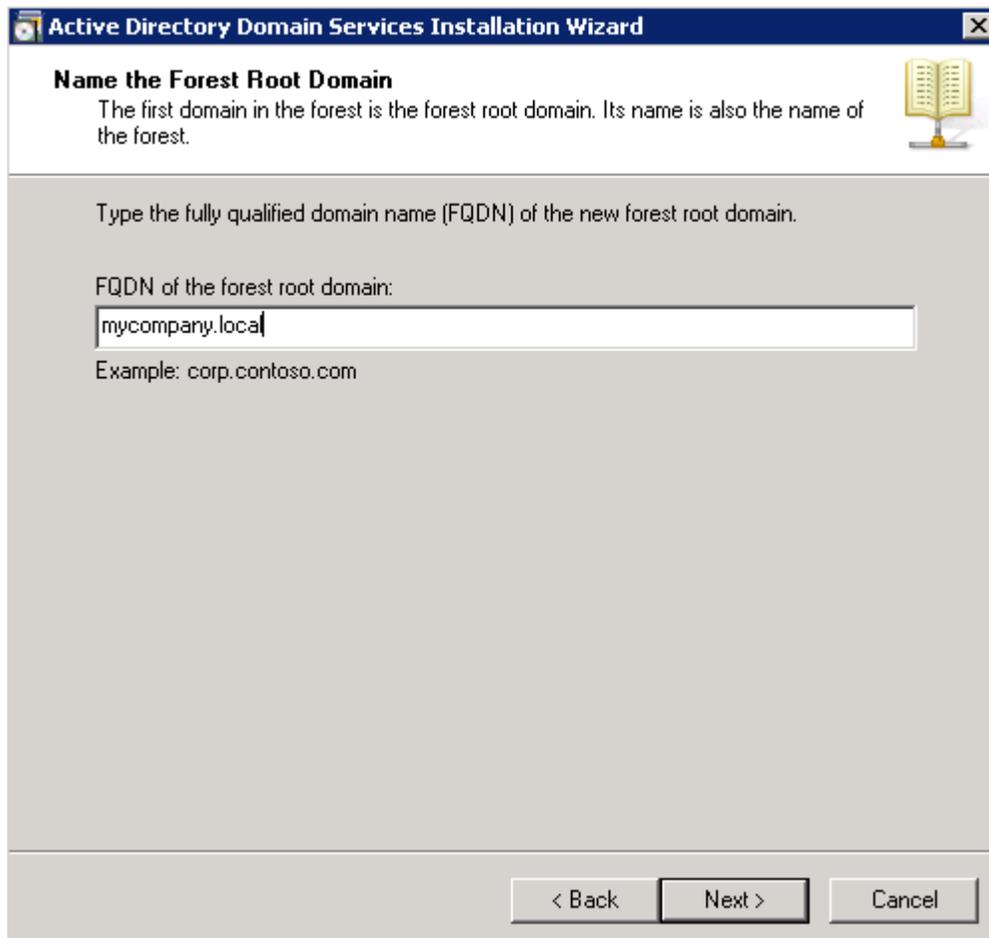
Creating the Active Directory domain

Now that you have installed the necessary roles and features, you can start creating the Active Directory domain. To do this:

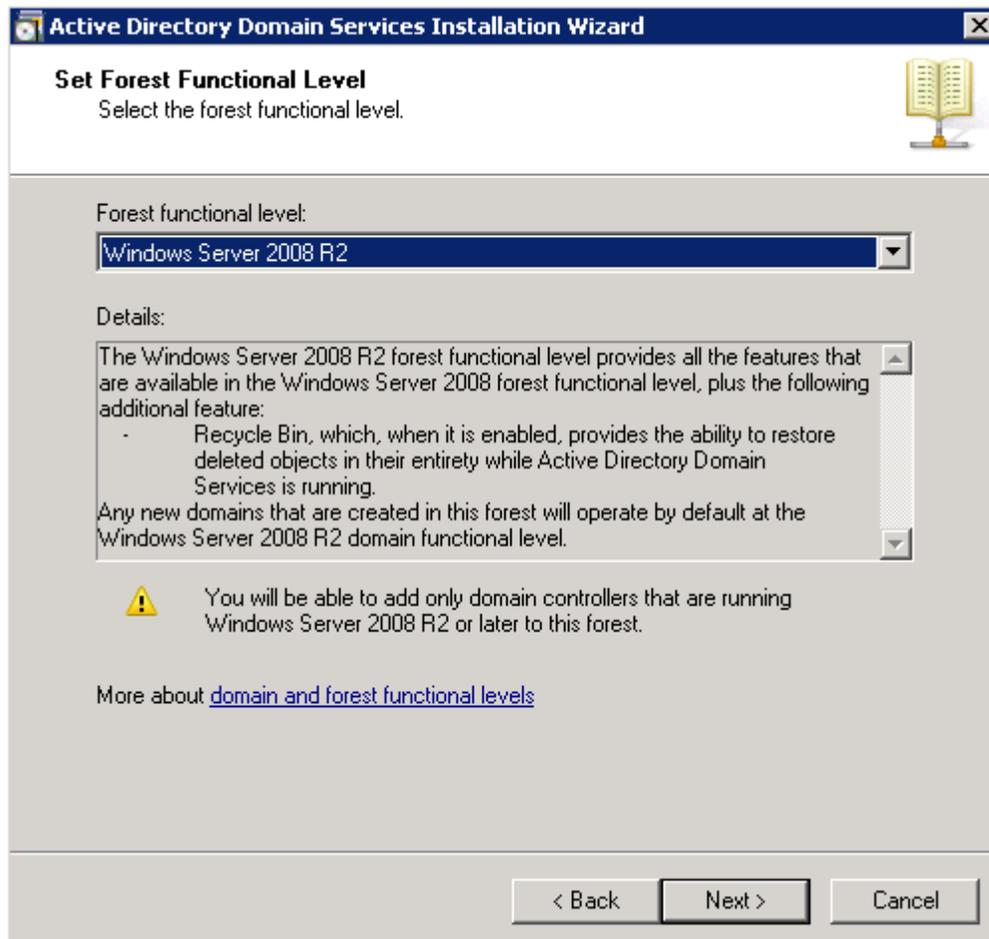
- 1 Click **Start > Run**, and execute the `dcpromo` command in the **Run** dialog to launch the **Active Directory Installation** wizard.
- 2 In the **Welcome** window, click **Next**.
- 3 In the **Operating System Compatibility** window, read the displayed information, and click **Next**.
- 4 In the **Choose a Deployment Configuration** window, select the **Create a new domain in a new forest** radio button, and click **Next**.



- 5 In the Name the Forest Root Domain window, specify a full DNS name for the new domain, and click Next. You can choose any descriptive name you like. In our example, we will assign the `mycompany.local` name to the domain.



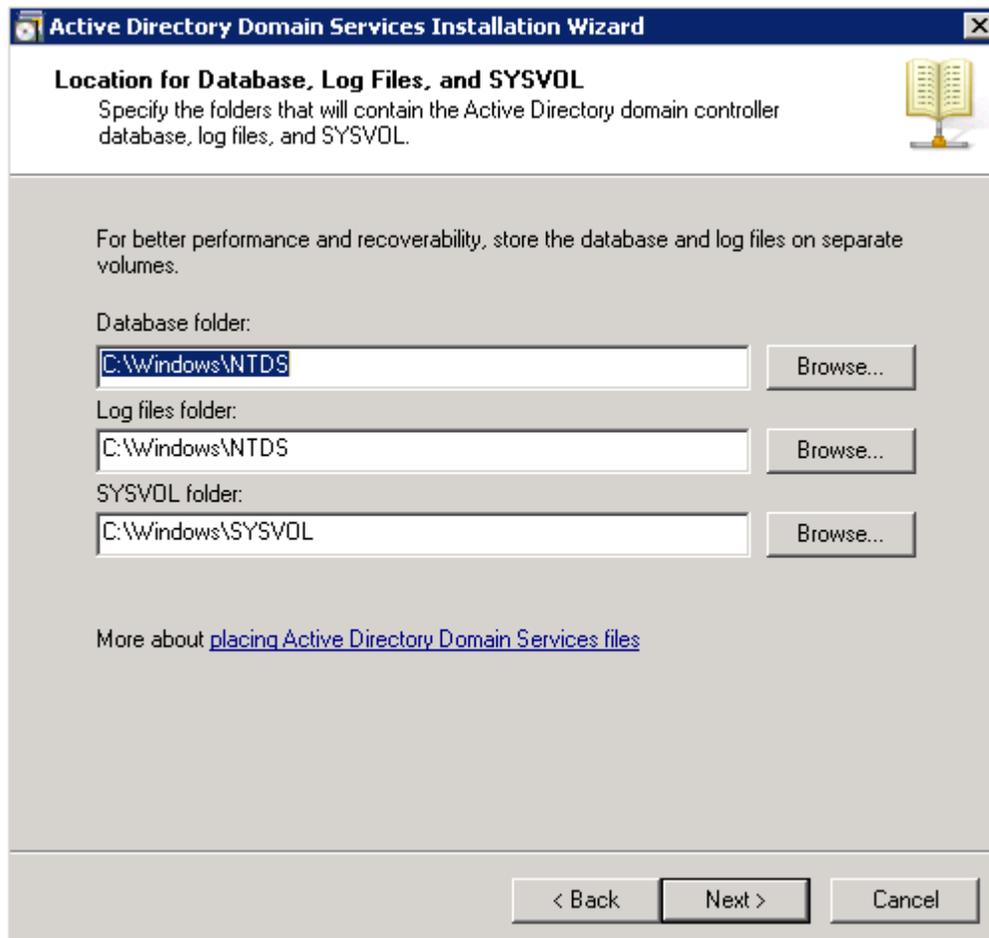
- 6 In the Set Forest Functional Level window, select the forest functional level. If your cluster will contain only nodes with Windows Server 2008 R2, select the Windows Server 2008 R2 option, and click Next.



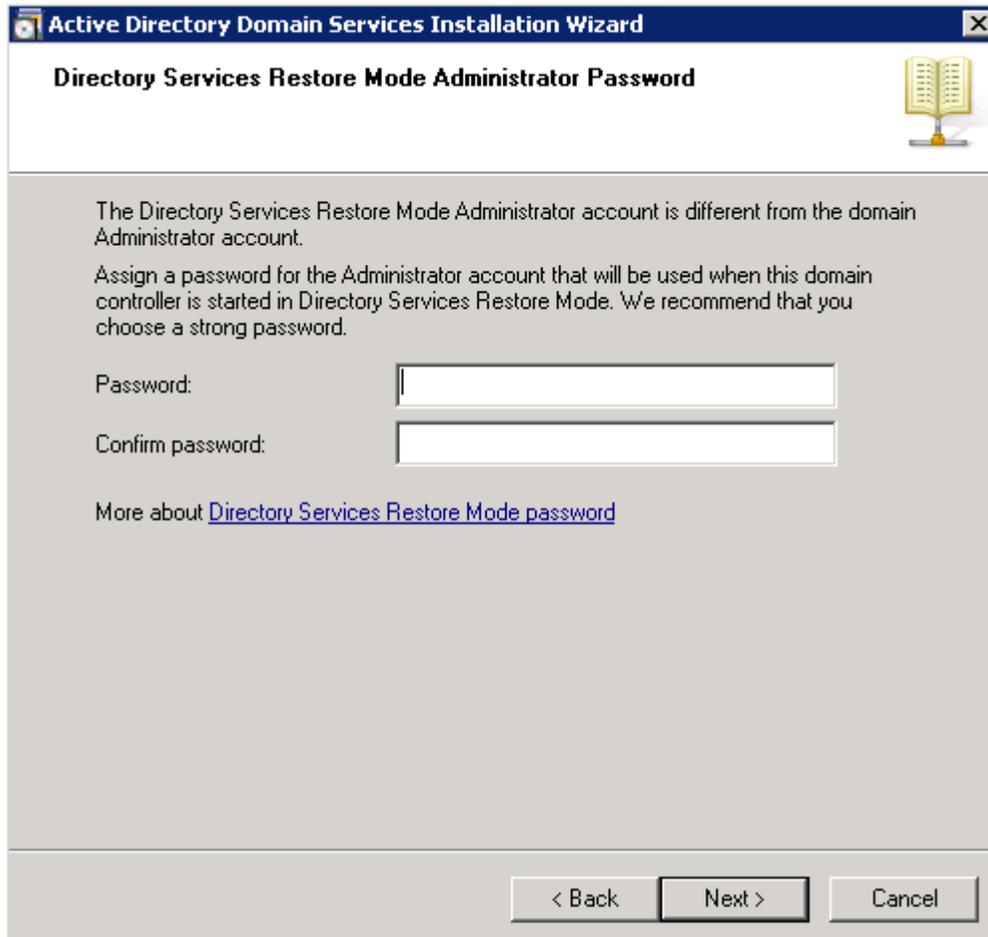
- 7 In the Additional Domain Controller Options window, make sure the DNS server option is selected, and click Next. If you see the following warning, ignore it by clicking Yes.



- 8 In the Location for Database, Log Files, and SYSVOL window, accept the default values offered by the wizard, and click Next.



- 9 In the **Directory Services Restore Mode Administrator Password** window, specify an arbitrary password for the Directory Services Restore Mode administrator. You will need this password to boot the domain controller in Directory Services Restore Mode (e.g., if the Active Directory service has failed or needs to be restored). Make sure the password meets the Windows length and complexity requirements. Click **Next**.



The screenshot shows a Windows wizard window titled "Active Directory Domain Services Installation Wizard". The current step is "Directory Services Restore Mode Administrator Password". The window contains the following text: "The Directory Services Restore Mode Administrator account is different from the domain Administrator account. Assign a password for the Administrator account that will be used when this domain controller is started in Directory Services Restore Mode. We recommend that you choose a strong password." Below this text are two input fields: "Password:" and "Confirm password:". At the bottom of the window are three buttons: "< Back", "Next >", and "Cancel".

10 The last window allows you to review the parameters provided by you in the previous steps of the wizard. To modify any parameters, click **Back**; otherwise, click **Next**.

11 A window will open showing you the operation progress. Select the **Reboot on completion** check box to automatically restart the server once the Active Directory domain is created.

For more information on creating and configuring Active Directory domains, refer to <http://support.microsoft.com/kb/324753/en-us>.

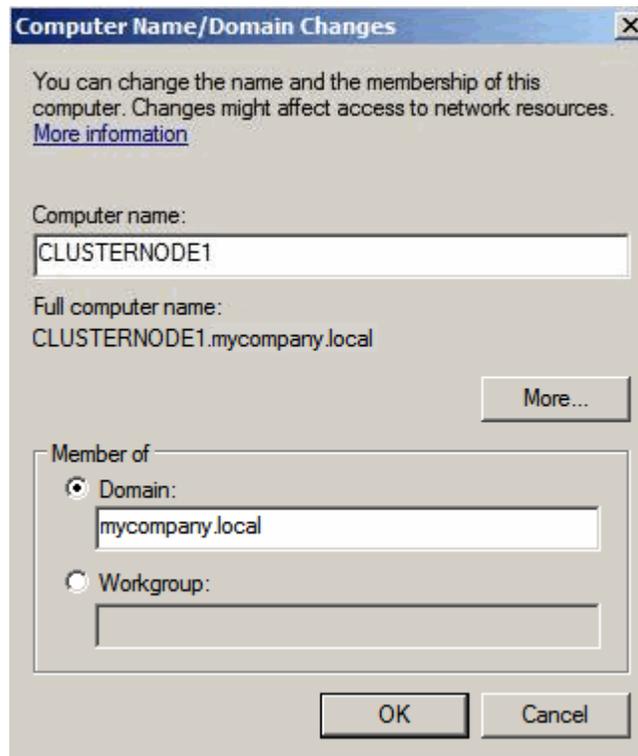
Adding Nodes to the Domain

Now that you have created the Active Directory domain and the domain administrator account, you can add your nodes to the domain. To do this:

- 1** Log in to the first node you want to add to the domain, and click **Start > Control Panel > System and Security > System**.
- 2** In the **Computer name, domain, and workgroup settings** section, click **Change settings**.
- 3** On the **Computer Name** tab of the displayed window, click **Change**.

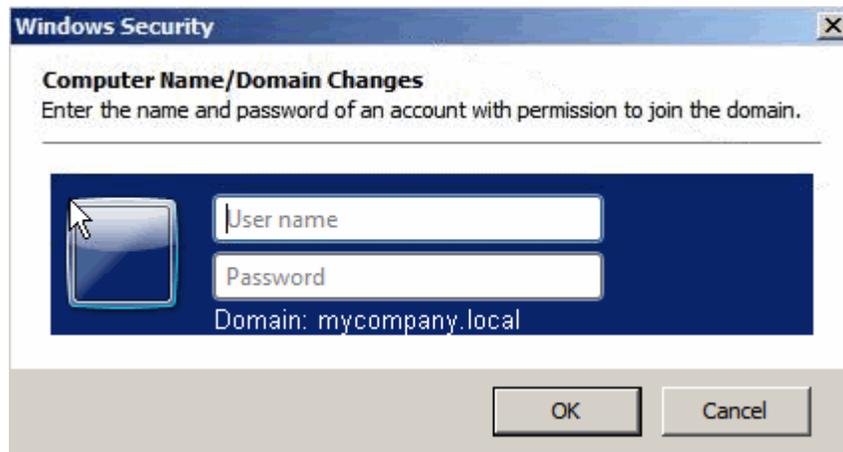
- 4 In the Computer Name Changes window, do the following:
 - In the **Computer name** field, specify a server hostname. This name will be used to uniquely identify the given node among other nodes in the cluster. By default, you are offered to use the hostname assigned to the node during the Windows Server installation. However, we recommend that you change this hostname to something more descriptive (e.g., CLUSTERNODE1).
 - Select the **Domain** radio button and type the domain DNS name (you specified this name during the Active Directory domain). In our example, the domain DNS name must be set to mycompany.local.

After providing the necessary information, your window should look like the following:



When you are ready, click OK.

- 5 In the Computer Name Changes window, type the username and password of the domain administrator account, and click OK.



- 6 Click OK to close the displayed message welcoming you to the domain, and then click OK once more to close the Computer Name Changes window.
- 7 Restart the node.
- 8 Perform Steps 1-7 for both remaining nodes.

Connecting Servers to a Shared Storage

Next, you should configure the shared storage and connect it to all your cluster nodes. The way of configuring storage devices and connecting them to cluster nodes varies from vendor to vendor. So we recommend that you refer to the vendor's guidelines and recommendations to properly set up your storage.

For your reference, we list below the general steps that you usually have to perform to configure a shared SCSI device:

- 1 Power off all your cluster nodes.
- 2 Prepare the necessary controllers for the shared storage bus.
- 3 Connect the SCSI device to the shared bus, and then join the bus to your cluster nodes.
- 4 Power on the first cluster node.
- 5 Format all the disks on the cluster storage device as NTFS, and configure them as basic disks.
- 6 Partition the disks on the cluster storage device.

Pay special attention to the disk to be allocated to the cluster quorum resource. The quorum resource maintains the cluster configuration data; so, we highly recommend that you use a separate disk (ca. 500 MB in size) for this resource.

- 7 Assign drive letters to the disks on the cluster storage device.

Note: All the disks on the cluster storage device must be assigned drive letters. If a disk does not have a letter, it cannot be added to a clustered service (known as a resource group in Windows Server 2003)

intended for hosting Parallels Containers resources. For information on clustered services, see [Configuring the Cluster](#) (p. 31, p. 46).

For more information about configuring shared storage devices in cluster environments, refer to <http://technet2.microsoft.com/windowsserver/en/library/1bce08cc-2caf-417c-b53d-035f57c4e78b1033.msp?mfr=true>.

Setting Up a Cluster

Now you are ready to set up a failover cluster that will ensure the high availability of your Parallels Containers installations. The process of setting up this kind of cluster includes the following steps:

- 1 Creating a Cluster Service user account on the domain controller.

Note: Creating a Cluster service account is not necessary if your failover cluster will include only nodes running Windows Server 2008 and Windows Server 2008 R2.

- 2 Creating the failover cluster and configuring the first cluster node.
- 3 Configuring all the remaining nodes in the cluster.

These steps are described in the following subsections in detail.

Creating a Cluster Service Account

Before creating a failover cluster, you need to create a domain user account under which the Cluster service (this service is responsible for managing all cluster-related activities) will run. You will need this account later on when making the failover cluster and joining your nodes to it.

Creating a Cluster service account is not necessary if your failover cluster will include only nodes running Windows Server 2008 and Windows Server 2008 R2. So, the example below demonstrates how to create a Cluster service account in Windows Server 2003:

- 1 Log in to the domain controller.
- 2 Click **Start**, point to **Administrative Tools**, and select **Active Directory Users and Computers**.
- 3 In the left pane of the **Active Directory Users and Computers** window, expand the contents of the domain.
- 4 Right-click the **Users** folder, point to **New**, and select **User**.
- 5 In the displayed window, do the following:
 - Type arbitrary names in the **First name** and **Last name** fields (the choice of these names is not important for the account being created).
 - In the **User logon name** field, type a descriptive name that will be used to log in to the cluster (e.g., `ClusterService`).

New Object - User

Create in: mycompany.local/Users

First name: Jack Initials:

Last name: Brown

Full name: Jack Brown

User logon name: ClusterService @mycompany.local

User logon name (pre-Windows 2000): MYCOMPANY\ ClusterService

< Back Next > Cancel

6 Specify a password for the domain account, and click Next.

New Object - User

Create in: mygroup.local/Users

Password:

Confirm password:

User must change password at next logon

User cannot change password

Password never expires

Account is disabled

< Back Next > Cancel

- 7 The last screen allows you to review the parameters provided in the previous steps. To modify any settings, click **Back**; otherwise, click **Finish** to create the Cluster service account.

Deploying a Failover Cluster in Windows Server 2008 and Windows Server 2008 R2

This section describes the process of setting up a failover cluster on servers running Windows Server 2008 and Windows Server 2008 R2. For information on deploying a failover cluster on servers running Windows Server 2003, see *Deploying a Failover Cluster in Windows Server 2003* (p. 38).

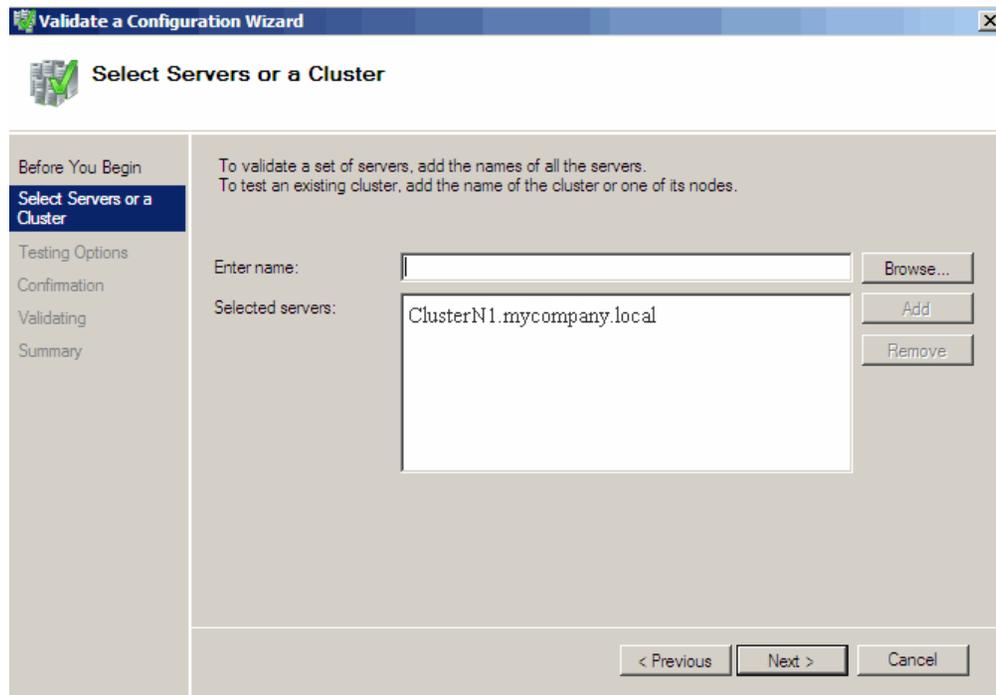
Validating Hardware

If you plan to create a failover cluster from servers running Windows Server 2008 and Windows Server 2008 R2, you first need to ensure that your servers hardware is supported for use with the Windows Server 2008 and Windows Server 2008 R2 clustering software. To do this:

- 1 Log in to any of your cluster nodes using the domain administrator account; you created this account in *Setting Up a Domain Administrator Account*.
- 2 Install the Failover Clustering feature on the node, if it is not yet installed.
- 3 Click **Start > Administrative Tools > Failover Cluster Manager** to open the Failover Cluster Manager snap-in.

Note: The Failover Cluster Manager snap-in is called Failover Cluster Management in Windows Server 2008 and can be accessed by clicking **Start > Administrative Tools > Failover Cluster Management**.

- 4 In the Management section, click **Validate a Configuration** to start the **Validate a Configuration** wizard.
- 5 In the **Before You Begin** window, read the provided information, and click **Next**.
- 6 In the **Select Servers or a Cluster** window, specify the server to be validated, and click **Next**. To save time, you can specify all your nodes you plan to include in the cluster.



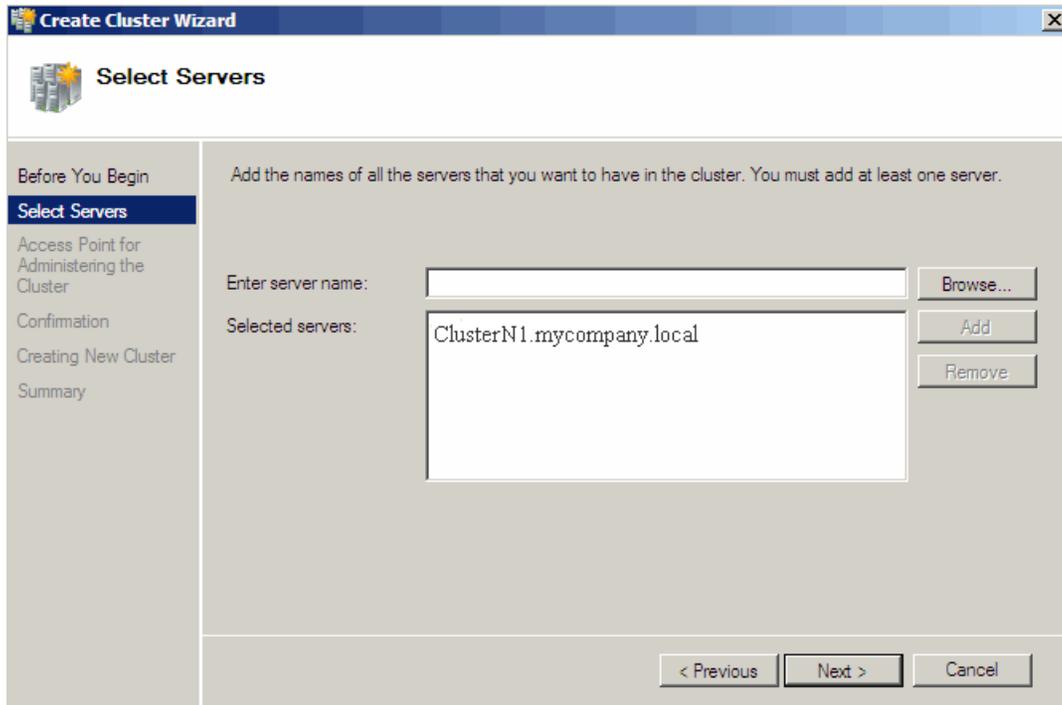
- 7 In the **Testing Options** window, click **Next**.
- 8 In the **Confirmation** window, review the servers to be validated and the tests to be run during the validation, and click **Next**.
You can view the progress in the **Validating** window.
- 9 After the validation is complete, review the validation results in the **Summary** window, and click **Finish** to close the wizard.

For more information on validating hardware for a failover cluster, see [http://technet.microsoft.com/en-us/library/cc732035\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc732035(WS.10).aspx).

Creating a Cluster in Windows Server 2008 and Windows Server 2008 R2

The processes of creating a failover cluster and configuring the first cluster node take place simultaneously and are described below:

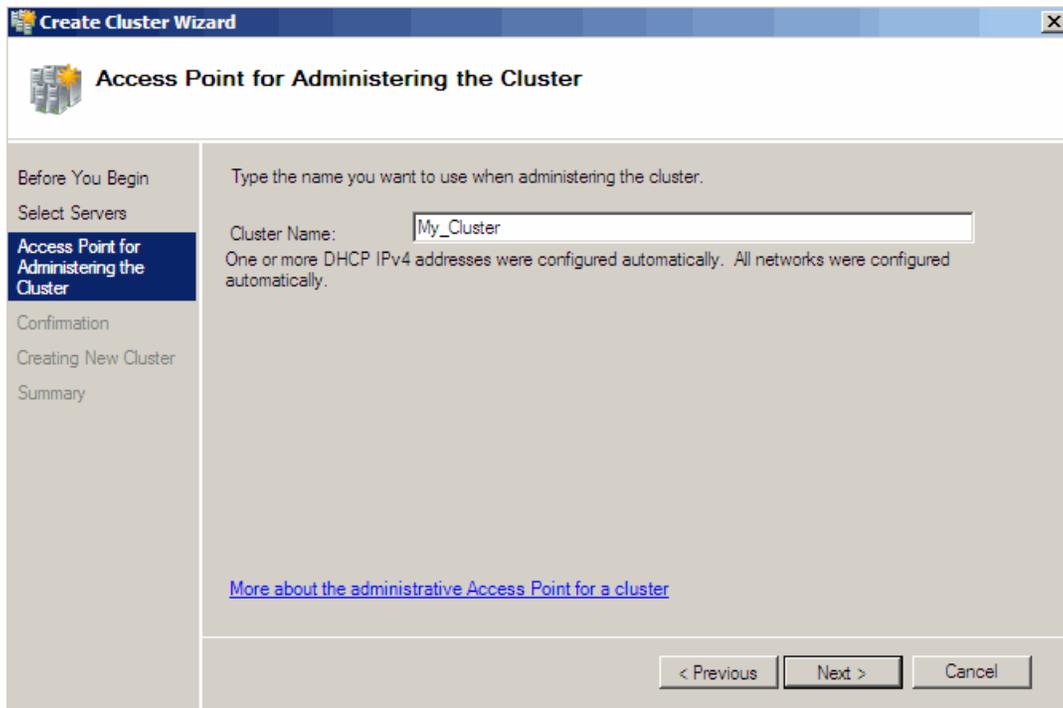
- 1 Log in to any of your cluster nodes using the domain administrator account (you created this account in **Setting Up a Domain Administrator Account**). This node will be registered as the first node in the cluster.
- 2 Click **Start > Administrative Tools > Failover Cluster Management** to open the Failover Cluster Management snap-in.
- 3 In the center pane under **Management**, click **Create a Cluster** to launch the **Create Cluster** wizard.
- 4 In the **Before You Begin** window, click **Next** to start validating your server hardware. This step is skipped if you have already validated the hardware.
- 5 In the **Select Servers** window, specify the node hostname or IP address, and click **Add**.



In this step, you can specify all nodes you want to add to the cluster. For the purposes of this guide, however, we add only one node in this step and will add all the other nodes after creating the cluster.

When you are ready, click **Next**.

- 6 Next, you are asked to assign a name to the cluster. Type a descriptive name for the cluster in the **Cluster Name** field (e.g., `My_Cluster`), and click **Next**. As the cluster administrator, you will then use this name to connect to the cluster with cluster management applications (e.g., Cluster Administrator or `cluster.exe`) and manage the cluster in different ways (e.g., configure some of its settings or monitor its activity).



- 7 In the Confirmation window, review the parameters that will be used for creating the cluster, and click Next.

You can view the progress in the Creating New Cluster window.

- 8 After the cluster has been successfully created, the Summary window appears. Click Finish to close the wizard.

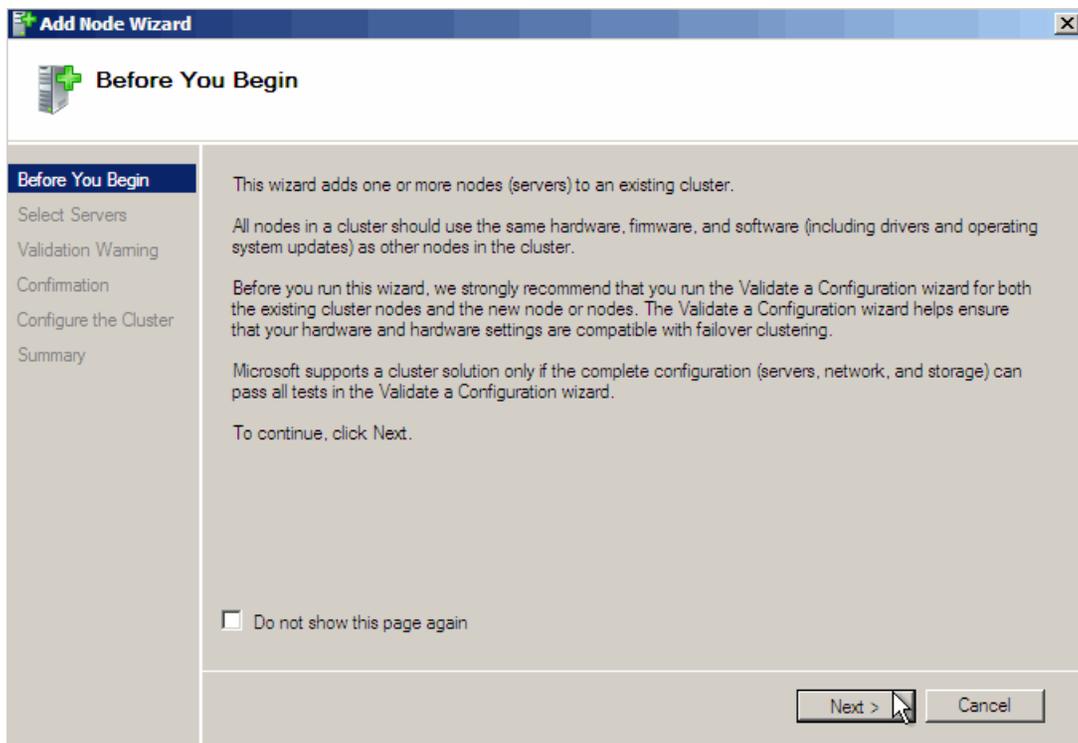
The cluster name appears in the left pane of the Failover Cluster Management snap-in.



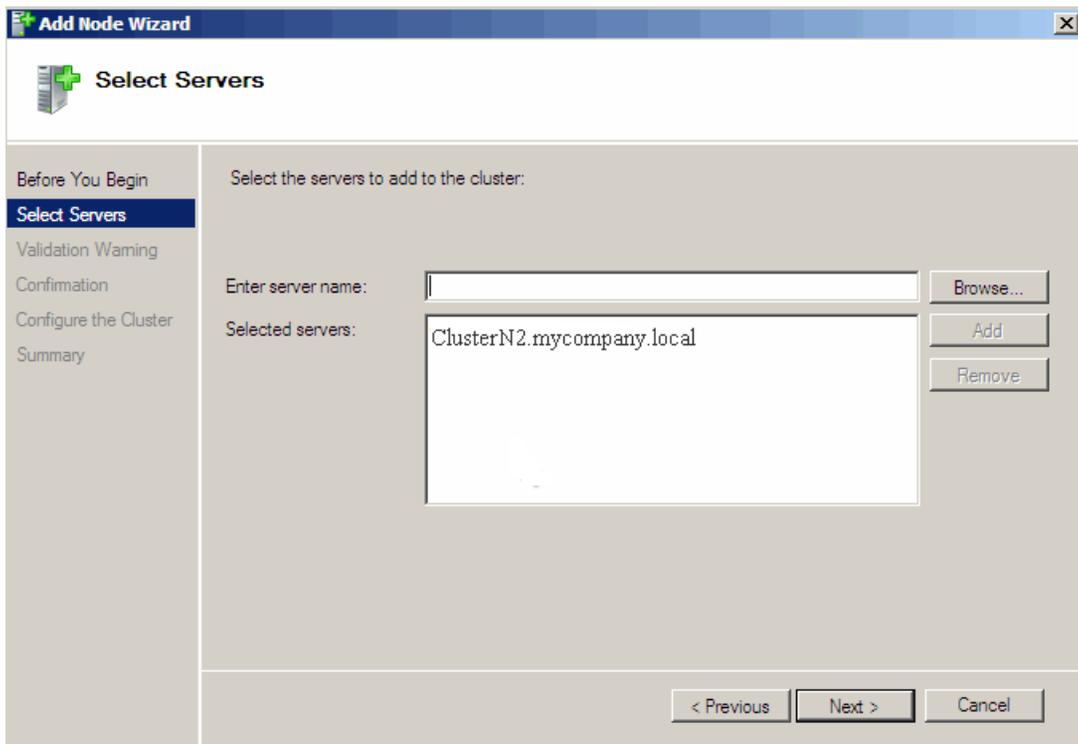
Adding Nodes to the Cluster

The process of adding the second and all the remaining nodes to the failover cluster is almost identical to that of setting up the first cluster node and is described below:

- 1 Log in as the domain administrator to the cluster node that you registered in the previous step.
- 2 Click **Start > Administrative Tools > Failover Cluster Management** to open the Failover Cluster Management snap-in.
- 3 In the left pane, right-click the cluster name, and select **Add Node**.
- 4 In the **Before You Begin** window, click **Next** to proceed with the wizard.

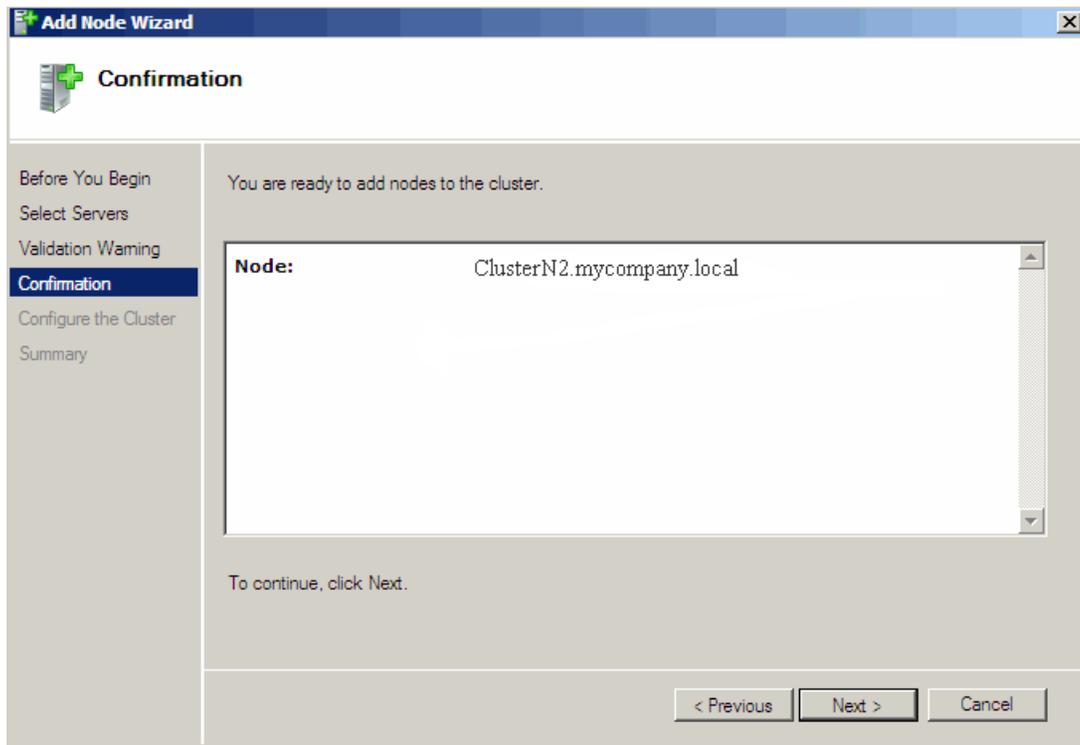


- 5 In the **Select Servers** window, specify the nodes to be added to the cluster.



To add a node, type its hostname or IP address in the **Enter server name** field, and click **Add**. After specifying all the necessary nodes, click **Next**.

- 6** In the **Validation Warning** window, read the information provided, and click **Next** to start validating your server hardware.
- 7** In the **Confirmation** window, review the nodes that will be added to the cluster, and click **Next**.



The progress of adding nodes to the cluster is shown in the **Configure the Cluster** window.

- 8 Once your nodes have been successfully added to the cluster, the **Summary** window appears. Click **Finish** to close the wizard.

Configuring the Cluster

After you have created the failover cluster and registered all nodes with it, you need to configure some of its settings to prepare the cluster for installing the Parallels Containers software. You need to do the following:

- 1 Create clustered services (called resource groups in Windows Server 2003).
- 2 Add disk resources to services.
- 3 Configure the quorum mode.

Detailed information on these configuration steps is given below.

Creating Clustered Services

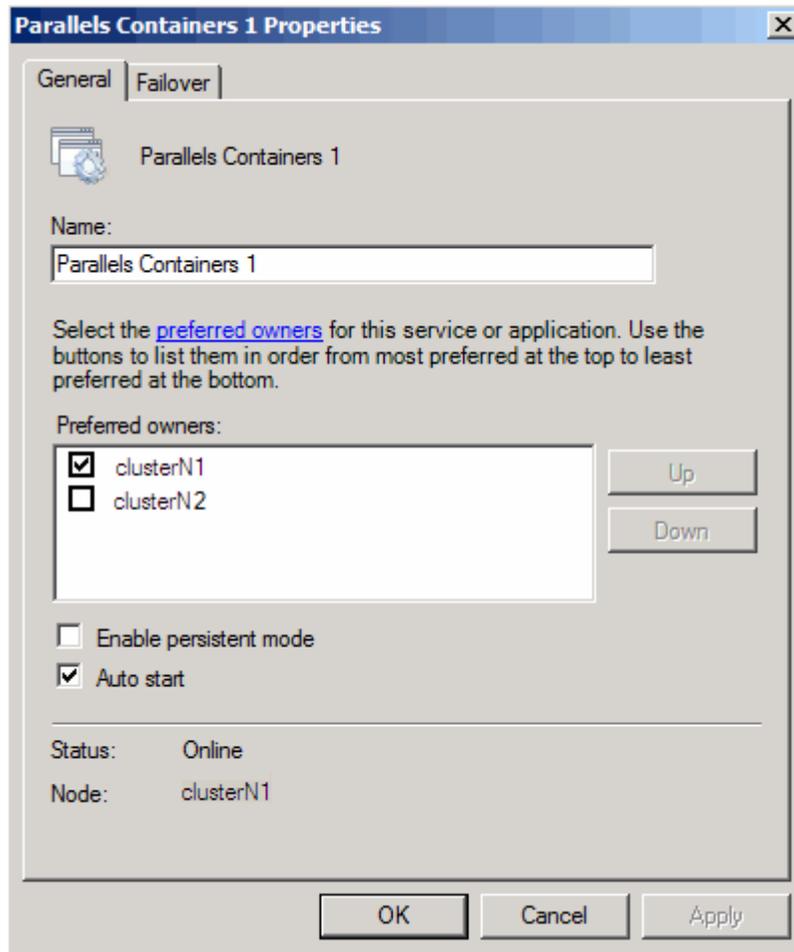
In the first step, you need to create clustered services. These services will contain all Parallels Containers resources and fail over to another node in the case of a failure. You must create a separate service for each active node in the cluster.

To create a clustered service:

- 1 In the left pane of the Failover Cluster Manager snap-in, right-click the **Service and Applications** item, and select **More Actions > Create Empty Service or Application**.

After a while, the new service appears in the **Services and Applications** pane.

- 2 Right-click the service, and choose **Properties**.
- 3 Type a descriptive name for the service in the **Name** field, and click **OK**.



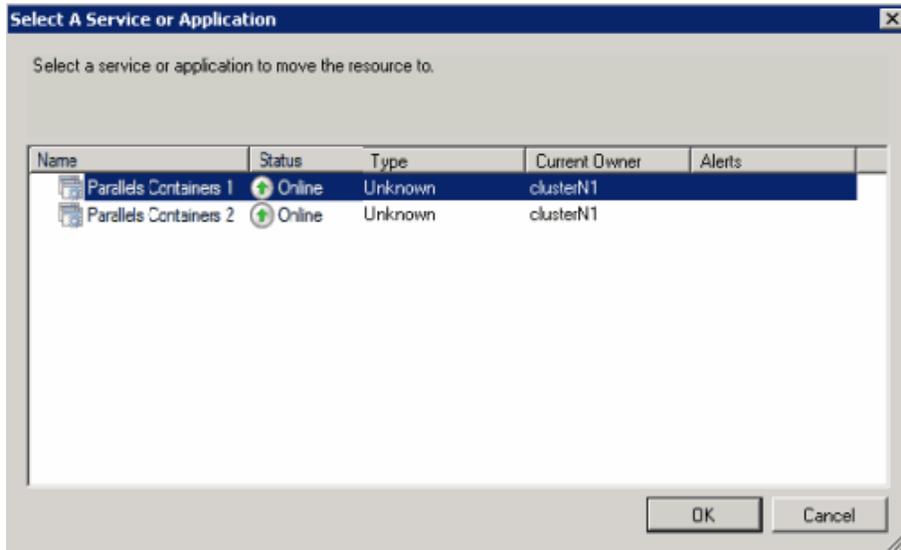
Adding Disk Resources to Resource Groups

Next, you need to add disk resources to the clustered services you created in the previous step. Each clustered service intended for hosting Parallels Containers resources must contain at least one disk resource. When you create a failover cluster, the clustering software automatically processes all disks on your shared storage and adds them to the failover cluster. You can view these disks by clicking the **Storage** item in the Failover Cluster Manager snap-in.

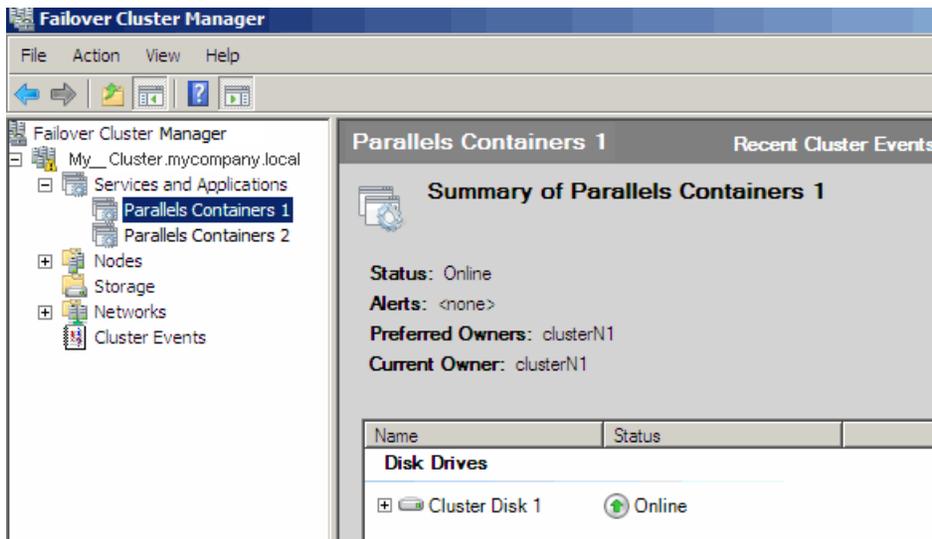
To add a disk resource to a service, do the following:

- 1 In the Failover Cluster Manager snap-in, double-click the cluster name, and click **Storage**.
- 2 Right-click the disk resource, and select **More Actions > Move this resource to another service or application**.

- 3 In the **Select a Service or Application** window, select the service you want to add the disk resource to, and click **OK**.



To check that the disk resource has been successfully added to the service, expand the **Services and Applications** item, and click the service name. The added disk should be listed under **Disk Drives**, and its status should be set to **Online**.



When working with disk resources, keep in mind the following:

- Each service intended for hosting Parallels Containers resources must contain at least one disk resource.
- A separate disk is also needed for hosting the *quorum* resource if you decide to store the quorum on a specific disk. For details, see **Configuring the Quorum Mode** below. In this case, you can calculate the number of disks using this formula: $disks_number = groups_number + 1$. For example, if you have 3 services, you need to create at least 4 disk resources.

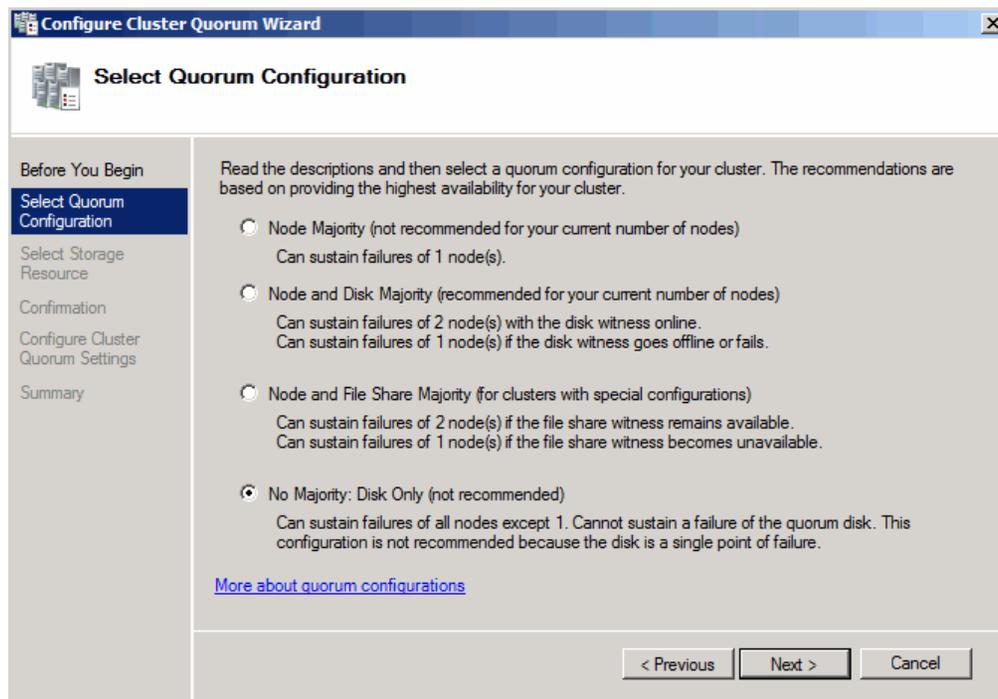
- If the number of clustered services exceeds the number of disk resources you currently have or you need a disk for hosting the quorum resource, you can add additional disk resources to the cluster. For information on how you can do this, see the **Adding Disks Resources to the Cluster** section below.

Configuring the Quorum Mode

In the last step, you need to choose the quorum mode for use in your cluster. Follow this link to find out what quorum mode is recommended for your failover cluster implementation: [http://technet.microsoft.com/en-us/library/dd197496\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd197496(WS.10).aspx).

After deciding on the right quorum mode, you can enable it by doing the following:

- 1 In the left pane of the Failover Cluster Manager snap-in, click the name of the cluster.
- 2 In the Actions pane, select **More Actions > Configure Cluster Quorum Settings** to launch the **Configure Cluster Quorum** wizard.
- 3 In the **Before You Begin** window, read the provided information, and click **Next**.
- 4 In the **Select Quorum Configuration** window, choose the quorum configuration that best meets your needs, and click **Next**.



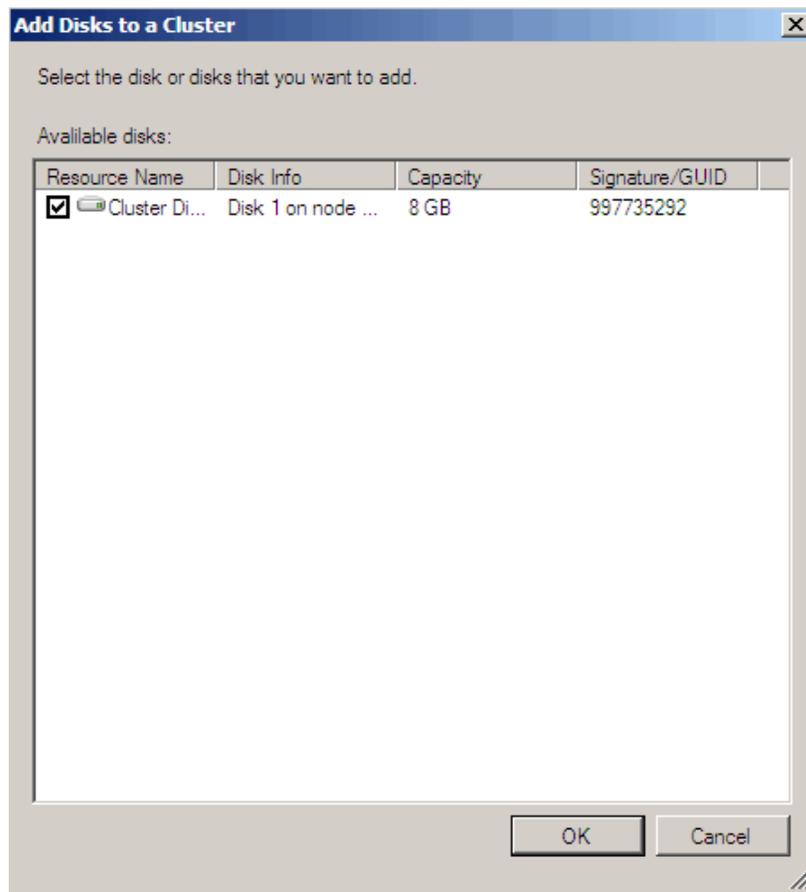
- 5 Follow the instructions of the wizard to complete the quorum configuration.

Adding Disk Resources to the Cluster

Sometimes, you may need to add an additional disk resource to the failover cluster. For example, this may be the case if the number of clustered services exceeds the number of disk resources you currently have or if you need a disk for hosting the quorum resource.

To add a disk resource to the cluster, do the following:

- 1 Make sure your shared storage has at least one free disk configured for including in the cluster. For more information on configuring shared storages and their disks, see [Connecting Servers to a Shared Storage](#) (p. 22).
- 2 In the Failover Cluster Manager snap-in, double-click the cluster name.
- 3 Right-click the **Storage** item, and select **Add a disk**.
- 4 In the **Add Disks to a Cluster** window, select the check boxes next to the disk you want to add to the cluster, and click **OK**.



After a while, the disk should appear in the center pane of the Failover Cluster Manager snap-in.

Checking Cluster Operation

You can use the procedures below to verify the correct operation of your failover cluster.

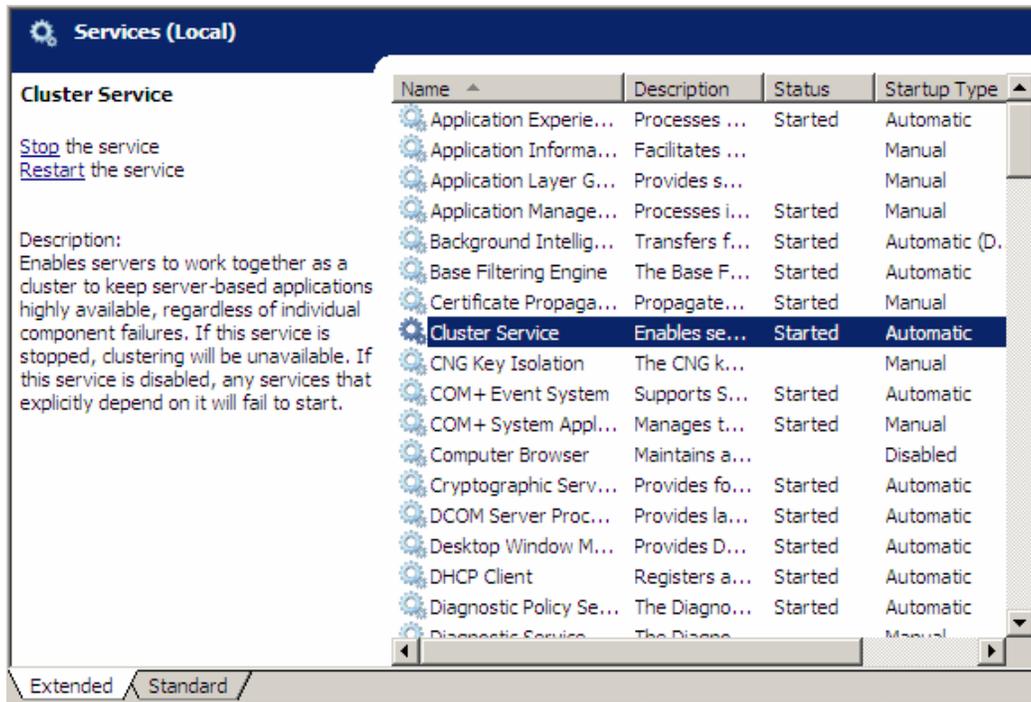
First Procedure (viewing the status of cluster nodes in Failover Cluster Manager):

- 1 Log in to any of your cluster nodes using the domain administrator account.
- 2 Click **Start > Administrative Tools > Failover Cluster Manager** to open the Failover Cluster Manager snap-in.

3 Verify that all nodes are listed under the corresponding cluster and their resources are online.

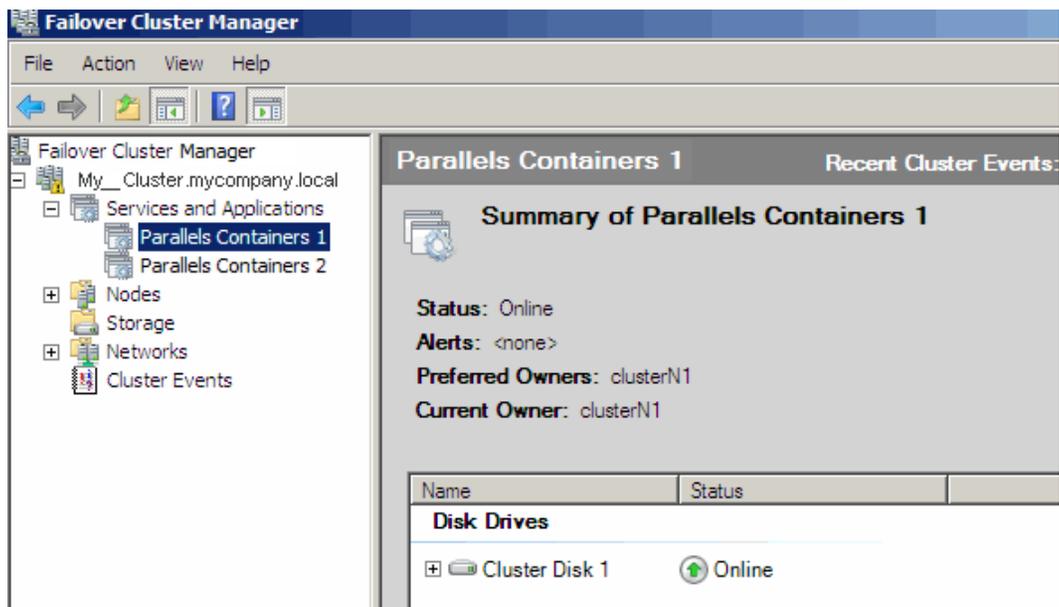
Second Procedure (checking the Cluster service availability):

- 1 On each cluster node, click **Start > Administrative Tools > Services** to launch the Services snap-in.
- 2 In the right pane of the displayed window, verify that the **Cluster Service** service is listed in the Services table and its status is set to **Started**.



Third Procedure (checking the failover of cluster resources):

- 1 Log in to any of your cluster nodes using the domain administrator account.
- 2 Click **Start > Administrative Tools > Failover Cluster Manager** to open the Failover Cluster Manager snap-in.
- 3 Expand the **Services and Applications** item, and click the name of a clustered service. The center pane will display all the resources currently included in the service.

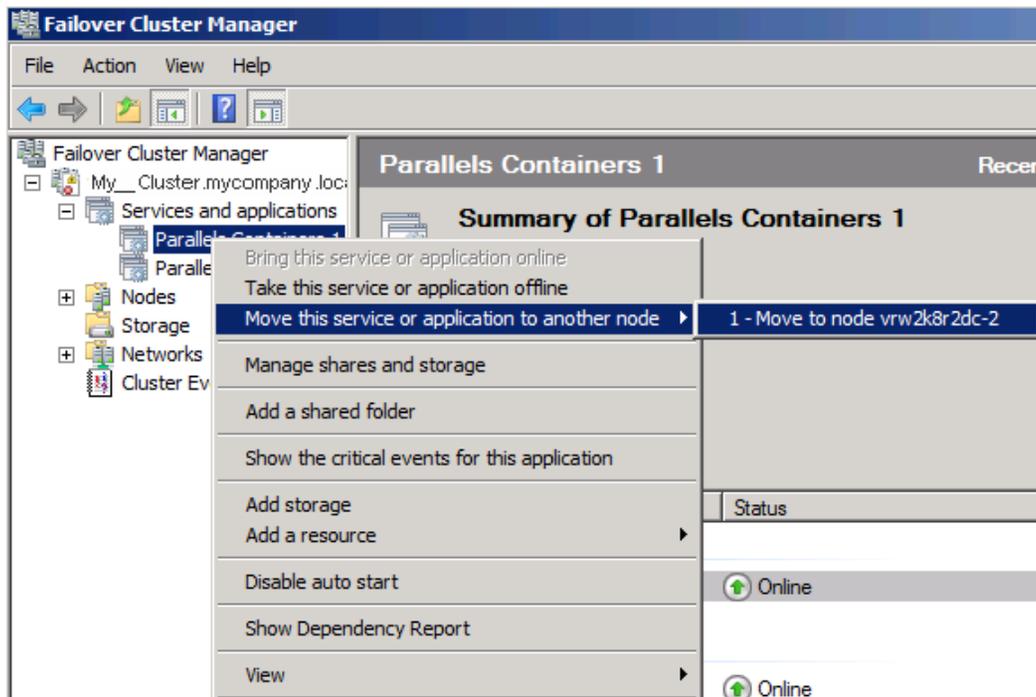


- 4 Choose a resource to use for simulating the failover, right-click it, and select **More Actions > Simulate failure of this resource**.

The resource should come back online either on the same or on another node in the cluster.

Fourth Procedure (checking the failover of clustered services):

- 1 Log in to any of your cluster nodes using the domain administrator account.
- 2 Click **Start > Administrative Tools > Failover Cluster Manager** to open the Failover Cluster Manager snap-in.
- 3 Expand the **Services and Applications** item, right-click a clustered service, and select **Move this service or application to another node > Move to node *Node_Name***.



Make sure that the clustered service and all its resources have come back online on the selected node.

- 4 Repeat Steps 1-3 for each clustered service.

Deploying a Failover Cluster in Windows Server 2003

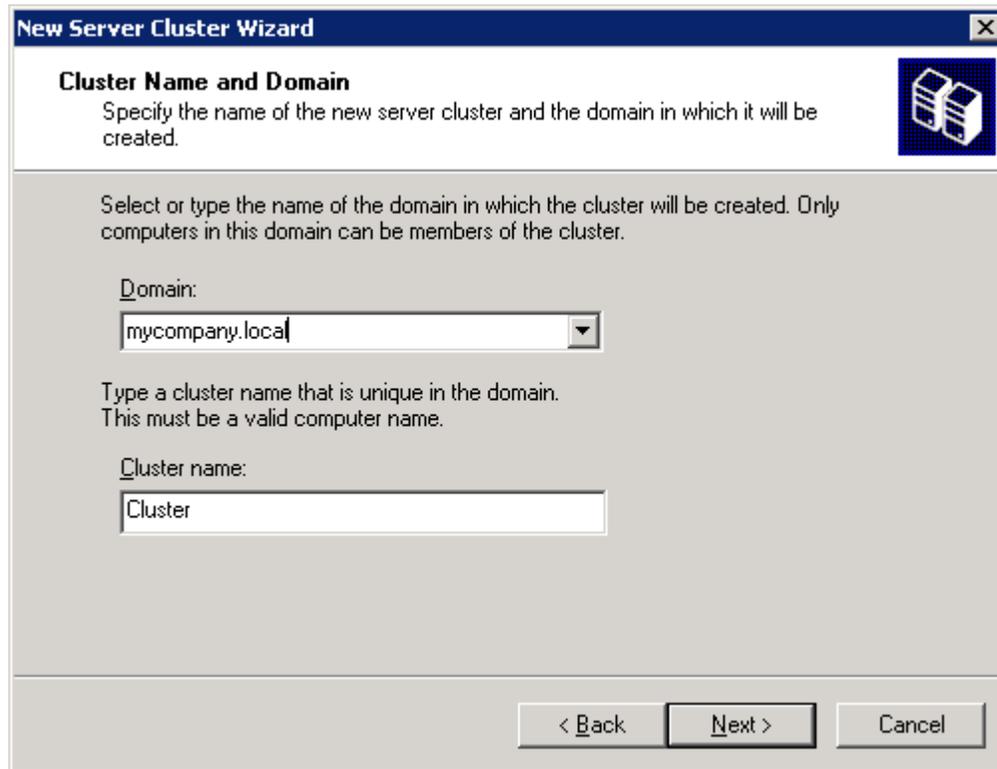
This section describes the process of setting up a server cluster on servers running Windows Server 2003. For information on deploying a failover cluster on servers running Windows Server 2008 and Windows Server 2008 R2, refer to *Deploying a Failover Cluster in Windows Server 2008 and Windows Server 2008 R2* (p. 25).

Creating the Cluster and Configuring the First Cluster Node

The processes of creating a failover cluster and configuring the first cluster node take place simultaneously and can be performed by doing the following:

- 1 Log in to any of your cluster nodes using the domain administrator account (you created this account in *Setting Up a Domain Administrator Account*). This node will be registered as the first node in the cluster.
- 2 Click **Start > Run**, and execute the `cluadmin` command in the **Run** dialog to launch the **Cluster Administrator** application.
- 3 In the **Open Connection to Cluster** dialog box, select the **Create new cluster** item on the **Action** menu, and click **OK** to launch the **New Server Cluster** wizard.
- 4 In the **Welcome** window, click **Next**.
- 5 In the displayed window, provide the following information:

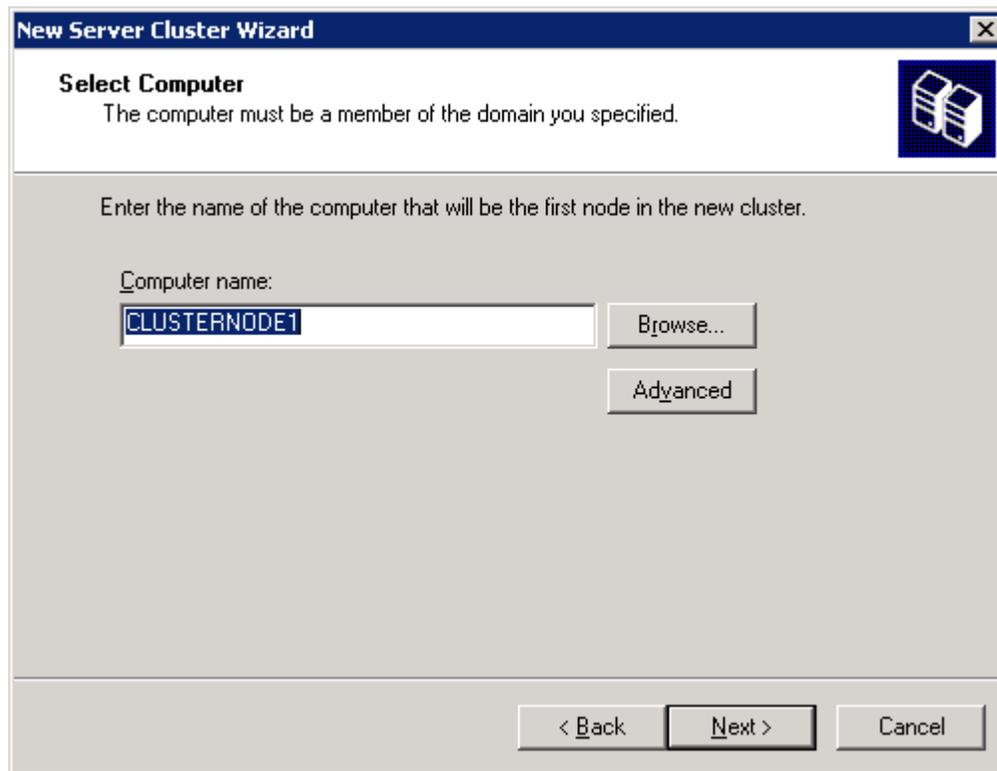
- In the **Domain** field, specify the name of the domain where the cluster will be located. You specified this name during the Active Directory domain creation (in our example it must be set to `mycompany.local`).
- In the **Cluster name** field, indicate an arbitrary name (e.g., `Cluster`) that will uniquely identify the cluster in the domain.



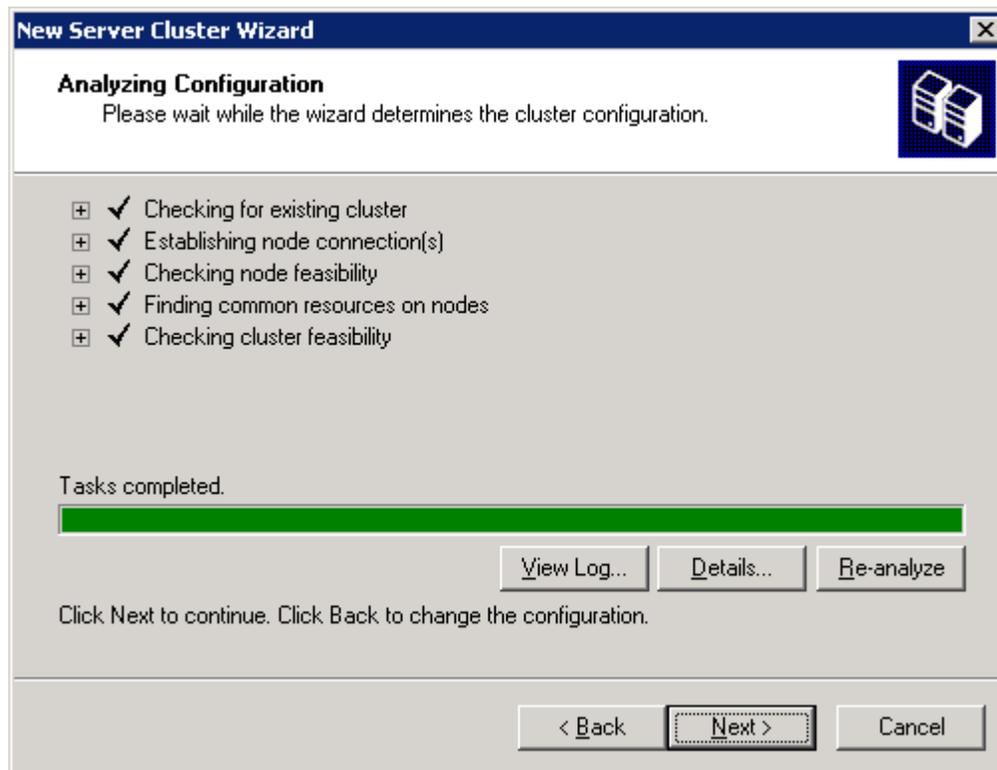
The screenshot shows a Windows-style dialog box titled "New Server Cluster Wizard". The main heading is "Cluster Name and Domain" with a sub-instruction: "Specify the name of the new server cluster and the domain in which it will be created." Below this, there is a text box for "Domain:" containing "mycompany.local". A second instruction reads: "Select or type the name of the domain in which the cluster will be created. Only computers in this domain can be members of the cluster." Below that is a text box for "Cluster name:" containing "Cluster". At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

When you are ready, click **Next**.

- 6 In the **Select Computer** window, provide the hostname of the server that will be registered as the first node in the cluster. By default, the wizard displays the hostname of the server you are currently logged in to. Click **Next** to accept the default setting and start analyzing your cluster configuration.

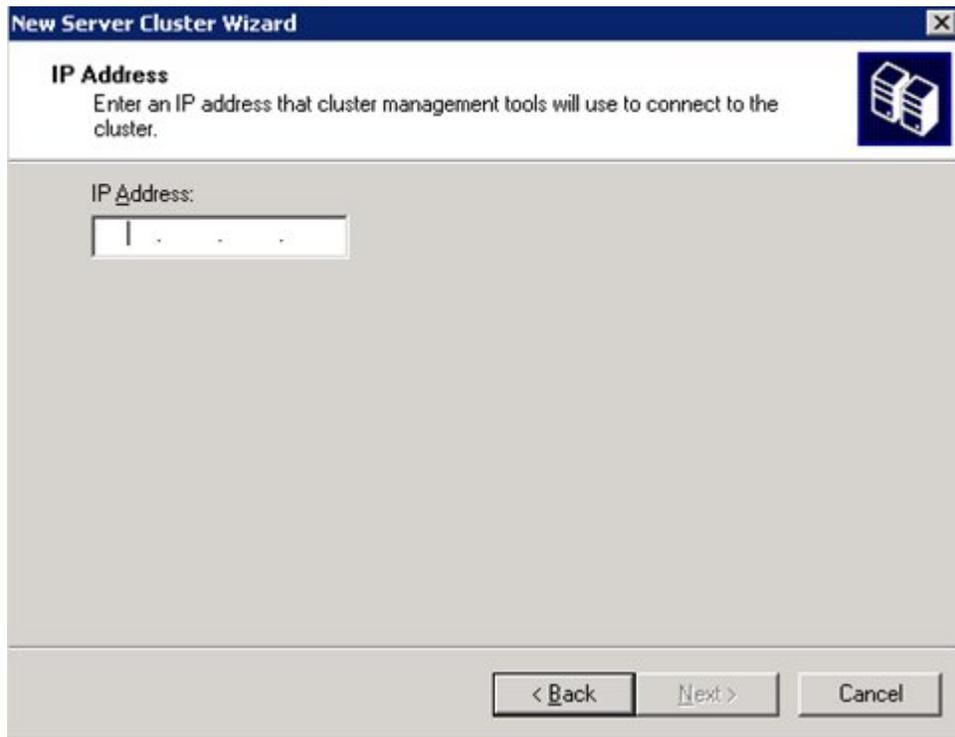


- 7 If the configuration is correct, the progress bar turns green and the Next button becomes available. Click Next to continue with your cluster configuration.



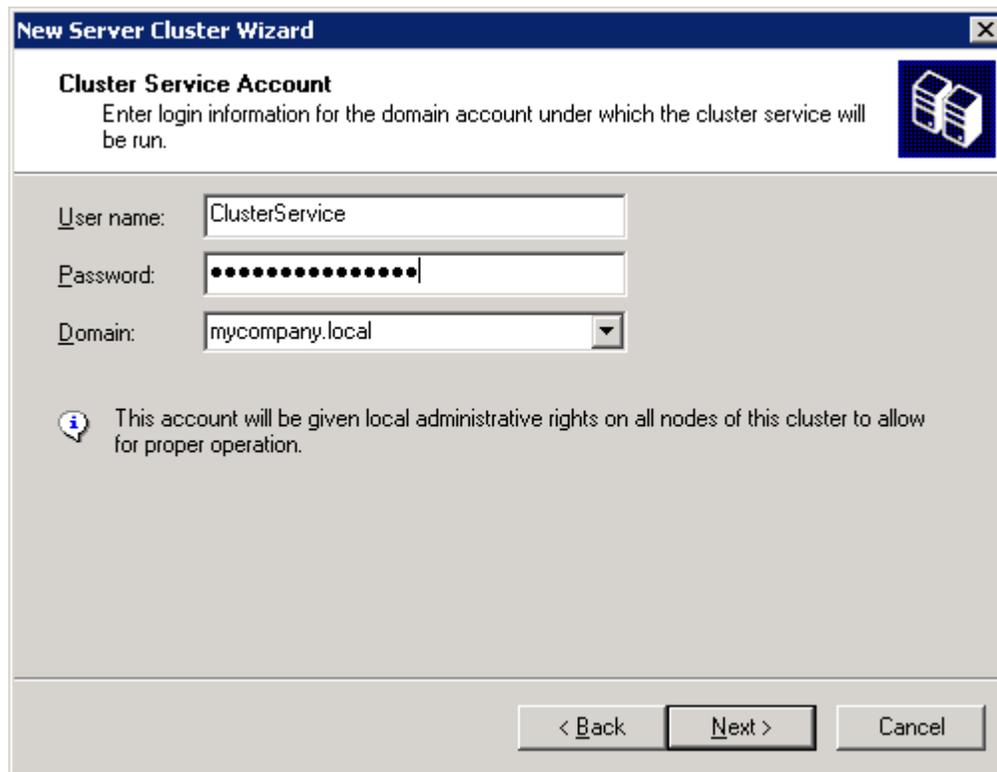
- 8** In the **IP Address** window, indicate an IP address for the cluster. The cluster administrator will use this IP address to connect to the cluster using cluster management applications (e.g. Cluster Administrator or `cluster.exe`) and to manage the cluster in different ways (e.g. configure some of its settings or monitor its activity). When choosing an IP address for the cluster, make sure that it is static, unique, and not assigned to any other computer on the network, including the cluster nodes themselves.

Note: The IP address you specify in this step will be associated with the cluster name you indicated in the **Cluster Name** field of the **Cluster Name and Domain** window (see **Step 5**). That means you can use either the IP address or the cluster name to connect to the cluster.



Click Next to continue.

- 9 In the **Cluster Service Account** window, enter the user name and password of the Cluster service account created by you in the previous subsection.



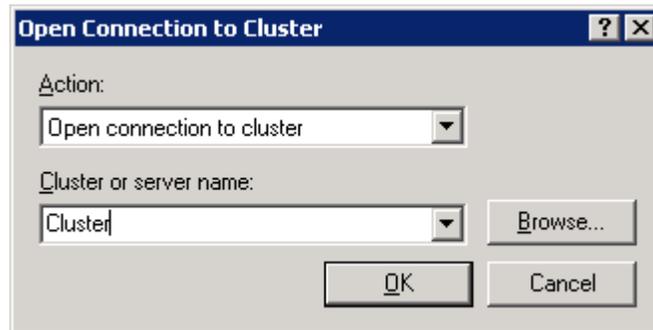
After providing the necessary information, click **Next** to log the node into the domain using the specified credentials.

- 10** In the **Proposed Cluster Configuration** window, review the configuration parameters provided by you in the previous steps of the wizard. If you are satisfied with the specified parameters, click **Next** to create the cluster and join the first node to it. Otherwise, click **Back** and modify the necessary parameters.
- 11** Click **Finish** to exit the wizard and open the **Cluster Administrator** application with the first cluster node registered.

Adding Nodes to the Cluster

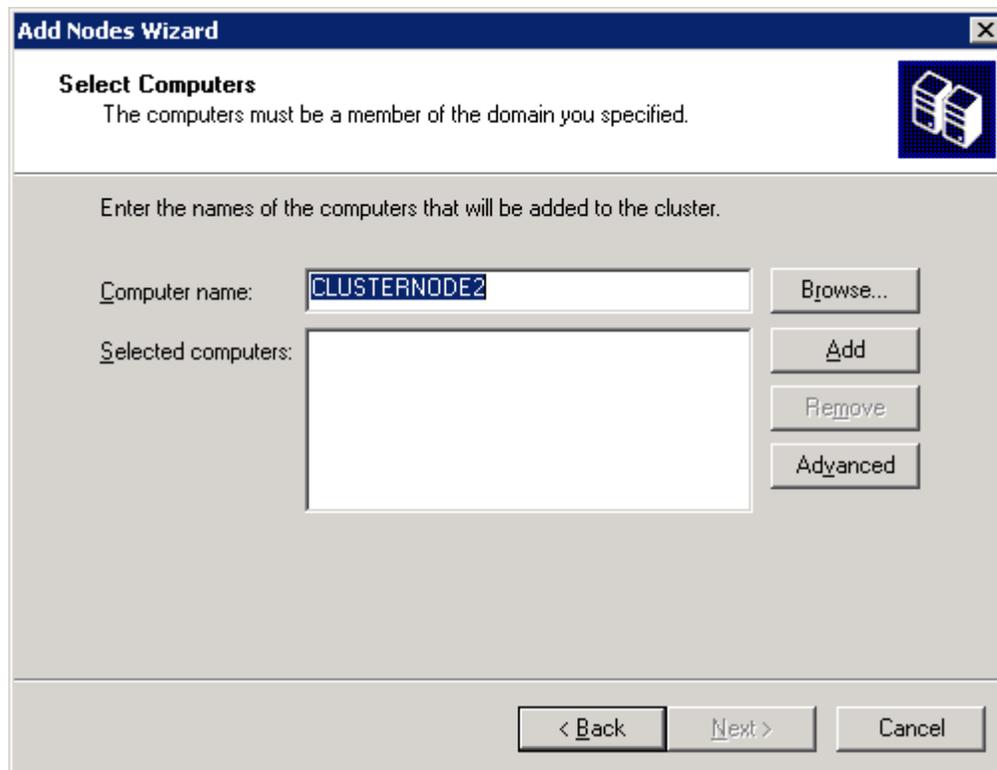
The process of adding the second and all the remaining nodes to the server cluster is almost identical to that of setting up the first cluster node and is described below:

- 1** Log in to the node you wish to register in the cluster and launch the **Cluster Administrator** application by clicking **Start > Run** and executing the `cluadmin` command in the **Run** dialog.
- 2** In the **Open Connection to Cluster** window, do the following:
 - Under **Action**, select the **Open connection to cluster** item on the **Action** menu.
 - Under **Cluster or server name**, specify the name of the cluster you created in the previous subsection.



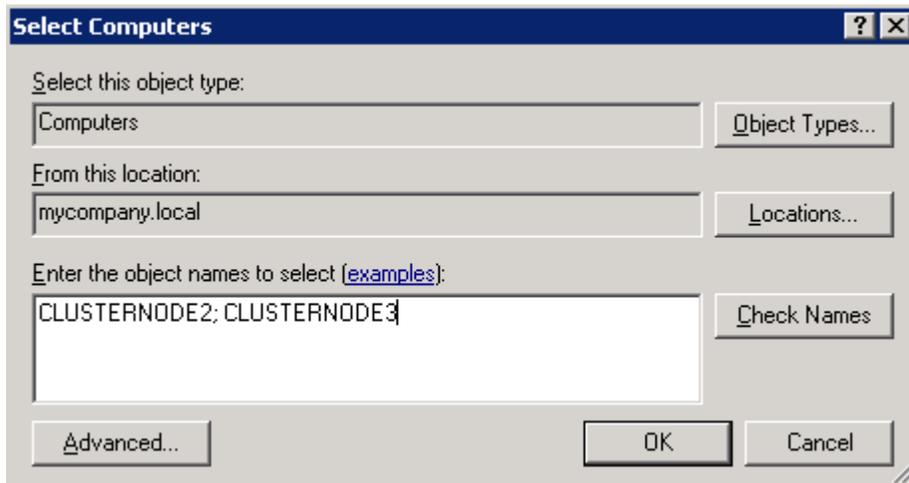
Click OK to connect to the cluster and display its configuration in Cluster Administrator.

- 3 Click File, point to New, and select Node to launch the Add Nodes wizard.
- 4 In the Welcome to ... window, click Next to proceed with the wizard.
- 5 In the Select Computers window, select the nodes to be added to the cluster.

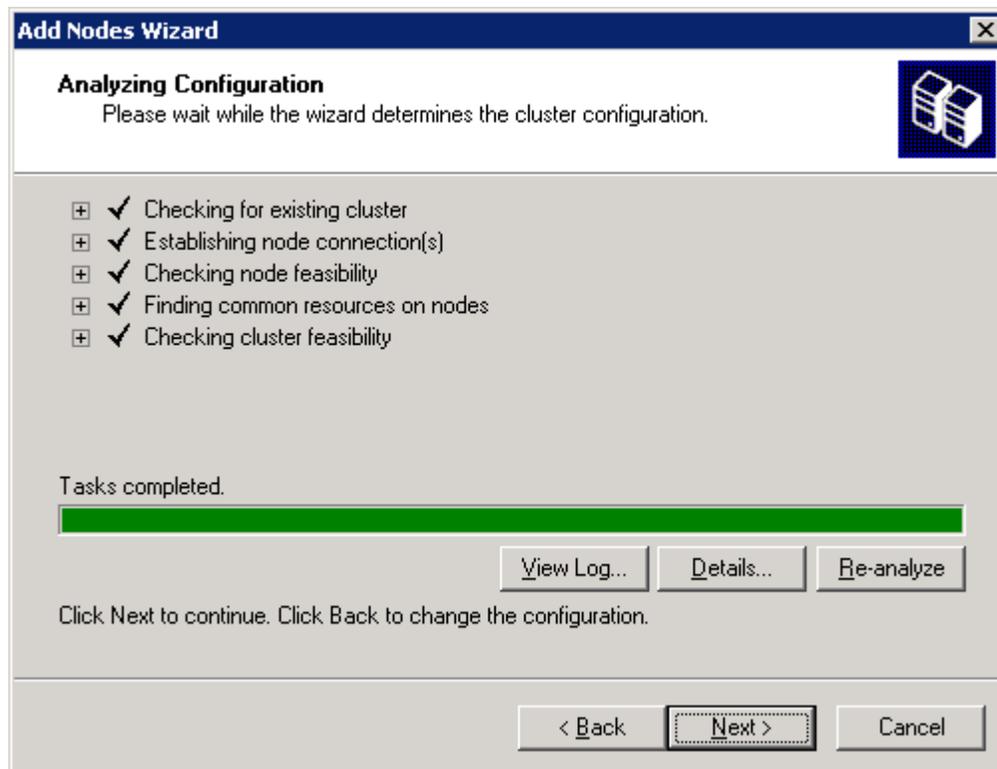


By default, the wizard offers to include in the cluster only the node you are currently logged in to. The hostname of this node is displayed in the **Computer name** field (in our case this node is assigned the hostname of CLUSTERNODE2). If you wish to add only the given node to the cluster, click **Add**, and then click **Next**. If you wish to join any other nodes to the cluster, do the following:

- Click **Browse** to display the **Select Computers** window.

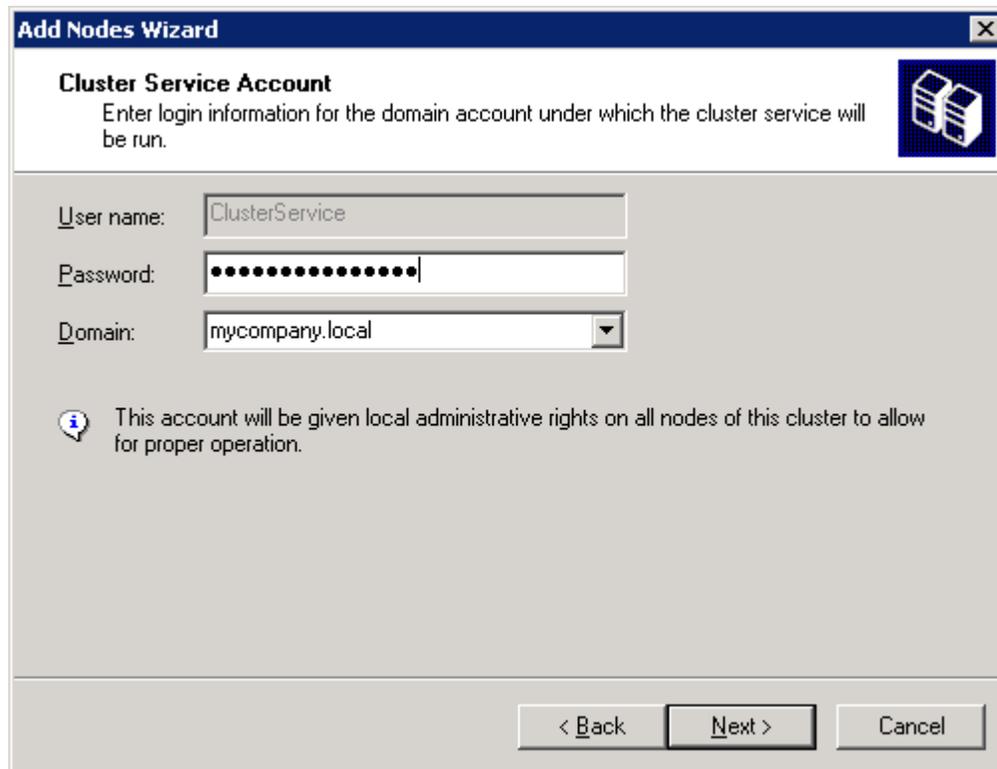


- In the Enter the object names to select field, type the hostnames of those nodes that you wish to add to the cluster and click OK. All the selected nodes will be displayed in the Computer name field of the Select Computers window.
 - Click Add, and then click Next.
- 6 The wizard starts checking the configuration of the selected nodes and the cluster itself. If everything is OK, the progress bar turns green and the Next button becomes available.



Click Next to continue.

- 7 In the Cluster Service Account window, type the password of the Cluster service account in the Password field, and click Next to log the node into the domain.



The screenshot shows the 'Add Nodes Wizard' dialog box with the 'Cluster Service Account' step selected. The title bar reads 'Add Nodes Wizard'. The main heading is 'Cluster Service Account' with a sub-instruction: 'Enter login information for the domain account under which the cluster service will be run.' There is an icon of two server racks on the right. The form contains three input fields: 'User name:' with the text 'ClusterService', 'Password:' with a masked password of 12 dots, and 'Domain:' with a dropdown menu showing 'mycompany.local'. Below the fields is an information icon and a message: 'This account will be given local administrative rights on all nodes of this cluster to allow for proper operation.' At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

- 8 In the Proposed Cluster Configuration window, review the parameters provided by you in the previous steps of the wizard. If you wish to modify any parameters, click Back; otherwise, click Next to start adding your nodes to the cluster. The process is displayed in the progress bar of the Adding Nodes to the Cluster window.
- 9 After the nodes have been successfully registered in the cluster, click Finish to exit the wizard and open the Cluster Administrator main window with all the nodes registered in the cluster.

Configuring the Cluster

Sometimes, after you have created the failover cluster and registered all nodes with it, you may need to configure some of its settings. Usually, such a configuration involves performing the following operations:

- Creating cluster resource groups (clustered services in Windows Server 2008 and Windows Server 2008 R2) that will store Parallels Containers resources. For example, creating a resource group may be necessary in the following cases:
 - You did not have enough shared disks when creating the cluster. As a rule, the clustering software automatically creates a resource group for each shared disk in the cluster. However, if the number of shared disks during the cluster creation was less than the number of Parallels Containers groups you plan to have, you need to manually create a resource group for each of the shared disks you will add to the cluster.

- You decided to include an additional node with Parallels Containers after creating the cluster. If you make up your mind to add one or more nodes running Parallels Containers to your cluster in the future, you need to create a resource group for each of the added nodes.
- Adding the necessary resources to existing resource groups. Each resource group must contain the following resources to be able to host Parallels Containers resources:
 - *Disk*. The Disk resource must be added to each resource group before starting the Parallels Containers installation.
 - *IP Address* and *Network Name*. The IP Address and Network Name resources can be added to groups either when configuring the cluster or when installing Parallels Containers on cluster nodes. However, we highly recommend that you add both resources to groups before starting the Parallels Containers installation. Doing so will ensure that the parameters of the IP Address and Network Name resources are properly configured and eliminate any errors that may be caused by providing incorrect parameters for these resource during the Parallels Containers installation.

The following subsections focus on all those tasks you may need to perform to configure your cluster. These tasks are listed below in the order of their execution:

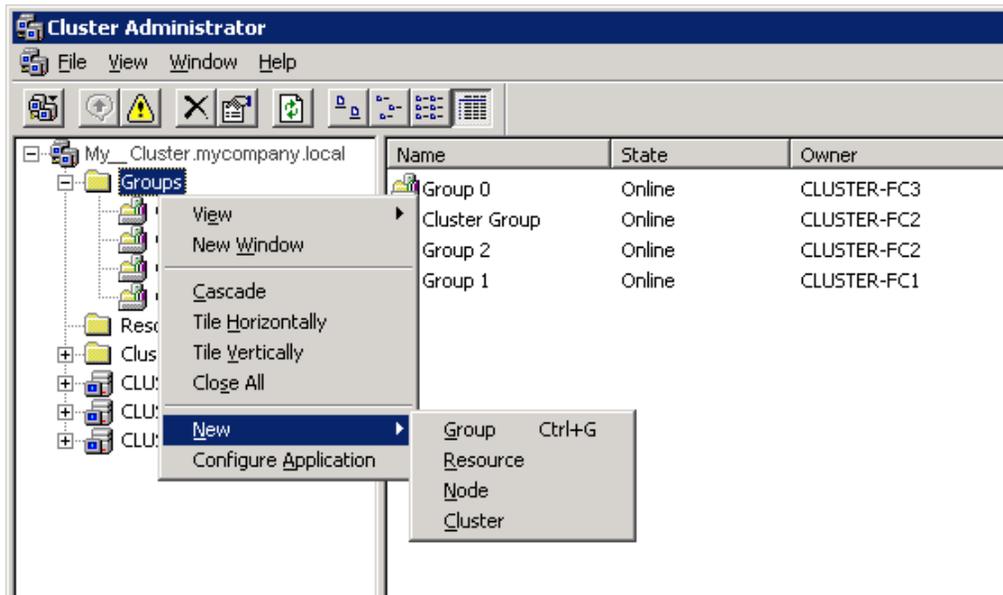
- 1 Creating a resource group.
- 2 Adding a Disk resource to a resource group.
- 3 Adding an IP Address resource to a resource group.
- 4 Adding a Network Name resource to a resource group.

Creating Cluster Resource Groups

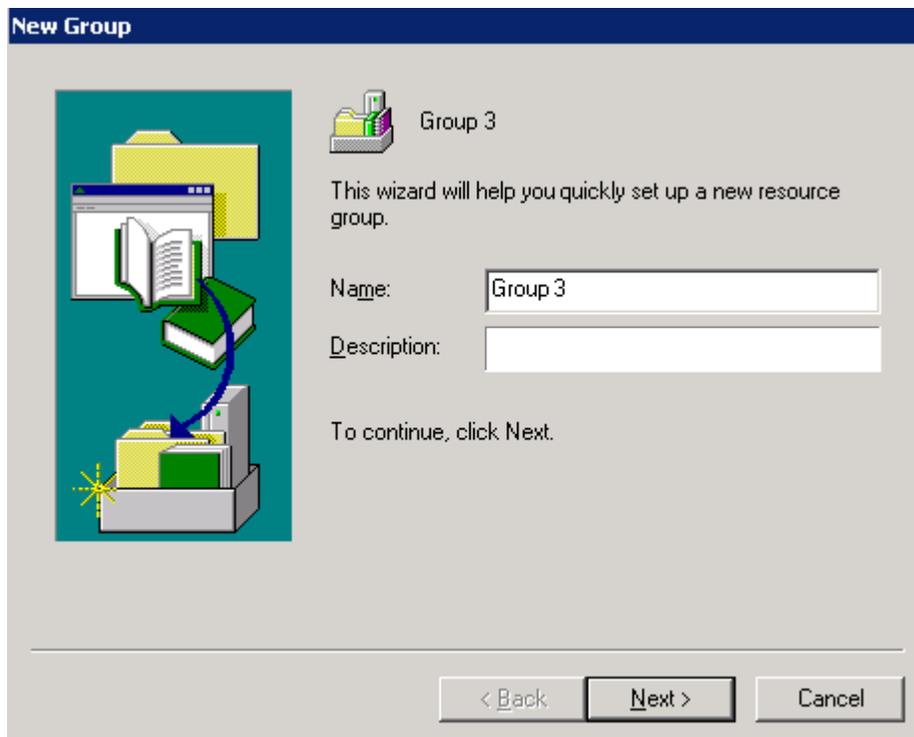
Each active node in the failover cluster must have a separate resource group. So, if your cluster has 4 active nodes, you must have at least 4 resource groups. Each resource group will contain Parallels Containers resources and fail over to another node in the case of a failure.

To create a resource group:

- 1 In the left pane of the Cluster Administrator snap-in, right-click the **Groups** item, and select **New > Group**.



- 2 In the New Group wizard, type a descriptive name for the group in the Name field (e.g. Group 3), and click Next.

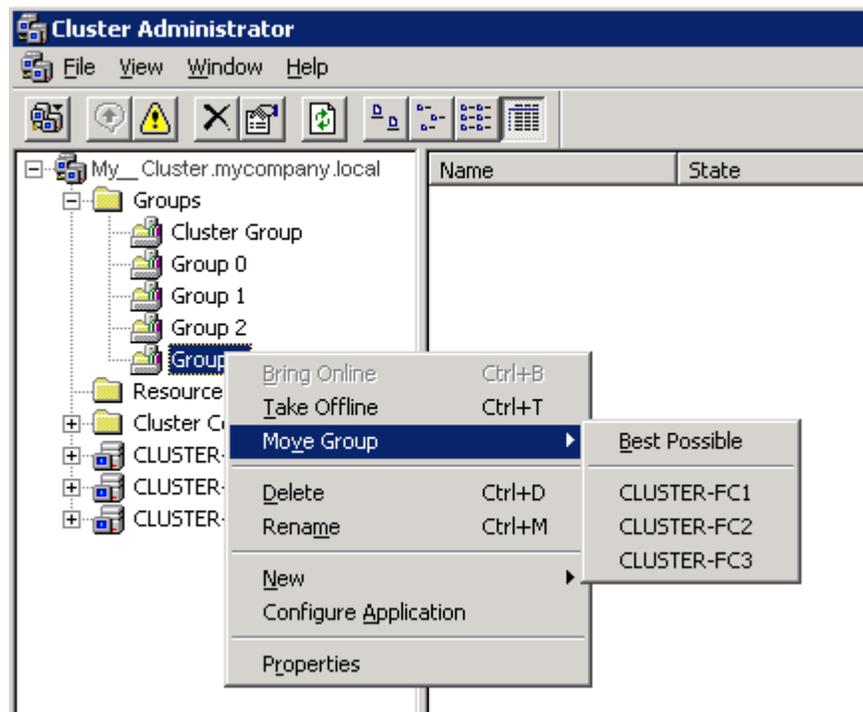


- 3 In the Preferred Owners window, click Finish.
- 4 The dialog appears informing you that the resource group has been successfully created. Click OK to close the dialog.

The new created resource group appears in the left pane of Cluster Administrator under the **Group** item. To activate the group, right-click it, and select **Bring Online**. After a while, the group status should change from **Offline** to **Online**.

Moving a Group

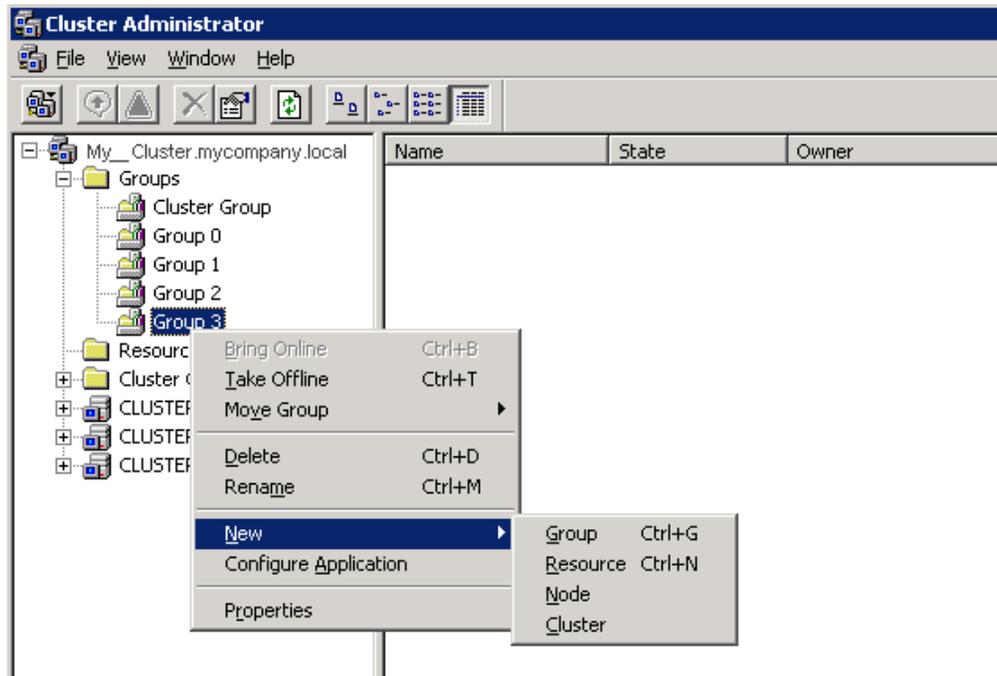
Once you have created all the necessary resource groups, make sure that they are hosted on the same node where you run the Cluster Administrator snap-in. To do this, click the **Groups** item and check the node name in the **Owner** column opposite the corresponding group. If a group is hosted on a different node, move it by right-clicking the group and selecting **Move Group > Node_Name** where *Node_Name* is the name of the node to move the group to.



Adding a Disk Resource

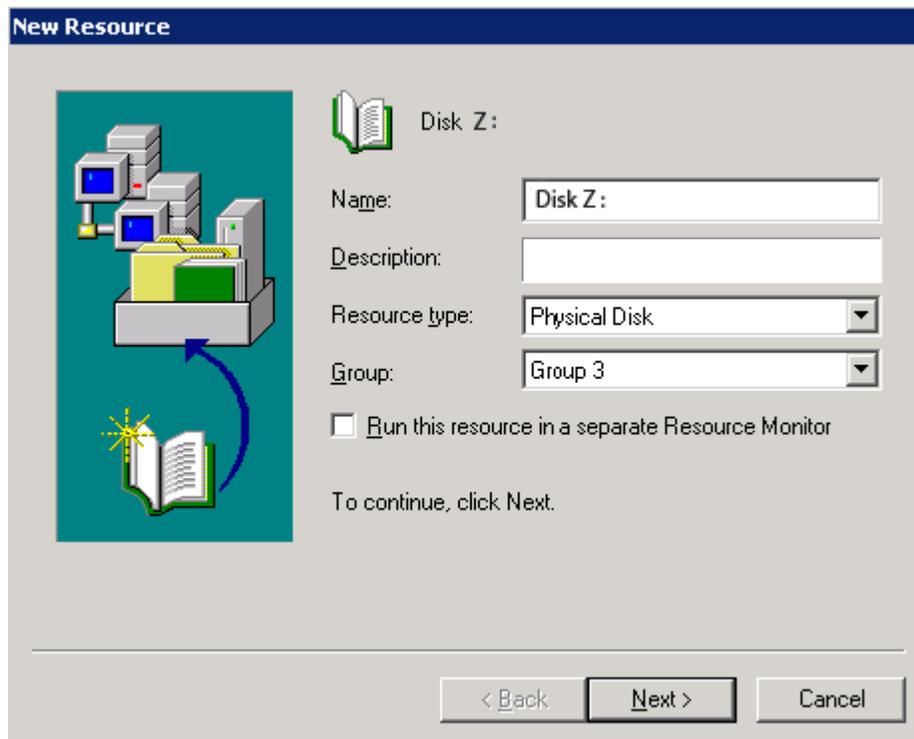
Each resource group intended for hosting the Parallels Containers resources must have a Disk resource. So in the next step, you need to create a Disk resource for each of your resource groups and add it to the corresponding group. To do this:

- 1 In the Cluster Administrator snap-in, expand the **Group** item, right-click the name of the group where you want to add the Disk resource, and select **New > Resource**.



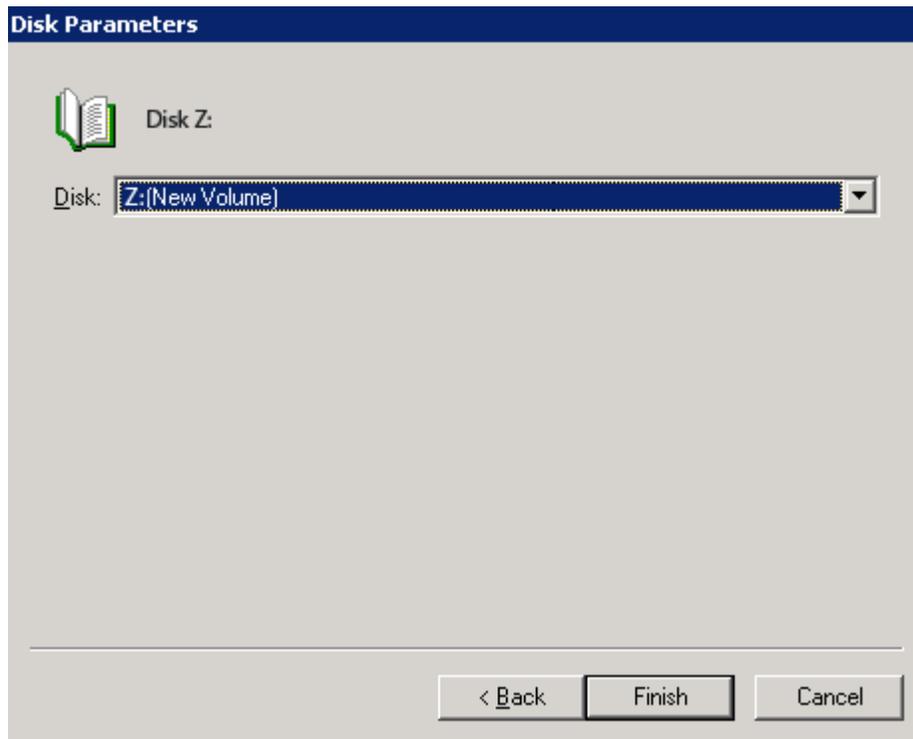
2 In the New Resource window, fill in the provided fields:

- In the Name field, specify a descriptive name for the Disk resource.
- In the Resource type field, choose Physical Disk from the drop-down menu.
- In the Group field, choose from the drop-down menu the group where you want to include the Disk resource.



When you are ready, click **Next**.

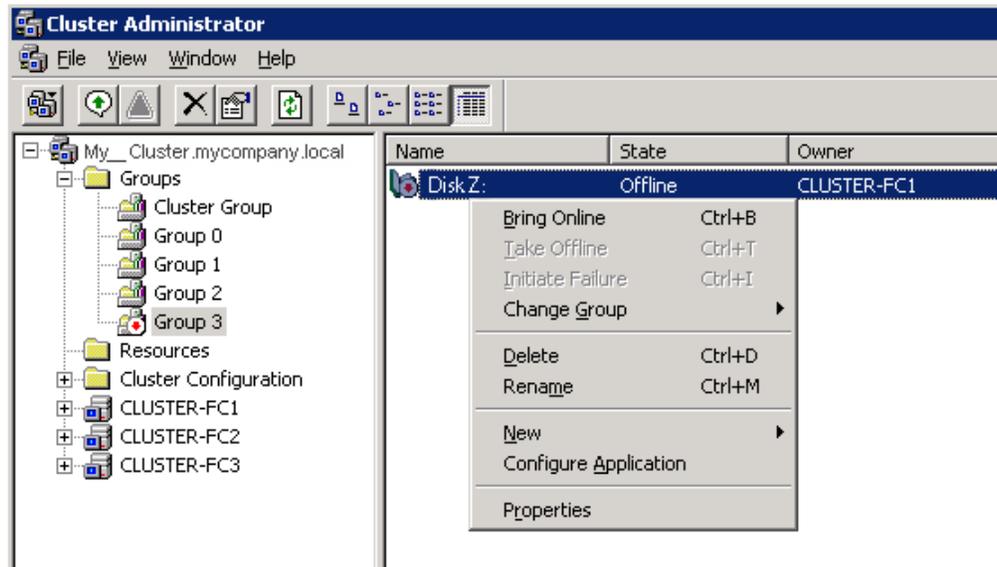
- 3** In the **Possible Owners** window, click **Next**.
- 4** In the **Dependencies** window, click **Next**.
- 5** In the **Disk Parameters** window, choose from the drop-down menu the drive letter assigned to the disk drive you want to use as the **Disk** resource, and click **Finish**.



Note: All the disks on the cluster storage device must be assigned drive letters. If a disk does not have a letter, it cannot be added to the resource group. For best practices for assigning drive letters to disks, see <http://support.microsoft.com/kb/318534>.

- 6 The dialog appears informing you that the disk resource has been successfully created. Click **OK** to close the dialog.

The newly created Disk resource appears in the center pane of the Cluster Administrator snap-in. To activate the Disk resource, right-click it, and select **Bring Online**.



After a while, the resource status should change from *Offline* to *Online*.

Adding IP Address and Network Name Resources

Along with the Disk resource, each resource group intended for hosting Parallels Containers resources must also contain an IP Address resource and a Network Name resource. You can add both resources to the corresponding resource groups

- when configuring the cluster or
- when installing Parallels Containers on your cluster nodes

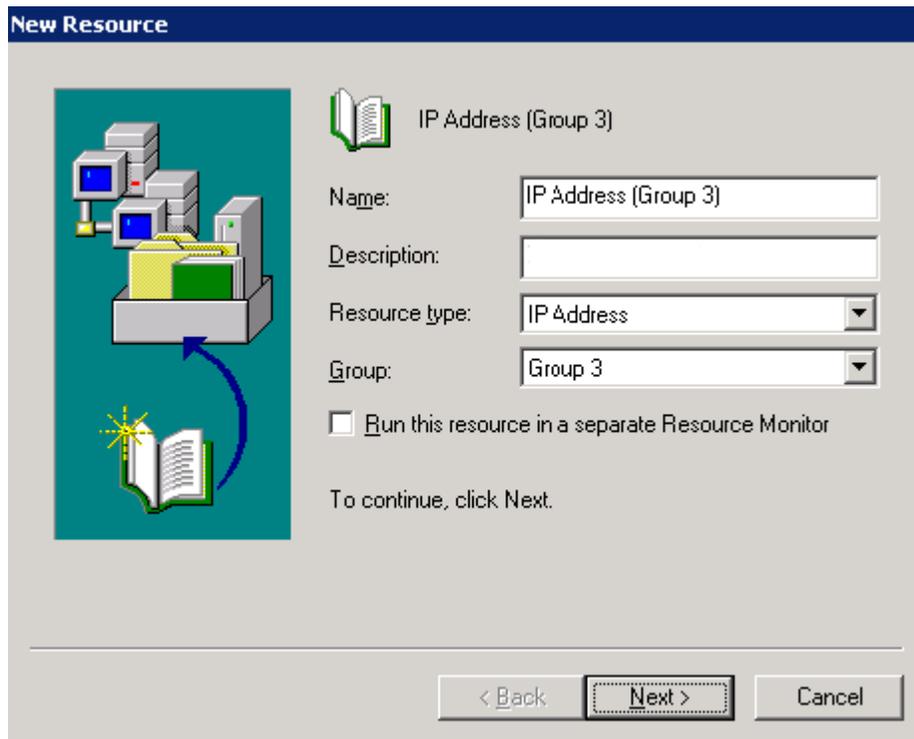
We highly recommend that you add these resources before starting the Parallels Containers installation. Doing so will ensure that all cluster parameters are configured properly and eliminate any errors that may be caused by providing incorrect parameters for IP Address and Network Name resources during the Parallels Containers installation.

The following subsections explain in detail how you can add IP Address and Network Name resources to your resource groups.

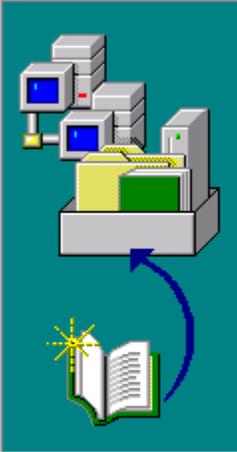
Adding an IP Address Resource

To add an IP Address resource to a resource group:

- 1 In the Cluster Administrator snap-in, expand the **Group** item, right-click the name of the group for which you want to create the IP Address resource, and select **New > Resource**.
- 2 In the New Resource window, fill in the provided fields:
 - In the **Name** field, specify a descriptive name for the IP Address resource.
 - In the **Resource type** field, choose **IP Address** from the drop-down menu.
 - In the **Group** field, choose from the drop-down menu the group where you want to include the IP address resource.



New Resource

 IP Address (Group 3)

Name:

Description:

Resource type:

Group:

Run this resource in a separate Resource Monitor

To continue, click Next.

< Back **Next >** Cancel

When you are ready, click **Next**.

- 3** In the **Possible Owners** window, click **Next**.
- 4** In the **Dependencies** window, click **Next**.
- 5** In the **TCP/IP Address Parameters** window, define the network parameters:
 - In the **Address** and **Subnet mask** fields, specify an IP address and network mask. Make sure that the indicated IP address is static, unique, and not assigned to any other computer on the network, including the other nodes in the cluster. Once the failover cluster is deployed, you will be able to use the indicated IP address to connect to the corresponding cluster node via the **Parallels Information Manager** application and manage this node and its **Containers** using this application.

When you are ready, click **Finish**.

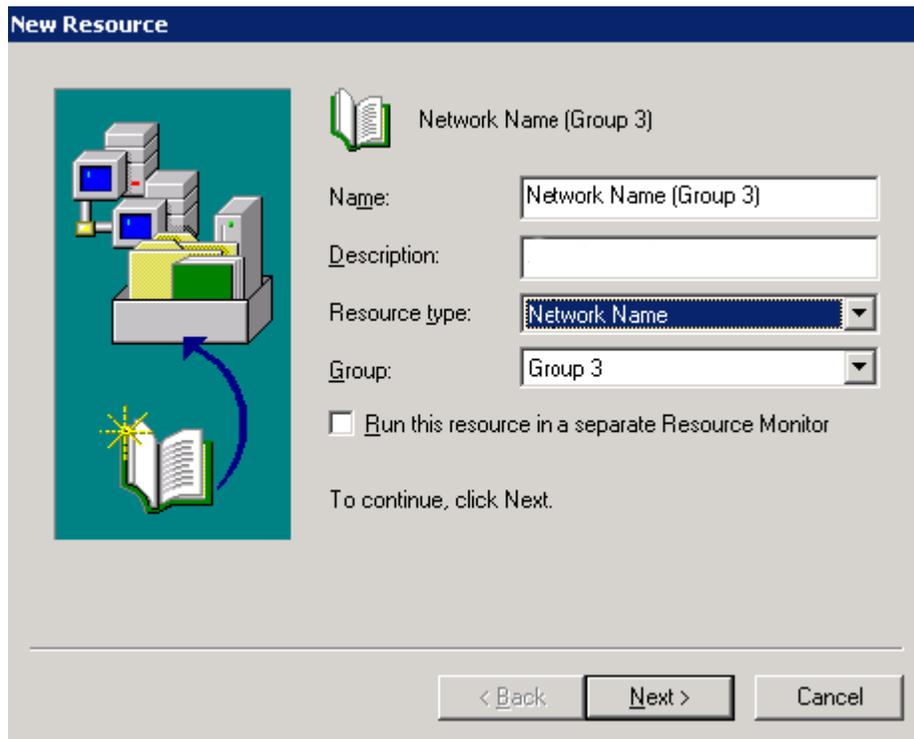
- 6 The dialog appears informing you that the IP Address resource has been successfully created. Click **OK** to close the dialog.

The newly created IP Address resource appears in the center pane of the Cluster Administrator snap-in. To activate the IP Address resource, right-click it, and select **Bring Online**. After a while, the resource status should change from **Offline** to **Online**.

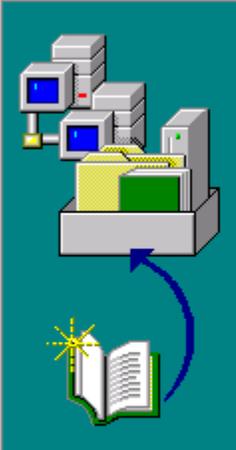
Adding a Network Name Resource

Once you have created the IP Address resource and added it to a resource group, you can include a Network Name resource in this group. To do this:

- 1 In the Cluster Administrator snap-in, expand the **Group** item, right-click the name of the group where you want to add the Network Name resource, and select **New > Resource**.
- 2 In the **New Resource** window, fill in the provided fields:
 - In the **Name** field, specify a descriptive name for the Network Name resource.
 - In the **Resource type** field, choose **Network Name** from the drop-down menu.
 - In the **Group** field, choose from the drop-down menu the group where you want to include the Network Name resource.



New Resource

 Network Name (Group 3)

Name:

Description:

Resource type:

Group:

Run this resource in a separate Resource Monitor

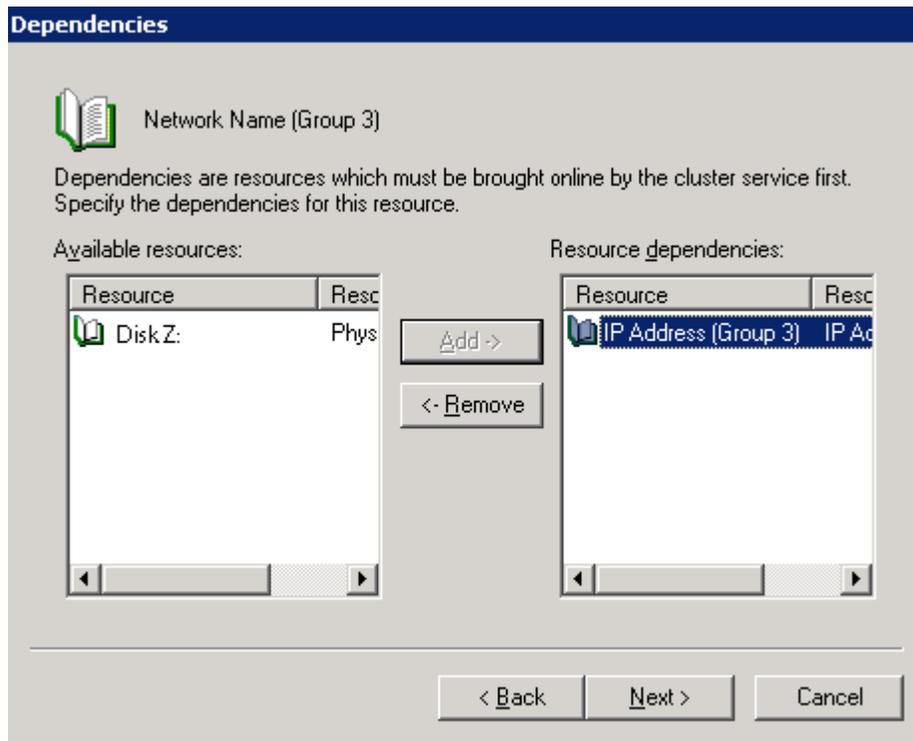
To continue, click Next.

When you are ready, click Next.

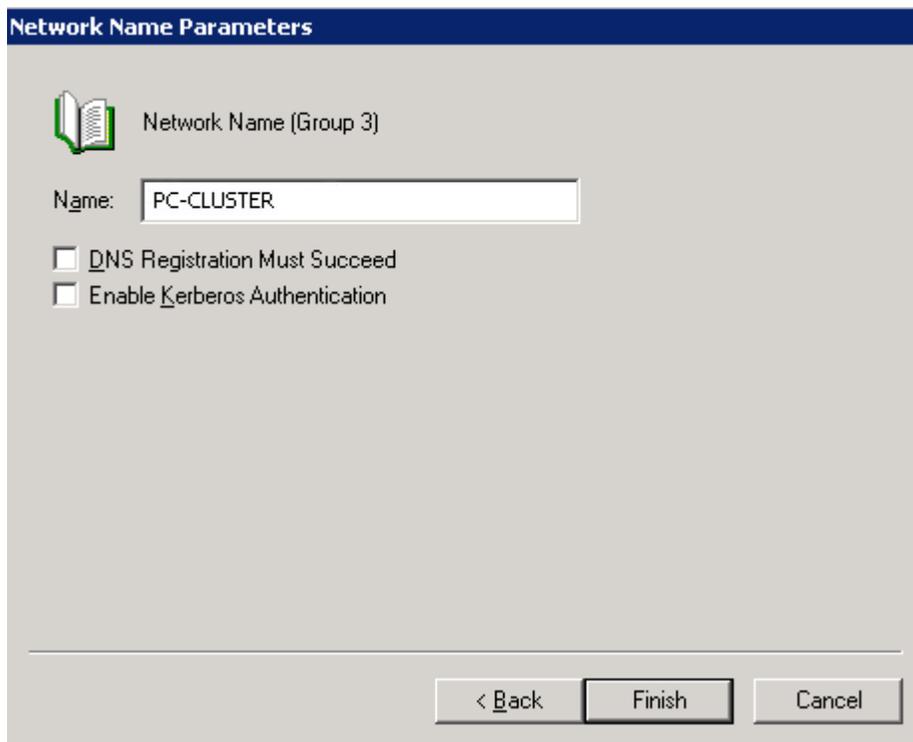
- 3 In the Possible Owners window, click Next.
- 4 In the Dependencies window, set the dependency for the Network Name resource on the IP Address resource (you created this IP Address resource in the previous step). To set the dependency, select IP Address in the Available resources table, and click Add. The IP Address resource should appear in the Resource dependencies table.

Note: For more information on resource dependencies, see

- [http://technet.microsoft.com/en-us/library/cc787307\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc787307(WS.10).aspx)
 - [http://technet.microsoft.com/en-us/library/cc783522\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc783522(WS.10).aspx)
-



- 5 In the Network Name Parameters window, specify a descriptive name for your Network Name resource, and click Finish. This name will be associated with the IP address you specify in the IP address field using the dependency you set in Step 4. Make sure that this network name is not already assigned to any other computer on the network, including the other nodes in the cluster.



- 6 The dialog appears informing you that the Network Name resource has been successfully created. Click **OK** to close the dialog.

The newly created Network Name resource appears in the center pane of the Cluster Administrator snap-in. To activate the resource, right-click the Network Name resource, and select **Bring Online**. After a while, the resource status should change from *Offline* to *Online*.

Checking Server Cluster Operation

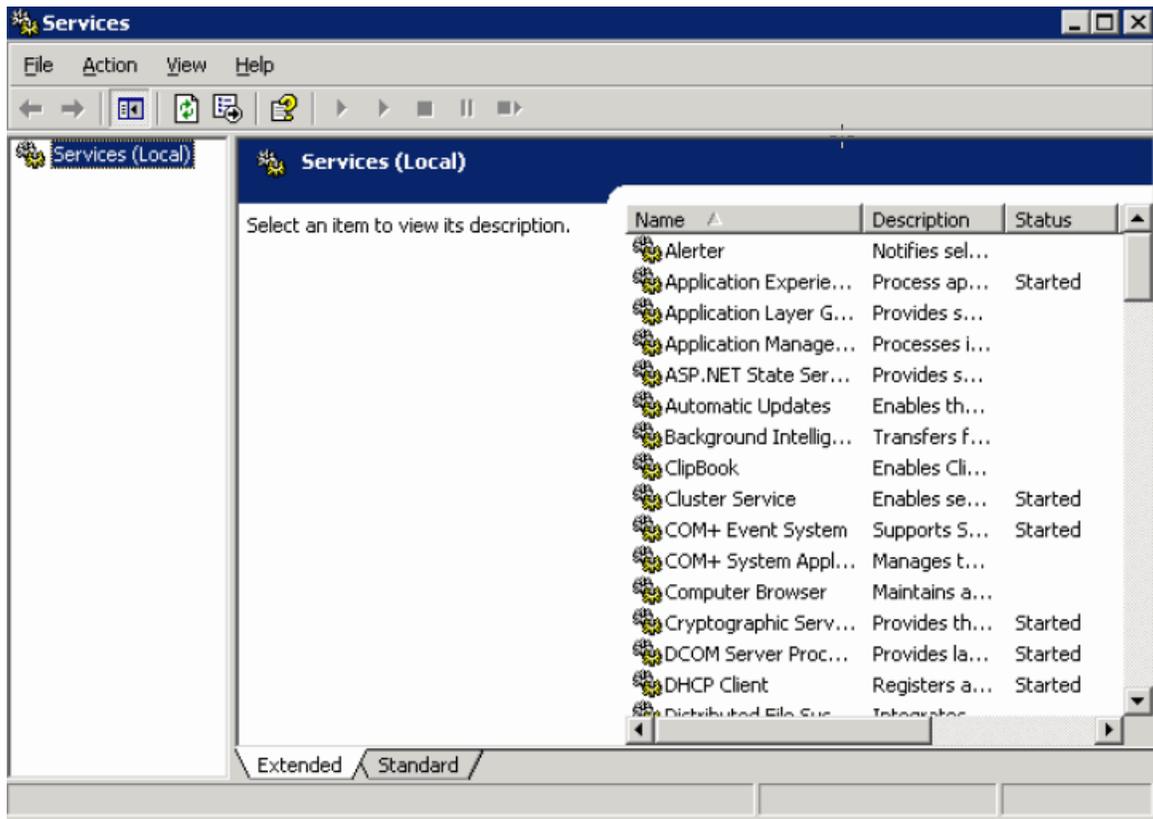
You can use the following procedures to verify the availability of the newly created cluster:

First Procedure (viewing cluster nodes status in Cluster Administrator):

- 1 Log in to any of your cluster nodes using the domain administrator account and launch the Cluster Administrator application by clicking **Start**, pointing to **Administrative Tools**, and selecting **Cluster Administrator**.
- 2 Connect to the newly created cluster:
 - In the **Open Connection to Cluster** window, select the **Open connection to cluster** item on the **Action** menu.
 - In the **Cluster or server name** field, type the name of the cluster, and click **OK**.
- 3 Verify that all nodes registered in the cluster are listed in the **Cluster Administrator** main window.

Second Procedure (checking the Cluster service availability):

- 1 On each cluster node, click **Start**, point to **Administrative Tools**, and select **Services** to launch the **Services** snap-in.
- 2 In the right pane of the displayed window, verify that the **Cluster Service** service is listed in the **Services** table and its status is set to **Started**.

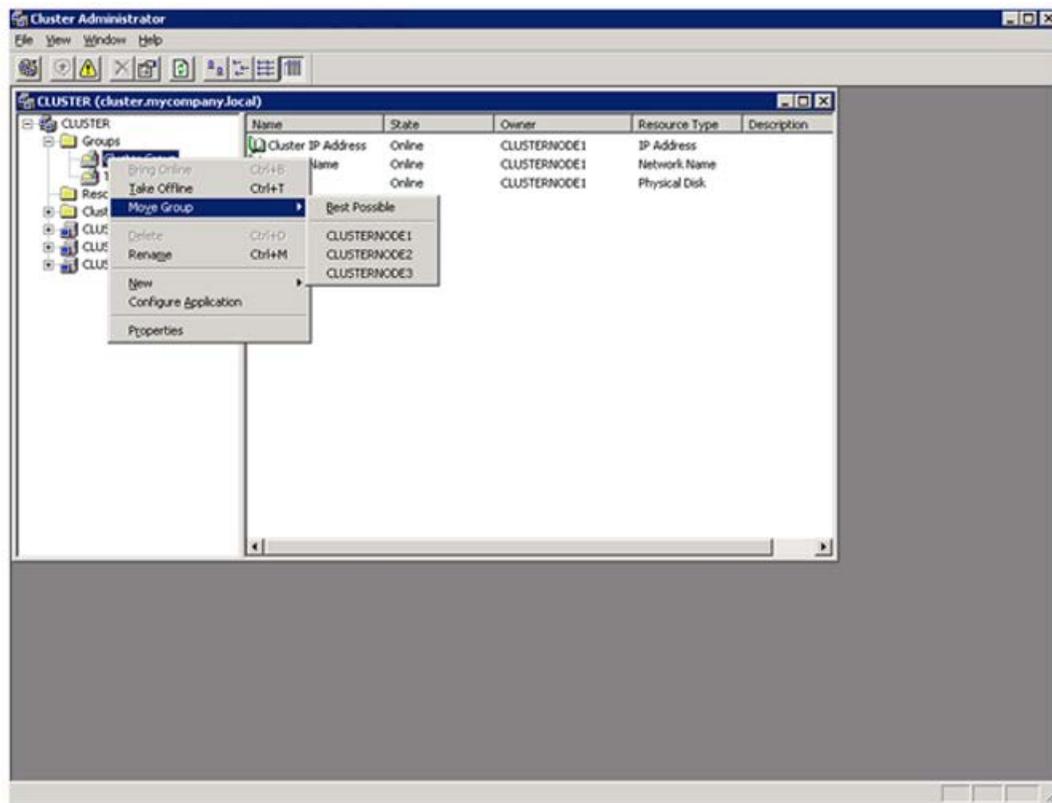


Third Procedure (checking the failover of cluster resources from one node to another):

- 1 Log in to any of your cluster nodes.
- 2 Click **Start**, point to **Administrative Tools**, and select **Cluster Administrator** to launch the Cluster Administrator application.
- 3 In the left pane of the displayed window, expand the **Groups** folder, and select **Cluster Group** to view the cluster core resources. The information on these resources is presented in the following table:

Column Name	Description
Name	The name assigned to the cluster resource.
State	The current state of the cluster resource.
Owner	The name of the cluster node currently hosting the cluster resource.
Resource Type	The type of the cluster resource.
Description	Additional information on the cluster resource.

- 4 Right-click the **Cluster Group** item, point to **Move Group**, and select the node where you want to move the cluster resources.



The failover process completes successfully if the cluster resources are brought online on the selected node. The name of this node should be displayed in the **Owner** field.

Installing Parallels Containers on Cluster Nodes

The procedure of installing Parallels Containers on servers (also known as nodes) participating in a failover cluster slightly differs from that of installing it on standalone servers. Use the guidelines below to correctly install Parallels Containers on your cluster nodes.

General Guidelines

Before starting the installation, be sure to read the following information:

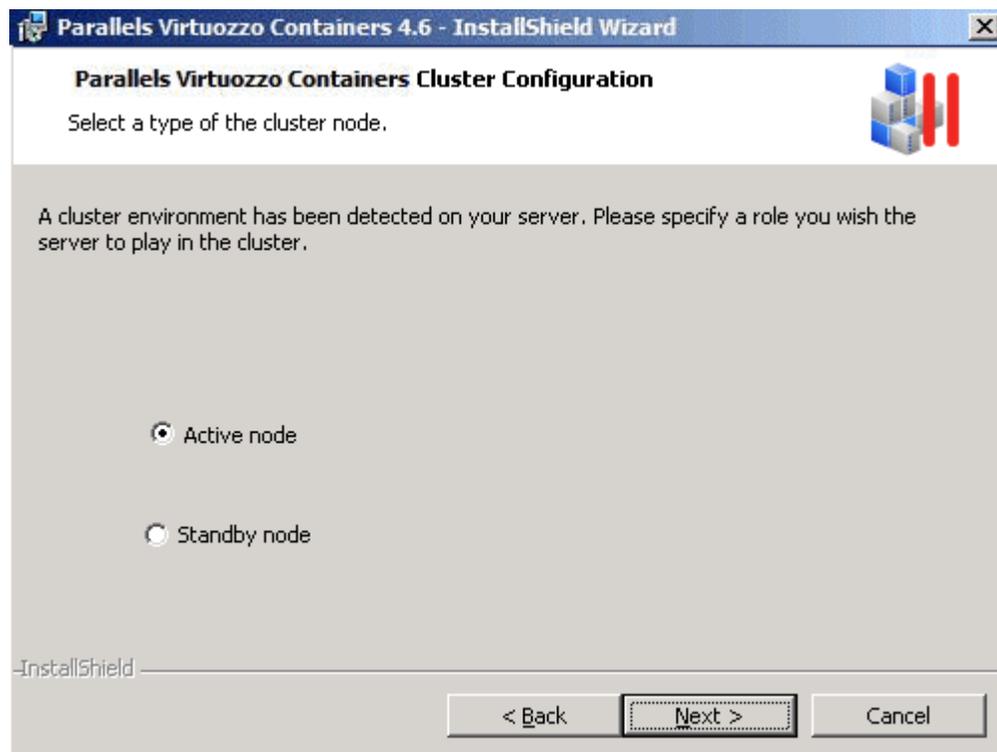
- You must install the Parallels Containers software on each node you plan to include in the failover cluster.
- You must install the Parallels Containers software on one cluster node at a time.
- The order of installing the Parallels Containers software on active and passive nodes is not important. It means that you can, for example, first install Parallels Containers on an active node, then proceed with a standby node, and then install the software on another active node.

- The Parallels Containers installation files must be located either on a local disk drive or on a CD or DVD inserted in the node's local CD/DVD-ROM drive. The Parallels Containers installation will fail if the installation files are located on a network drive or clustered disk.

Installation Guidelines

To install Parallels Containers on a cluster node, do the following:

- 1 Follow the installation instructions in the *Parallels Containers 4.6 Installation Guide* until the Parallels Containers Cluster Configuration window is displayed. In this window you are asked to specify a role you want the Hardware Node to play in the cluster.



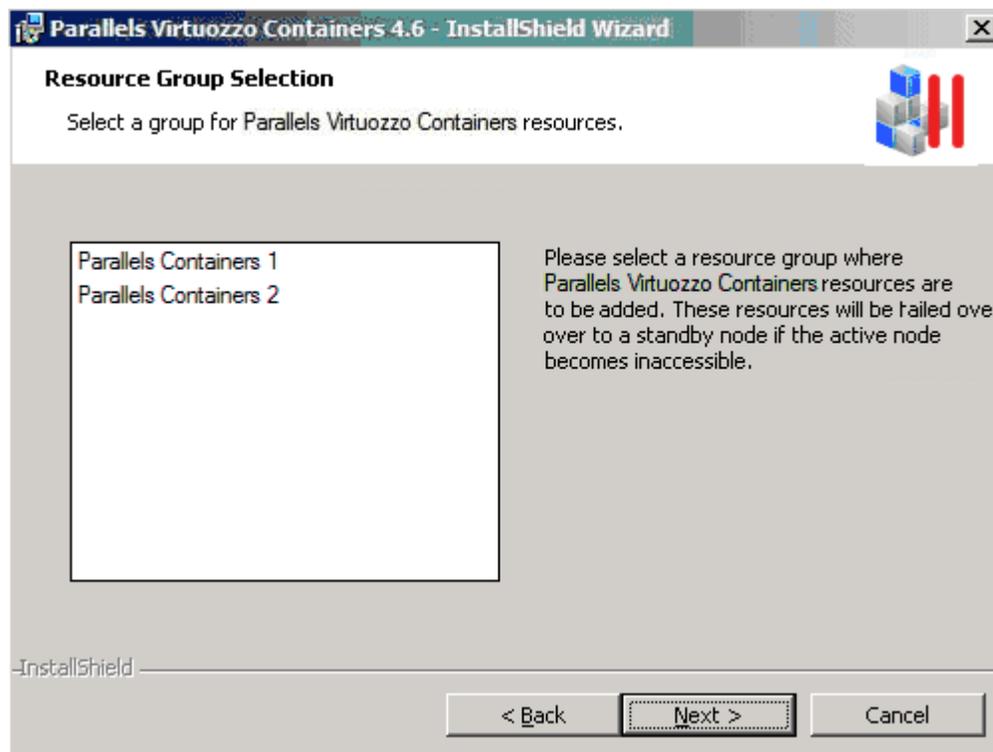
You can choose one of the following options:

- Select the **Active node** radio button to make the server operate as an active node in the cluster. This node will host all relevant Parallels Containers resources and fail them over to a standby node in the case of a failure. If you are going to deploy an active/active failover cluster, you must select the **Active node** radio button for each node that will participate in this cluster. For more information on active/active clusters, see **Active/Active Failover Clusters** (p. 9).
- Select the **Standby node** radio button to make the server operate as a standby node in the cluster. This node will take control over the Parallels Containers software and its resources in the case of an active node failure.

When assigning roles to your nodes in the Parallels failover cluster, keep in mind the following:

- a You can choose any of the aforementioned options irrespective of the role played by the server in the cluster before installing Parallels Containers on it.

- b A role is not a permanent property of a cluster node, i.e. one and the same node can act as both active and standby in different periods of time, irrespective of the role assigned to the node during the Parallels Containers installation. For example, a standby node can become active when it takes control over the group responsible for keeping Parallels Containers-specific resources in the event of an active node failure. The failed active node, in turn, automatically becomes standby when it is taken online after its failure provided that the Failback feature is not enabled for Parallels Containers resources on this node (for more information on the Failback feature, see the next section).
- 2 On the Resource Group Selection screen, choose a resource group (known as a clustered service in Windows Server 2008 and Windows Server 2008 R2 and as a resource group in Windows Server 2003) for Parallels Containers-specific resources that will be able to fail over to a standby node in the case of a failure.



The table on this screen lists all groups existing in your cluster and containing the *physical disk* resource. The latter must be present in a group to successfully add the Parallels Containers resources to this group. The main Parallels Containers resources which are failed over to a standby node in the event of an active node failure are given below:

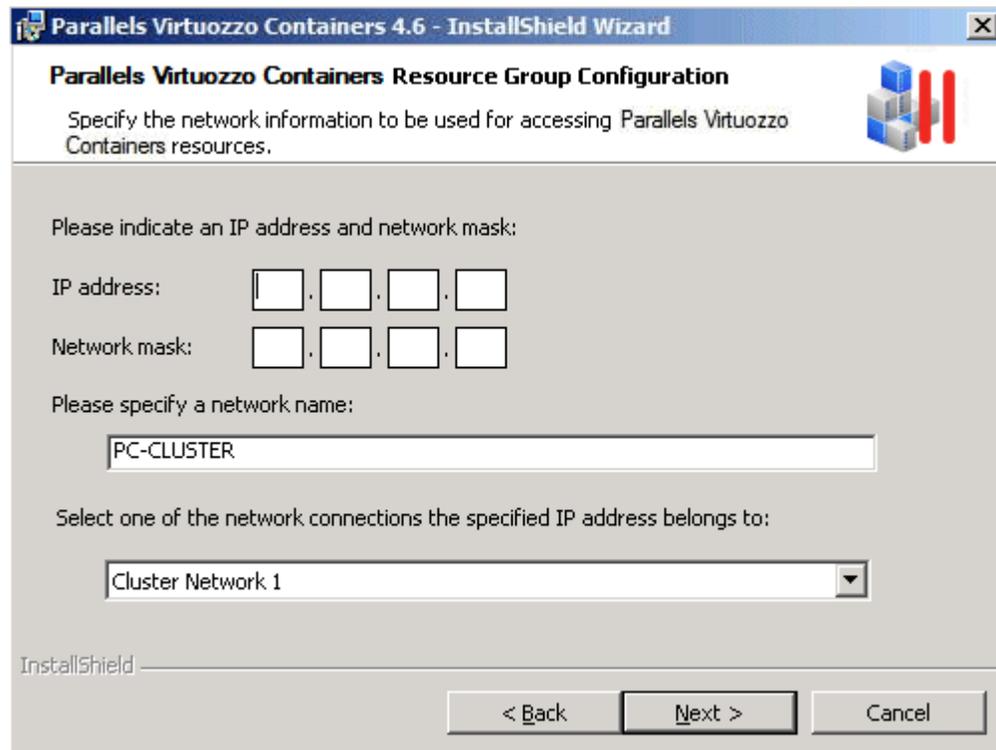
- Virtuozzo VZFS Database
- Virtuozzo Server
- Parallels Agent
- Parallels Agent Compat
- Container1

Besides, any Container is registered as the `Virtuozzo Container` resource in the cluster during its startup on the active node.

Note: This step is skipped if you are deploying Parallels Containers on a standby node.

- 3** When specifying the location of Parallels Containers program files, Container data, and backups on the **Locations of Parallels Containers Data and Program Files** screen, keep in mind the following:
 - All nodes in the cluster (both active and standby) must have the Parallels program files installed in one and the same folder (e.g., `C:\Program Files\Parallels\Containers`).
 - For each active node in the cluster, Parallels Containers allows you to store Container data and backups only on one of those disks that are registered in the resource group you specified on the **Resource Group Selection** screen (see **Step 2** above).
 - When choosing a disk for storing Container data and backups, make sure that this disk has a drive letter assigned to it.
 - For all standby nodes in the cluster, you are not allowed to modify the location of Container data and backups since these nodes make use of the corresponding folders created during the Parallels Containers installation on the active node and only in the case of an active node failure.
- 4** On the **Parallels Containers Resource Group Configuration** screen, configure the network parameters for the resources group you chose for hosting Parallels Containers-specific resources (see **Step 2** above):
 - In the **IP address** and **Network mask** fields, specify an IP address and network mask. Make sure that the indicated IP address is static, unique, and not assigned to any other computer on the network, including the other nodes in the cluster.
 - In the **Specify a network name...** field, indicate a network name. This name will be associated with the IP address you specify in the **IP address** field above. Make sure that this network name is not already assigned to any other node in the cluster.
 - In the **Select one of the network connections...** field, choose a network connection the IP address specified in the **IP address** field belongs to.

After installing Parallels Containers, you can use the indicated IP address or network name to connect to the cluster node via the Parallels Virtual Automation application and manage this node and its Containers using this application.



Notes:

1. This step is skipped if you are deploying Parallels Containers on a standby node.
 2. The fields on the **Parallels Containers Resource Group Configuration** screen are automatically filled in by the installer if you added all the necessary network parameters before starting the Parallels Containers installation. See **Configuring the Cluster** (p. 31, p. 46) for details.
-

The other peculiarities that should be taken into account when installing Parallels Containers on cluster nodes are the following:

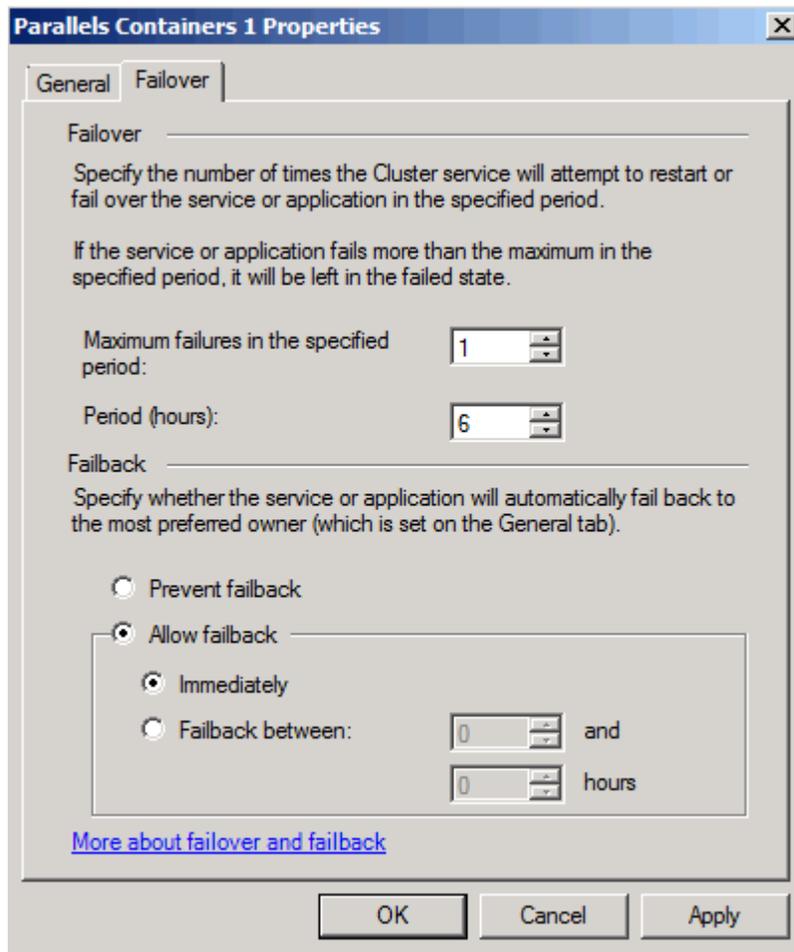
- A number of operations which are normally performed during the Parallels Containers installation on a standalone server are skipped when installing Parallels Containers on a standby node (e.g. the step of installing Parallels Containers components and creating the Service Container). This is explained by the fact that standby nodes do not have their own Parallels Containers resources. They only can make use of the resources controlled by the corresponding active node (e.g., in the case of an active node failure or when the active node is taken offline as part of your maintenance plan).
- Make sure that a valid Parallels Containers license is installed not only on the active nodes but also on all standby nodes participating in the cluster. Otherwise, your standby nodes will not be able to take control over the Parallels Containers software and its resources in the case of an active node failure.
- We highly recommend that you install the same set of Parallels Containers and Windows Server updates on all nodes in the failover cluster.

Configuring Automatic Failback

By default, if an active node fails, a standby node becomes active taking control over the group hosted on the active node and keeping all Parallels Containers-related resources. After the failure associated with the active node has been repaired and the node has been brought online again, this node becomes standby. You can change this behavior by enabling the *Failback* feature, which allows the group to automatically fail back from a standby node to an active one when the latter is taken online again. You can turn on the Failback feature as follows.

On nodes running Windows Server 2008 and Windows Server 2008 R2:

- 1** Log in to any of your cluster nodes, and click **Start > All Programs > Administrative Tools > Failover Cluster Manager** to open the Failover Cluster Manager snap-in.
- 2** Expand the **Services and applications** item, right-click the clustered service for which you want to enable the Failback feature, and choose **Properties**.
- 3** On the **Failover** tab of the displayed window, select the **Allow failback** radio button, and do one of the following:
 - Select the **Immediately** radio button if you want the clustered service to immediately fail back to an originally active node when it is brought online again.
 - Select the **Failback between...** radio button, and set the time interval after which the clustered service will fail back to an originally active node.



On nodes running Windows Server 2003:

- 1 Log in to any of your cluster nodes, and open the Cluster Administrator application by clicking Start, pointing to Administrative Tools, and selecting Cluster Administrator.
- 2 In the Cluster Administrator window, select the group for which you want to enable the Failback feature, and click the Properties icon.
- 3 On the Failback tab of the displayed window, select the Allow failback radio button, and do one of the following:
 - Select the Immediately radio button if you want the group to immediately fail back to an originally active node when it is brought online again.
 - Select the Failback between... radio button, and set the time interval after which the group will fail back to an originally active node.

Adding Nodes to the Failover Cluster

Sometimes, you may wish to join additional nodes to your failover cluster. For example, this may be the case if you wish to increase the number of active or standby nodes participating in the cluster or need to

replace a broken node with a new one. The process of adding a new node to a failover cluster is identical to that of configuring any of its original nodes and includes the following basic operations:

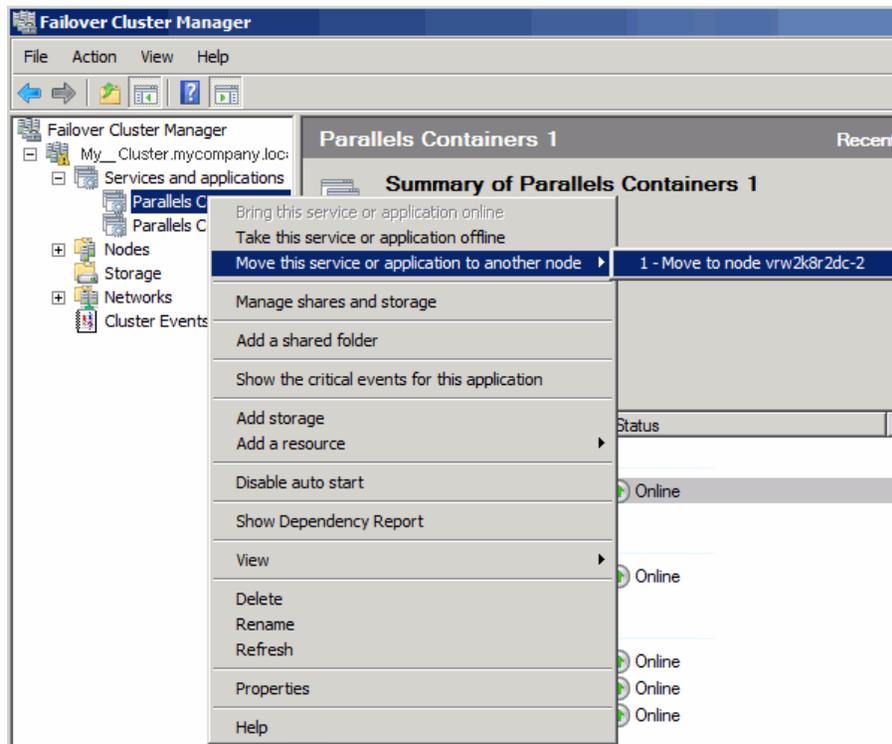
- 1 Preparing all the necessary hardware for the node. See the **Setting Up Hardware for the Cluster** section for details (p. 11).
- 2 Configuring the node. See the **Configuring Physical Servers** section for details (p. 12).
- 3 Adding the node to the Active Directory domain. See the **Adding Nodes to the Domain** section for details (p. 20).
- 4 Connecting the node to your shared storage. See the **Connecting Servers to a Shared Storage** section for details (p. 22).
- 5 Adding the node to the failover cluster. See the **Setting Up the Cluster** section for details (p. 23).
- 6 Installing Parallels Containers on the newly configured cluster node. See the **Installing Parallels Containers on Cluster Nodes** section for details (p. 60).

Testing the Failover Cluster

You can test the operation of your failover cluster by completing the following tasks:

For clusters consisting of nodes with Windows Server 2008 or Windows Server 2008 R2:

- 1 Log in to any of your cluster nodes using the domain administrator account.
- 2 Click **Start > Administrative Tools > Failover Cluster Management** to open the Failover Cluster Management snap-in.
- 3 Expand the **Services and Applications** item, right-click the group with Parallels Containers resources, and select **Move this service or application to another node > Node_Name** (where *Node_Name* is the name of the cluster node where the group will be moved).



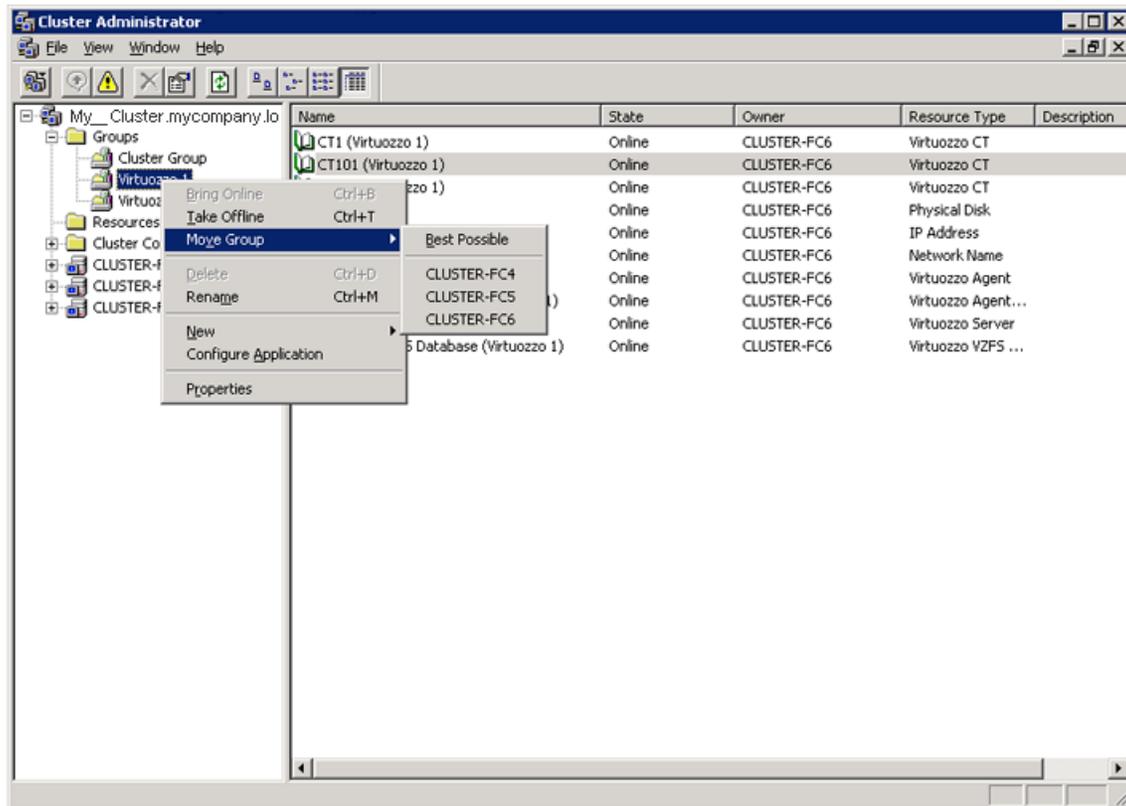
- 4 Make sure that the group has been successfully moved to the selected node and all resources are online.

For clusters consisting of nodes with Windows Server 2003:

- 1 Log in to any of your cluster nodes.
- 2 Click **Start**, point to **Administrative Tools**, and select **Cluster Administrator** to launch the Cluster Administrator application.
- 3 In the left pane of the Cluster Administrator window, expand the **Groups** folder, and click the name of the group storing Parallels Containers-related resources. The information on the displayed resources is presented in the following table:

Column Name	Description
Name	The name assigned to the resource.
State	The current state of the resource.
Owner	The name of the cluster node currently hosting the resource.
Resource Type	The type of the resource.
Description	Additional information on the resource.

- 4 Right-click the group, point to **Move Group**, and select the node where you want to move the resources.



The failover process completes successfully if all Parallels Containers-related resources are brought online on the selected node (the name of this node should be displayed in the **Owner** field).

Troubleshooting Failover Clusters

This section describes some problems you may encounter when working with failover clusters and provides possible solutions to these problems.

Failure to bring a resource group offline or online

You try to bring a Parallels Containers resource group offline or online, but the operation takes a long time and eventually fails. The reason for this problem may be as follows.

Some Containers, each treated as a Parallels Containers resource, may require a long time to start or shut down. If this time exceeds the values of the Pending Timeout and Deadlock Timeout parameters set for the resource, the operation fails and the clustering software puts the resource group to the **Failed** state.

Solution

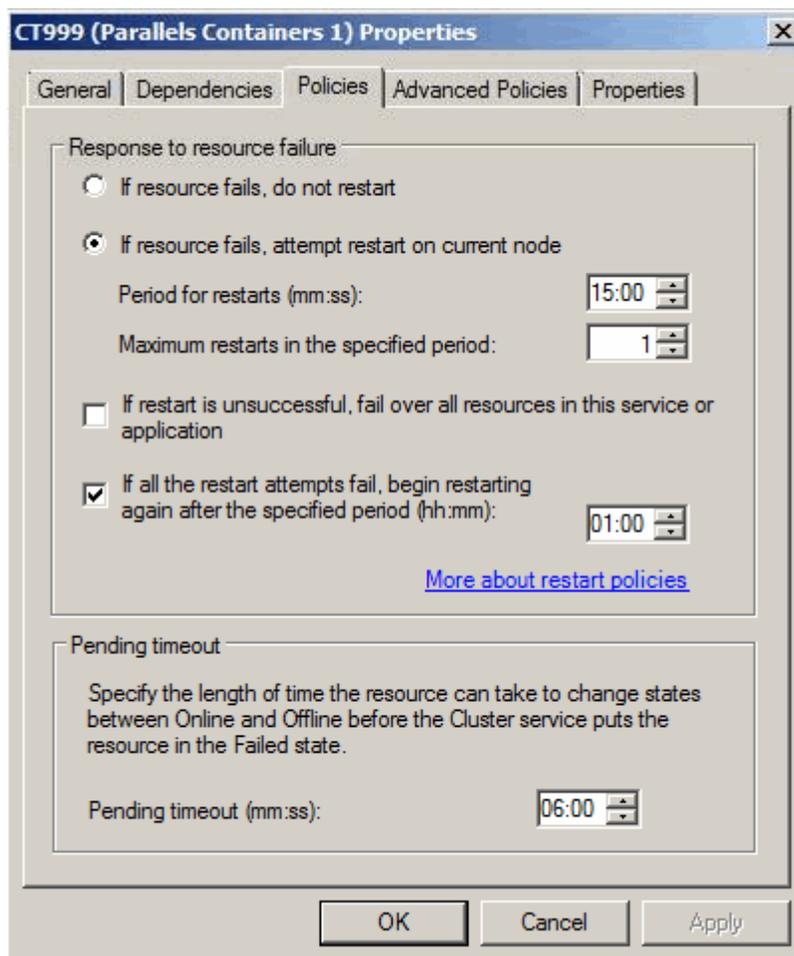
To solve the problem, make sure that the time needed to stop and start your Containers does not exceed the values set for the Pending Timeout and Deadlock Timeout parameters. Both parameters are

configurable per Container so be sure to check them for each Container. If necessary, tune these parameters to meet your demands.

Configuring the Pending Timeout parameter

To configure the pending timeout for a Parallels Containers resource:

- 1 In the Failover Cluster Manager snap-in (the Failover Cluster Management snap-in in Windows Server 2003), expand the console tree under the cluster you want to configure.
- 2 Expand the Services and applications folder.
- 3 Click the resource group hosting the Container for which you want to change the pending timeout.
- 4 In the center pane, right-click the resource (that is, the necessary Container), and click Properties.
- 5 Click the Policies tab.



- 6 In the Pending timeout section, specify the desired pending timeout, and click OK.

Configuring the Deadlock Timeout parameter

The Deadlock Timeout parameter is available (and therefore can be configured) only for Hardware Nodes running Windows Server 2008 or Windows Server 2008 R2. You can configure this parameter using the `cluster res` command. For example, to increase the Deadlock Timeout value to 5 minutes (which equals 300 000 milliseconds) for Container 999 that belongs to the VZ1 group, you can run this command:

```
C:\...\Administrator>cluster res "999 (VZ1)" -prop DeadlockTimeout=300000
```

Note: You need to be a member of the local Administrators group to complete the operations above.

Creating Clusters From Containers

Parallels Containers allows you to include Containers in a failover cluster so that they can act as full participants (nodes) in the cluster and provide high availability for your applications and services. In Parallels Containers, you can create a cluster in one of the following ways:

- using a SAN (Storage Area Network)-based shared storage device
- using a 'loopback file'-based shared storage device

As is evident from their names, these ways differ from each other in the type of the shared storage device to be used in the cluster. While the first way assumes using a shared physical disk drive as a shared storage device, the second way requires the usage of a special loopback file created on the Hardware Node and emulating a shared physical disk drive in the cluster. Both ways of the cluster creation are described in the following subsections in detail.

Note: The examples used in the following subsections show only several possible cluster configurations you can to deploy in your working environments. You can hold to these configurations or work out your own ones and build your own cluster architecture.

In This Chapter

Setting Up a SAN-Based Cluster	72
Setting Up a 'Loopback File'-Based Cluster.....	74

Setting Up a SAN-Based Cluster

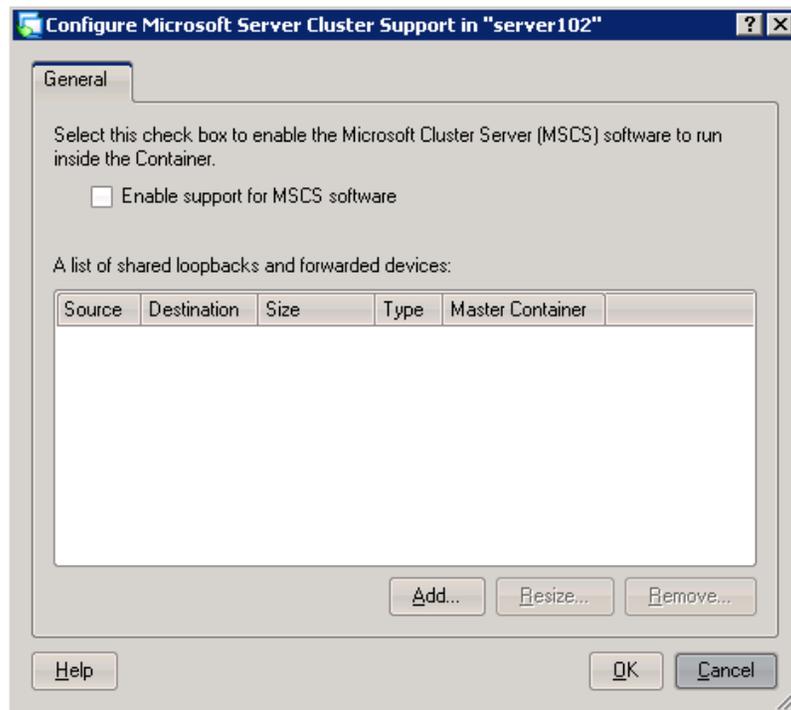
In Parallels Containers, you can include Containers, along with standalone physical servers, in clusters to provide higher availability for your applications and services (for example, for the MS SQL Server application). When creating such a cluster, keep in mind that all Containers participating in the cluster must reside on different Hardware Nodes since one and the same shared storage device on a Node (usually, a shared SCSI disk) can be forwarded to only one Container on this Node.

Notes:

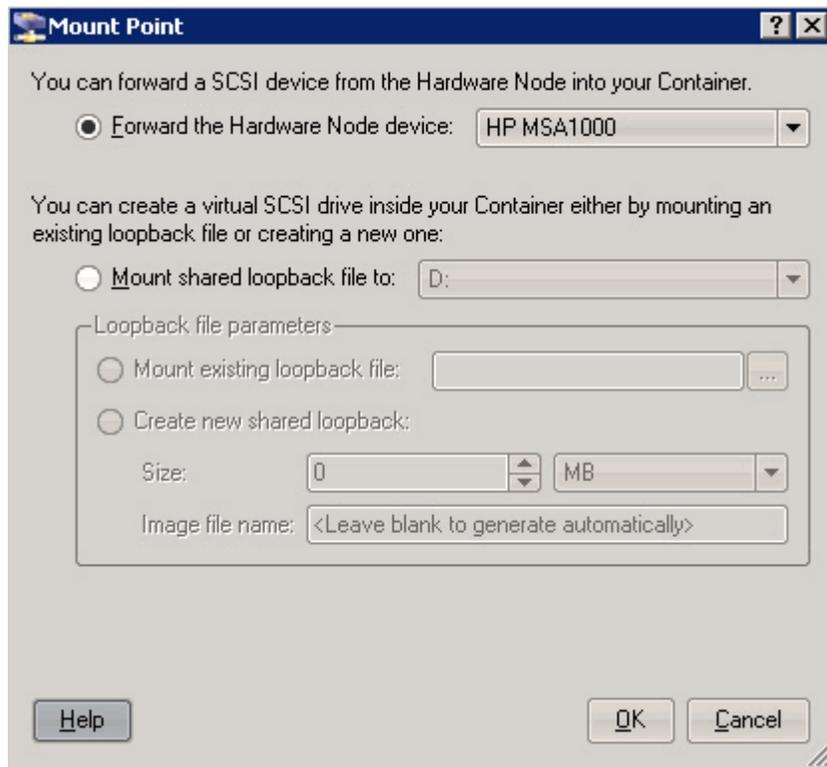
1. It is a good practice to make a standalone physical server act as an active node in the cluster and a Container as the standby one.
 2. Consult the documentation for the application you want to cluster to learn how many nodes you need to include in the cluster and what roles (active or passive) the nodes will play in it.
-

The following example demonstrates how to create a cluster that will include two Containers and use a SAN-based shared storage device (a SCSI disk drive) as a cluster storage—the place where your high-availability applications and all their data will be hosted:

- 1 Enable the clustering support (which is disabled by default for all Containers on the Node) inside the Container you wish to act as the active node in the cluster. To do this:
 - a Select the **Parallels Containers** item under the corresponding **Hardware Node** name.
 - b Right-click the Container to be included in the cluster, and select **Tasks > Configure Support for Windows Cluster** on the context menu:



- c In the displayed window, select the **Enable support for MSCS software** check box.
 - d Click **OK**.
 - e If the Container is running, restart it for the changes to take effect.
- 2 Forward the shared SCSI disk to be used as a cluster storage to the Container:
 - a Select the **Parallels Containers** item under the corresponding **Hardware Node** name.
 - b Right-click the Container which is to act as the active node in the cluster, and select **Tasks > Configure Support for Windows Cluster**.
 - c In the displayed window, click the **Add** button:



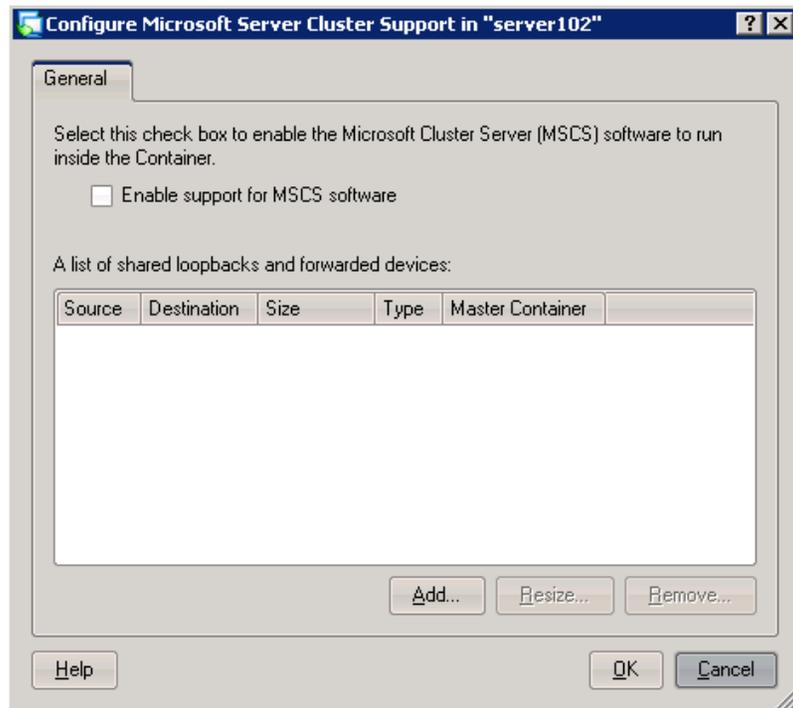
- d Select the **Forward the Hardware Node device** radio button, and choose the needed SCSI device on the drop-down menu.
 - e Click **OK**.
- 3 Perform **Steps 1-2** for each standby Container to be included in the cluster.
 - 4 Log in to the Container via RDP, create the cluster, and register both Containers in the cluster. For information on how to create clusters, refer to the **Setting Up a Cluster** section (p. 23).

Note: All SCSI disk drives forwarded from the Hardware Node to a Container are not kept during the Container migration.

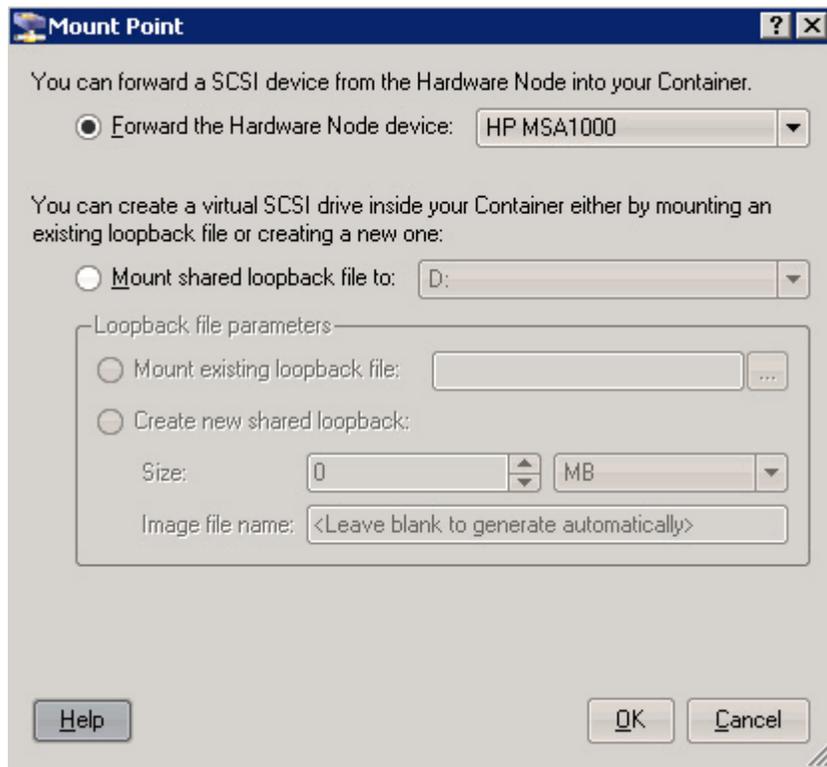
Setting Up a 'Loopback File'-Based Cluster

Another way of creating a cluster in Parallels Containers-based systems is to use loopback files as shared cluster storage devices. These files are mounted inside the corresponding Containers and emulate SCSI devices on a single SCSI bus inside these Containers. To create a cluster using a loopback file as a cluster storage device, perform the following operations:

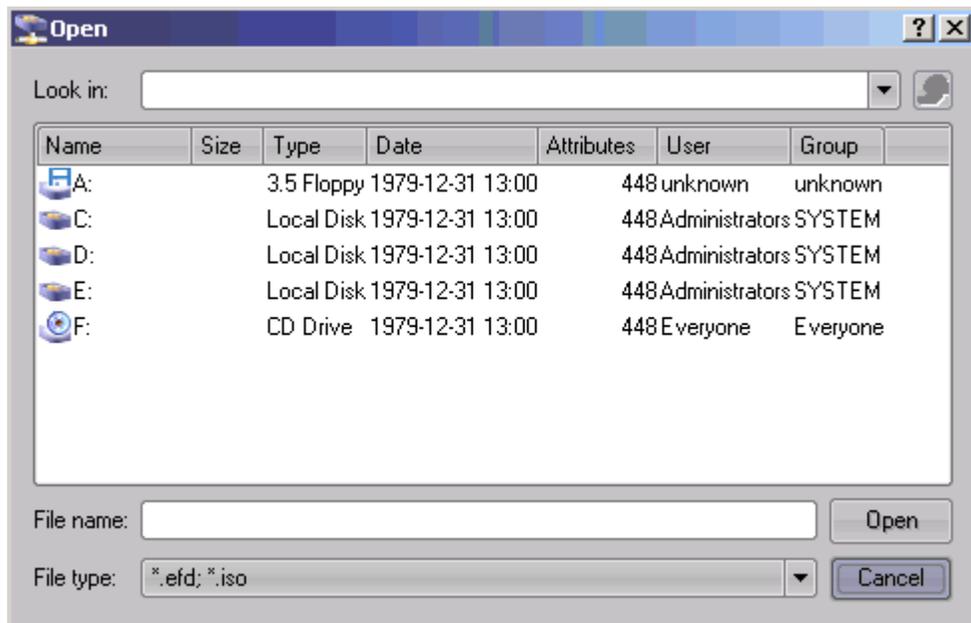
- 1 Decide on the Containers to be included in the cluster. Keep in mind that all Containers to participate in the cluster must reside on one and the same Hardware Node, i.e. on the Node where the corresponding loopback file is located.
- 2 Enable the clustering support (which is disabled by default for all Containers on the Node) inside the Container you want to act as the active node in the cluster. To do this:
 - a Select the **Parallels Containers** item under the corresponding Hardware Node name.
 - b Right-click the Container to be included in the cluster, and select **Tasks > Configure Support for Windows Cluster**.



- c In the displayed window, select the **Enable support for MSCS software** check box.
 - d Click **OK**.
 - e If the Container is running, restart it for the changes to take effect.
- 3 Create a loopback file to be used as the shared SCSI disk in the cluster:
 - a Select the **Parallels Containers** item under the corresponding Hardware Node name.
 - b Right-click the Container which is to act as the active node in the cluster, and select **Tasks > Configure Support for Windows Cluster** on the context menu.
 - c In the displayed window, click the **Add** button:



- d In the **Mount Point** window, select the **Mount shared loopback file to** radio button and, on the drop-down menu, choose a drive letter under which the loopback file emulating the SCSI device will be mounted inside the Container. Decide on the loopback file to be used as the shared storage device:
 - Create a new loopback file inside the Container by selecting the **Create new shared loopback** radio button and specifying the size and the name of the loopback file.
 - Mount an existing loopback file from the Hardware Node by selecting the **Mount existing loopback file** radio button, clicking the ... button, and specifying the path to the needed loopback file in the displayed window.
- e Click **OK**.
- 4 Log in to the Container via RDP, and create the cluster. For information on how to create clusters, refer to the **Setting Up a Cluster** section (p. 23).
- 5 Perform Steps 2-3 for each 'passive' Container to be included in the cluster. When executing Step 3, mount the same loopback file(s) as the one(s) mounted inside the active Container:
 - a Select the **Mount existing loopback file** radio button in the **Mount Point** window.
 - b Click the ... button, and in the displayed window, specify the path to the needed loopback file, for example:



Loopback files are located in the `X:\vz\private\CT_ID` folder on the Hardware Node where `CT_ID` denotes the ID of the Container where the corresponding loopback file was created.

- c** Click **Open**, and then click **OK**.
- 6** Add all standby Containers to the cluster. For information on how to add servers to an existing cluster, see the **Setting Up a Cluster** section (p. 23).

Notes:

1. A Container with a shared loopback file created inside cannot be migrated.
 2. If a Container has one or more shared loopback files mounted (e.g., all standby Containers in the example above), these loopback files will not be kept during the Container migration.
-

Setting Up an NLB Cluster

In Parallels Containers, any Container can act as a full participant in a Network Load Balancing (NLB) cluster. This allows you to deploy the following NLB cluster configurations in your network:

- NLB clusters that include two or more Containers residing on the same Hardware Node.
- NLB clusters that include two or more Containers residing on different Hardware Nodes.
- NLB clusters that include one or more Containers and physical servers.

The NLB clusters containing a number of Containers provide you with the same functionality as the clusters uniting standalone servers only. For example, you can use the created NLB cluster to balance incoming web-based traffic by splitting its load among several Containers to improve response times and reliability, add new Containers to the cluster and remove existing ones, etc. However, the procedure of creating an NLB cluster in Parallels Containers-based systems slightly differs from that of setting up an NLB cluster containing dedicated physical hosts only and is described below.

The following example demonstrates how to create an NLB cluster containing two Containers - Container 101 and Container 102. So, if you experience the failure of, say, Container 101 in the NLB cluster, you can still serve user requests on the remaining NLB cluster member - Container 102. To create such a cluster, you should perform the following operations:

Note: If Containers (or physical servers) you plan to include in the cluster are running Windows Server 2008 or Windows Server 2008 R2, install the Network Load Balancing feature on each of these servers before starting the cluster creation.

- 1 Make sure that Container 101 and Container 102 are operating in the bridged mode and bound to one and the same physical adapter on the Hardware Node. If they are not, refer to the **Connecting Bridged Containers to Network** section of the *Parallels Containers 4.6 User's Guide* for the information on how you can do it.

Note: All participants of an NLB cluster, irrespective of whether it is a Container or a physical server, should belong to the same broadcast domain. For Containers residing on the same Hardware Node, this can be done by simply configuring their network adapters to operate in the bridged mode. If Containers reside on different Nodes, you should additionally ensure that the physical adapters where the corresponding Container virtual network adapters are bridged belong to the same broadcast domain.

- 2 Decide on the Container operation mode in your NLB cluster. By default, all Containers are set to function in the unicast mode. To change the Container default mode to multicast, proceed as follows:

- To set the multicast mode for Container 101, run the following command:

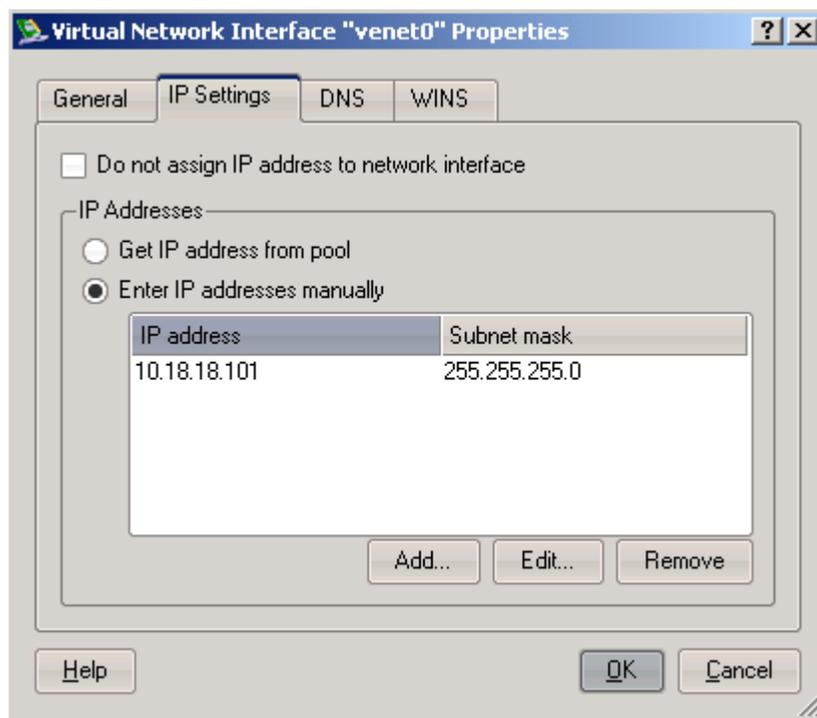
```
C:\...\Administrator>vzctl set 101 --nlb_mode multicast
```

- To set the multicast mode for Container 102, run the following command:

```
C:\...\Administrator>vzctl set 102 --nlb_mode multicast
```

Note: If your NLB cluster is operating in the multicast mode and client computers access it through a router, make sure that the router can map the cluster IP address to the MAC address associated with this IP address. If the router cannot, create the corresponding static ARP entry in the router. For more information about NLB clusters operating in the multicast mode, refer to the *Network Load Balancing parameters: Cluster Operation Mode* document available at <http://technet2.microsoft.com/windowsserver/en/library/57c24429-0268-4ed8-afdf-fd4b0b6539b71033.mspx?mfr=true>.

- 3 Assign valid IP addresses to Container 101 and Container 102, if you have not already done so before. Let Container 101 have the IP address of 192.168.1.156 and Container 102 the IP address of 192.168.1.157. For example, you can do it by completing the following operations in Parallels Management Console:
 - Select the Parallels Containers item under the corresponding Hardware Node name.
 - Right-click Container 101, and choose Properties.
 - On the Network tab of the displayed window, select the Network Adapters item.
 - In the Interfaces table, select the Container virtual adapter, and click the Properties button.
 - Go to the IP Settings tab of the Virtual Network Interface Properties window. For example:



- Select the Enter IP addresses manually radio button, click the Add button, and type 192.168.1.156 in the IP address field of the displayed window. You can also configure the subnet mask to be assigned to Container 101 by providing the necessary value in the Subnet mask field.
- Click OK.

- Perform the same operations for Container 102 and assign the IP address of 192.168.1.157 to it.

4 Add a 'virtual' (Cluster) IP address to Container 101 and Container 102. Let this 'virtual' IP address be 192.168.1.200:

- Add the cluster IP address to Container 101:

```
C:\...\Administrator>vzctl set 101 --cluster_ip 192.168.1.200 --save  
Command 'set' is successfully finished
```

- Add the cluster IP address to Container 102:

```
C:\...\Administrator>vzctl set 102 --cluster_ip 192.168.1.200 --save  
Command 'set' is successfully finished
```

5 Start the NLB service inside Container 101 and 102:

- For Container 101, issue the following command:

```
C:\...\Administrator>vzctl set 101 --nlb yes --save  
Command 'set' is successfully finished
```

- For Container 102, execute the following command:

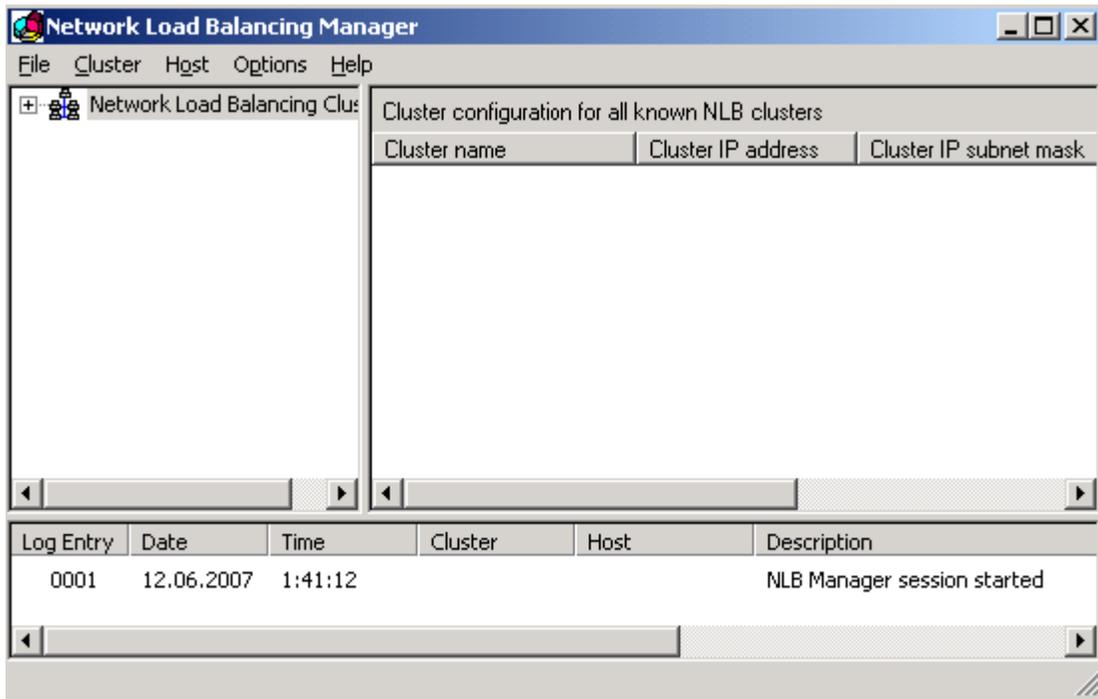
```
C:\...\Administrator>vzctl set 102 --nlb yes --save  
Command 'set' is successfully finished
```

Note: If you are planning to include more than 2 Containers in your NLB cluster, you should execute Steps 2 and 4 for each Container to be joined to the cluster.

6 Start Container 101 and Container 102 (e.g., using Parallels Management Console or the `vzctl start` command).

7 Log in to Container 101 via RDP.

8 Inside Container 101, select **Start > Run** and execute the `nlbmgr.exe` command to open the Network Load Balancing (NLB) Manager.



- 9 Create an NLB cluster using Network Load Balancing Manager and register two hosts (Container 101 and Container 102) in this cluster in the same way you would do it on any other standalone server.

After the second node has been successfully added to the NLB cluster, you can start using the cluster in the same way you would use any other NLB cluster containing standalone computers only.

Index

A

- Active/Active Failover Clusters - 9
- Active/Passive Failover Clusters - 7
- Adding a Disk Resource - 49
- Adding a Network Name Resource - 55
- Adding an IP Address Resource - 53
- Adding IP Address and Network Name Resources - 53
- Adding Nodes to the Cluster - 29, 43
- Adding Nodes to the Domain - 20
- Adding Nodes to the Failover Cluster - 66

C

- Checking Cluster Operation - 35
- Checking Server Cluster Operation - 58
- Configuring Automatic Failback - 65
- Configuring Physical Servers - 12
- Configuring the Cluster - 31, 46
- Connecting Servers to a Shared Storage - 22
- Creating a Cluster in Windows Server 2008 and Windows Server 2008 R2 - 26
- Creating a Cluster Service Account - 23
- Creating a Domain - 13
- Creating Cluster Resource Groups - 47
- Creating Clusters From Containers - 72
- Creating the Cluster and Configuring the First Cluster Node - 38

D

- Deploying a Failover Cluster in Windows Server 2003 - 38
- Deploying a Failover Cluster in Windows Server 2008 and Windows Server 2008 R2 - 25
- Deploying Parallels Containers Failover Clusters - 6
- Designing a Cluster Configuration - 7

F

- Failover Clusters - 5

I

- Installing Parallels Containers on Cluster Nodes - 60

M

- Microsoft Cluster Types - 5

N

- NLB Clusters - 5

S

- Setting Up a Cluster - 23
- Setting Up a Domain - 12
- Setting Up a 'Loopback File'-Based Cluster - 74
- Setting Up a SAN-Based Cluster - 72
- Setting Up an NLB Cluster - 78
- Setting Up Hardware for the Cluster - 11

T

- Testing the Failover Cluster - 67
- Troubleshooting Failover Clusters - 69

V

- Validating Hardware - 25