

Setup for Failover Clustering and Microsoft Cluster Service

ESXi 5.0
vCenter Server 5.0

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-000628-01

vmware[®]

You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2006–2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

Setup for Failover Clustering and Microsoft Cluster Service	5
Updated Information	5
Getting Started with MSCS	5
Cluster Virtual Machines on One Physical Host	10
Cluster Virtual Machines Across Physical Hosts	13
Cluster Physical and Virtual Machines	17
Use MSCS in an vSphere HA and vSphere DRS Environment	20
vSphere MSCS Setup Checklist	23
Index	27

Setup for Failover Clustering and Microsoft Cluster Service

Setup for Failover Clustering and Microsoft Cluster Service describes the types of clusters you can implement using virtual machines with Microsoft Cluster Service for Windows Server 2003 and Failover Clustering for Windows Server 2008. You get step-by-step instructions for each type of cluster and a checklist of clustering requirements and recommendations.

Unless stated otherwise, the term Microsoft Cluster Service (MSCS) applies to Microsoft Cluster Service with Windows Server 2003 and Failover Clustering with Windows Server 2008.

Setup for Failover Clustering and Microsoft Cluster Service covers ESXi and VMware® vCenter® Server.

Intended Audience

This information is for system administrators who are familiar with VMware technology and Microsoft Cluster Service.

NOTE This is not a guide to using Microsoft Cluster Service or Failover Clustering. Use your Microsoft documentation for information about installation and configuration of Microsoft Cluster Service or Failover Clustering.

Updated Information

This *Setup for Failover Clustering and Microsoft Cluster Service* documentation is updated with each release of the product or when necessary.

This table provides the update history of the *Setup for Failover Clustering and Microsoft Cluster Service* documentation.

Revision	Description
EN-000628-01	■ Added information to “vSphere MSCS Setup Limitations,” on page 9 about hosts that use memory overcommitment and their lack of suitability for MSCS virtual machines.
EN-000628-00	Initial release.

Getting Started with MSCS

VMware® vSphere® supports clustering using MSCS across virtual machines. Clustering virtual machines can reduce the hardware costs of traditional high-availability clusters.

NOTE vSphere High Availability (vSphere HA) supports a clustering solution in conjunction with vCenter Server clusters. *vSphere Availability* describes vSphere HA functionality.

Clustering Configuration Overview

Several applications use clustering, including stateless applications such as Web servers, and applications with built-in recovery features such as database servers. You can set up MSCS clusters in several configurations, depending on your environment.

A typical clustering setup includes:

- Disks that are shared between nodes. A shared disk is required as a quorum disk. In a cluster of virtual machines across physical hosts, the shared disk must be on a Fibre Channel (FC) SAN.
- A private heartbeat network between nodes.

You can set up the shared disks and private heartbeat using one of several clustering configurations.

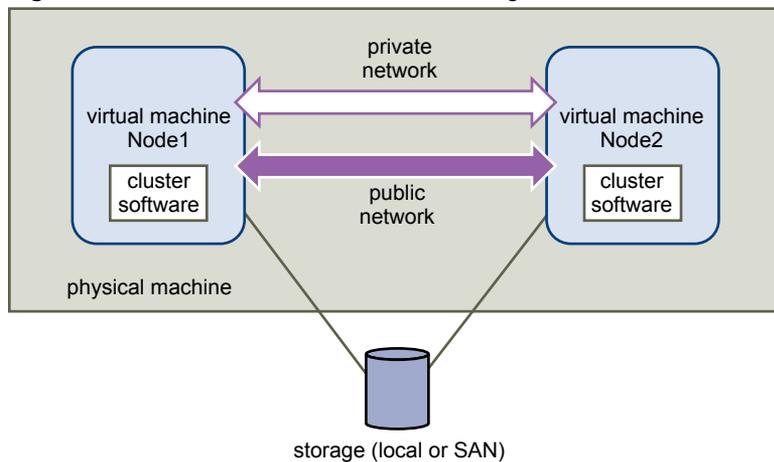
Clustering MSCS Virtual Machines on a Single Host

A cluster of MSCS virtual machines on a single host (also known as a cluster in a box) consists of two clustered virtual machines on the same ESXi host. The virtual machines are connected to the same storage, either local or remote. This configuration protects against failures at the operating system and application level, but it does not protect against hardware failures.

The following figure shows a cluster in a box setup.

- Two virtual machines on the same physical machine (ESXi host) run clustering software.
- The virtual machines share a private network connection for the private heartbeat and a public network connection.
- Each virtual machine is connected to shared storage, which can be local or on a SAN.

Figure 1. Virtual Machines Clustered on a Single Host

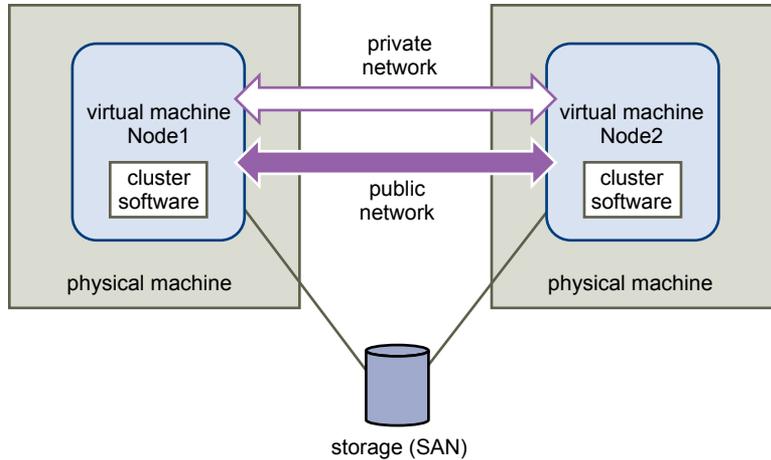


Clustering Virtual Machines Across Physical Hosts

A cluster of virtual machines across physical hosts (also known as a cluster across boxes) protects against software failures and hardware failures on the physical machine by placing the cluster nodes on separate ESXi hosts. This configuration requires shared storage on an Fibre Channel SAN for the quorum disk.

The following figure shows a cluster-across-boxes setup.

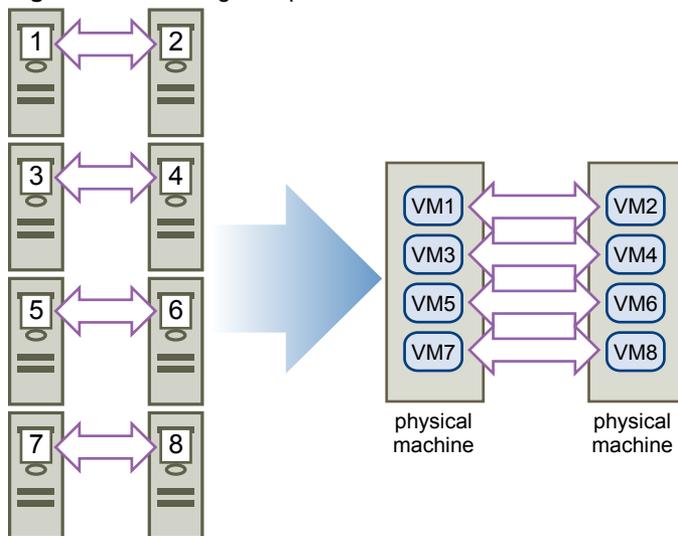
- Two virtual machines on two different physical machines (ESXi hosts) run clustering software.
- The virtual machines share a private network connection for the private heartbeat and a public network connection.
- Each virtual machine is connected to shared storage, which must be on a SAN.

Figure 2. Virtual Machines Clustered Across Hosts

This setup provides significant hardware cost savings.

You can expand the cluster-across-boxes model and place multiple virtual machines on multiple physical machines. For example, you can consolidate four clusters of two physical machines each to two physical machines with four virtual machines each.

The following figure shows how you can move four two-node clusters from eight physical machines to two.

Figure 3. Clustering Multiple Virtual Machines Across Hosts

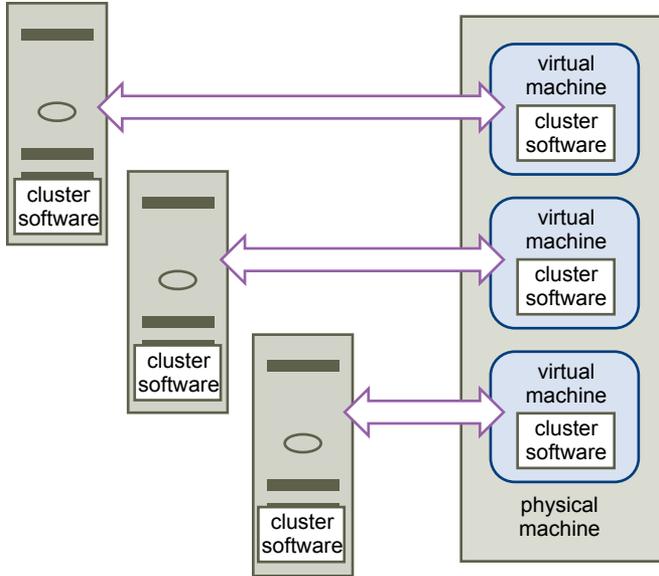
Clustering Physical Machines with Virtual Machines

For a simple MSCS clustering solution with low hardware requirements, you might choose to have one standby host.

Set up your system to have a virtual machine corresponding to each physical machine on the standby host, and create clusters, one each for each physical machine and its corresponding virtual machine. In case of hardware failure in one of the physical machines, the virtual machine on the standby host can take over for that physical host.

The following figure shows a standby host using three virtual machines on a single physical machine. Each virtual machine is running clustering software.

Figure 4. Clustering Physical and Virtual Machines



Hardware and Software Requirements for Clustering

All vSphere MSCS configurations require certain hardware and software components.

The following table lists hardware and software requirements that apply to all vSphere MSCS configurations.

Table 1. Clustering Requirements

Component	Requirement
Virtual SCSI adapter	LSI Logic Parallel for Windows Server 2003 LSI Logic SAS for Windows Server 2008
Operating system	Windows Server 2003 SP2 or Windows Server 2008 R2
Virtual NIC	Use the default type for all guest operating systems.
I/O timeout	Set to 60 seconds or more. Modify HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue. The system might reset this I/O timeout value if you re-create a cluster. You must reset the value in that case.
Disk format	Select Thick Provision to create disks in <code>eagerzeroedthick</code> format.
Disk and networking setup	Add networking before disks. Refer to the VMware knowledge base article at http://kb.vmware.com/kb/1513 if you encounter any errors.
Number of nodes	Use two-node clustering.
NTP server	Synchronize domain controllers and cluster nodes with a common NTP server, and disable host-based time synchronization when using clustering in the guest.

Supported Shared Storage Configurations

Different MSCS cluster setups support different types of shared storage configurations. Some setups support more than one type. Select the recommended type of shared storage for best results.

Table 2. Shared Storage Requirements

Storage Type	Clusters on One Physical Machine (Cluster in a Box)	Clusters Across Physical Machines (Cluster Across Boxes)	Clusters of Physical and Virtual Machines (Standby Host Clustering)
Virtual disks	Yes (recommended)	No	No
Pass-through RDM (physical compatibility mode)	No	Yes (recommended)	Yes
Non-pass-through RDM (virtual compatibility mode)	Yes	Yes	No

Use of software iSCSI initiators within guest operating systems configured with MSCS, in any configuration supported by Microsoft, is transparent to ESXi hosts and there is no need for explicit support statements from VMware.

NOTE Clusters across physical machines with non-pass-through RDM is supported only for clustering with Windows Server 2003. It is not supported for clustering with Windows Server 2008.

vSphere MSCS Setup Limitations

Before you set up MSCS, review the list of functions that are not supported for this release, and requirements and recommendations that apply to your configuration.

The following environments and functions are not supported for MSCS setups with this release of vSphere:

- Clustering on iSCSI, FCoE, and NFS disks.
- Mixed environments, such as configurations where one cluster node is running a different version of ESXi than another cluster node.
- Use of MSCS in conjunction with vSphere Fault Tolerance (FT).
- Migration with vSphere vMotion[®] of clustered virtual machines.
- N-Port ID Virtualization (NPIV)
- With native multipathing (NMP), clustering is not supported when the path policy is set to round robin. Third-party multipathing plug-ins might support round robin or other load balancing behavior with Microsoft clusters. Support of third-party multipathing plug-ins is provided by the plug-in vendor.
- ESXi hosts that use memory overcommitment are not suitable for deploying MSCS virtual machines. Memory overcommitment can cause virtual machines to stall for short durations. This can be significantly disruptive as the MSCS clustering mechanism is time-sensitive and timing delays can cause the virtual machines to behave incorrectly.

MSCS and Booting from a SAN

You can put the boot disk of a virtual machine on a SAN-based VMFS volume.

Booting from a SAN is complex. Problems that you encounter in physical environments extend to virtual environments. For general information about booting from a SAN, see the *vSphere Storage* documentation.

Follow these guidelines when you place the boot disk of a virtual machine on a SAN-based VMFS volume:

- Consider the best practices for boot-from-SAN that Microsoft publishes in the following knowledge base article: <http://support.microsoft.com/kb/305547/en-us>.
- Use StorPort LSI Logic drivers instead of SCSIport drivers when running Microsoft Cluster Service for Windows Server 2003 or 2008 guest operating systems.
- Test clustered configurations in different failover scenarios before you put them into production environments.

Setting up Clustered Continuous Replication or Database Availability Groups with Exchange 2010

You can set up Clustered Continuous Replication (CCR) with Exchange 2007 or Database Availability Groups (DAG) with Exchange 2010 in your vSphere environment.

When working in a vSphere environment:

- Use virtual machines instead of physical machines as the cluster components.
- If the boot disks of the CCR or DAG virtual machines are on a SAN, see “[MSCS and Booting from a SAN](#),” on page 9.

For more information, see Microsoft’s documentation for CCR or DAG on the Microsoft Web site.

Cluster Virtual Machines on One Physical Host

You can create a two-node MSCS cluster on a single ESXi host.

A cluster of virtual machines on one physical machine requires a host with one physical network adapter for the VMkernel. Use a separate physical network adapter for clustered virtual machines to connect with external hosts.

Create the First Node for Clusters on One Physical Host

To create the first node, you create and configure a virtual machine with two virtual network adapters and install a guest operating system on the virtual machine.

You configure the virtual network adapters to handle virtual machine traffic for the cluster: a private network connection for the private heartbeat and a public network connection.

Procedure

- 1 Open the vSphere Client and connect to the ESXi host or a vCenter Server system.
Use the user name and password of the user who will have administrator permissions the virtual machine.
- 2 In the vSphere Client inventory, right-click the host and select **New Virtual Machine**.
- 3 Proceed through the wizard to create the virtual machine.

Page	Action
Configuration	Select Typical .
Name and Location	Enter a name and select a location.
Datastore	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Guest Operating System	Select the Windows Server operating system that you intend to install.

Page	Action
Network Connections	Select the number of NICs to connect (1), the adapter type, and the network label. The network label refers to the network to be used for virtual machine traffic for the cluster, public or private.
Create Disk	Enter the appropriate value in Virtual disk size . Select Thick Provision to create a disk in <code>eagerzeroedthick</code> format.

- Click **Finish** to complete creating the virtual machine.

NOTE Do not add shared cluster disks at this time.

- In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.
The Virtual Machine Properties dialog box appears.
- On the Hardware tab of the Virtual Machine Properties dialog box, click **Add**.
- From the list of devices, select **Ethernet Adapter** and click **Next**.
- Select the adapter type and network label.
 - If you selected a private network for the first network adapter, you must select a public network for this network adapter.
 - If you selected a public network for the first network adapter, you must select a private network adapter.
- Click **Next**, and click **Finish** to complete creating the device.
- Install a Windows Server operating system on the virtual machine.

Create the Second Node for Clusters on One Physical Host

Create a template from the first virtual machine and deploy the second node from that template.



CAUTION If you clone a virtual machine with an RDM setup, the cloning process converts the RDMs to virtual disks. Unmap all RDMs before cloning, and remap them after cloning is complete.

Procedure

- In the vSphere Client inventory, right-click the first virtual machine you created and select **Template > Clone to Template**.
- Proceed through the wizard to create the virtual machine template.

Page	Action
Name and Location	Enter a name (for example, Node2_Template) and select a location.
Host / Cluster	Select the host or cluster where you will run the virtual machine.
Datastore	Select a datastore as the location for the virtual machine configuration file and the <code>.vmdk</code> file.
Disk Format	Select Same format as source .
Ready to Complete	Click Finish to create the virtual machine template.

- Right-click the virtual machine template and select **Deploy Virtual Machine from this Template**.
- Proceed through the deployment wizard to deploy the virtual machine.

Page	Action
Name and Location	Enter a name (for example, Node2) and select a location.
Host / Cluster	Select the host or cluster where you will run the virtual machine.

Page	Action
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Disk Format	Select Same format as source .
Customization	Select Customize using the Customization Wizard .

- 5 Proceed through the guest operating system Customization Wizard.
 - a On the Operating Systems Options page, select **Generate New Security ID (SID)** to generate a new security identity.
 - b Click **Finish** to exit the Customization Wizard.
- 6 Click **Finish** to deploy the virtual machine.

Add Hard Disks to the First Node for Clusters on One Physical Host

In an MSCS cluster, storage disks are shared between nodes. You set up a quorum disk and an optional shared storage disk.

Procedure

- 1 In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.
The Virtual Machine Properties dialog box appears.
- 2 Click **Add**, select **Hard Disk**, and click **Next**.
- 3 Select **Create a new virtual disk** and click **Next**.
- 4 Select the disk size.
- 5 Under Disk Provisioning, select **Thick Provision**.
You can also use a mapped SAN LUN set to virtual compatibility mode.
- 6 Click **Next**.
- 7 From the **Virtual Device Node** drop-down menu, select a new SCSI controller (for example, **SCSI (1:0)**).

NOTE You must select a new virtual device node. You cannot use SCSI 0.

- 8 Click **Next**, and click **Finish**.
The wizard creates a new hard disk and a new SCSI controller.
- 9 In the Virtual Machine Properties dialog box, select the new SCSI controller and click **Change Type**.
The Change SCSI Controller Type dialog box appears.
- 10 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003	LSI Logic Parallel
Windows Server 2008	LSI Logic SAS

- 11 Click **OK**.
- 12 On the Virtual Machine Properties dialog box, set SCSI Bus Sharing to **Virtual** and click **OK**.

Add Hard Disks to the Second Node for Clusters on One Physical Host

To allow shared access to clustered services and data, point the quorum disk of the second node to the same location as the first node's quorum disk. Point shared storage disks to the same location as the first node's shared storage disks.

Prerequisites

Before you begin, obtain the following information:

- Which virtual device node is for the first virtual machine's shared storage disks (for example, SCSI (1:0)).
- The location of the quorum disk specified for the first node.

Procedure

- 1 In the vSphere Client inventory, select the second virtual machine that you created and select **Edit Settings**.

The Virtual Machine Properties dialog box appears.

- 2 Click **Add**, select **Hard Disk**, and click **Next**.
- 3 Select **Use an existing virtual disk** and click **Next**.
- 4 Select the same virtual device node you chose for the first virtual machine's shared storage disks (for example, **SCSI (1:0)**), and click **Next**.

NOTE The location of the virtual device node for this virtual machine's shared storage must match the corresponding virtual device node for the first virtual machine.

- 5 In Disk File Path, browse to the location of the quorum disk specified for the first node.

Cluster Virtual Machines Across Physical Hosts

You can create a MSCS cluster that consists of two virtual machines on two ESXi hosts.

A cluster across physical hosts requires specific hardware and software.

- ESXi hosts that have the following:
 - Two physical network adapters dedicated to the MSCS cluster and to the public and private networks.
 - One physical network adapter dedicated to the VMkernel.
- Fibre Channel (FC) SAN. Shared storage must be on an FC SAN.
- RDM in physical compatibility (pass-through) or virtual compatibility (non-pass-through) mode. VMware recommends physical compatibility mode. The cluster cannot use virtual disks for shared storage.

Failover clustering with Windows Server 2008 is not supported with virtual compatibility mode (non-pass-through) RDMs.

Create the First Node for MSCS Clusters Across Physical Hosts

To create the first node, you create and configure a virtual machine with two virtual network adapters and install a guest operating system on the virtual machine.

You configure the virtual network adapters to handle virtual machine traffic for the cluster: a private network connection for the private heartbeat and a public network connection.

Procedure

- 1 Open the vSphere Client and connect to the ESXi host or a vCenter Server system.
Use the user name and password of the user who will have administrator permissions the virtual machine.
- 2 In the vSphere Client inventory, right-click the host and select **New Virtual Machine**.
- 3 Proceed through the wizard to create the virtual machine.

Page	Action
Configuration	Select Typical .
Name and Location	Enter a name and select a location.
Datastore	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Guest Operating System	Select the Windows Server operating system that you intend to install.
Network Connections	Select the number of NICs to connect (1), the adapter type, and the network label. The network label refers to the network to be used for virtual machine traffic for the cluster, public or private.
Create Disk	Enter the appropriate value in Virtual disk size . Select Thick Provision to create a disk in eagerzeroedthick format.

- 4 Click **Finish** to complete creating the virtual machine.

NOTE Do not add shared cluster disks at this time.

- 5 In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.
The Virtual Machine Properties dialog box appears.
- 6 On the Hardware tab of the Virtual Machine Properties dialog box, click **Add**.
- 7 From the list of devices, select **Ethernet Adapter** and click **Next**.
- 8 Select the adapter type and network label.
 - If you selected a private network for the first network adapter, you must select a public network for this network adapter.
 - If you selected a public network for the first network adapter, you must select a private network adapter.
- 9 Click **Next**, and click **Finish** to complete creating the device.
- 10 Install a Windows Server operating system on the virtual machine.

Create the Second Node for Clusters Across Physical Hosts

To create the second node in a cluster of virtual machines across physical hosts, you create a template of the first virtual machine and use it to deploy a second virtual machine onto a second ESXi host.



CAUTION If you clone a virtual machine with an RDM setup, the cloning process converts the RDMs to virtual disks. Unmap all RDMs before cloning, and remap them after cloning is complete.

Procedure

- 1 In the vSphere Client inventory, right-click the first virtual machine you created and select **Template > Clone to Template**.

- 2 Proceed through the wizard to create the virtual machine template.

Page	Action
Name and Location	Enter a name and select a location.
Host or Cluster	Select the second host for the cluster setup.
Resource Partition	Select the resource pool for the virtual machine, or select the host if no resource pools exist.
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Ready to Complete	Click OK to create the virtual machine template.

- 3 Right-click the virtual machine template and select **Deploy Virtual Machine from this Template**.

- 4 Proceed through the deployment wizard to deploy the virtual machine.

Page	Action
Name and Location	Enter a name (for example, Node2) and select a location.
Host / Cluster	Select the host or cluster where you will run the virtual machine.
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Disk Format	Select Same format as source .
Customization	Select Customize using the Customization Wizard .

- 5 Proceed through the guest operating system Customization Wizard.

- a On the Operating Systems Options page, select **Generate New Security ID (SID)** to generate a new security identity.
- b Click **Finish** to exit the Customization Wizard.

- 6 Click **Finish** to deploy the virtual machine.

Add Hard Disks to the First Node for Clusters Across Physical Hosts

In an MSCS cluster, storage disks are shared between nodes. You set up a quorum disk and an optional shared storage disk.

Prerequisites

Before you add hard disks to the first node, complete the following tasks:

- For each virtual machine, configure the guest operating system's private and public IP addresses.
- Ask your SAN administrator for the location of unformatted SAN LUNs. The hard disks you create in this task must point to SAN LUNs.

NOTE Use RDMs in physical compatibility mode. The procedure below uses physical compatibility mode.

Procedure

- 1 In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.
The Virtual Machine Properties dialog box appears.
- 2 Click **Add**, select **Hard Disk**, and click **Next**.
- 3 Select **Raw Device Mappings** and click **Next**.
- 4 Select an unformatted LUN and click **Next**.

- 5 Select a datastore and click **Next**.

This datastore must be on a SAN because you need a single shared RDM file for each shared LUN on the SAN.

- 6 Select **Physical** as the compatibility mode and click **Next**.

A SCSI controller is created when the virtual hard disk is created.

- 7 Select a new virtual device node (for example, select **SCSI (1:0)**), and click **Next**.

NOTE This must be a new SCSI controller. You cannot use SCSI 0.

- 8 Click **Finish** to complete creating the disk.

The wizard creates a new SCSI controller and a new hard disk.

- 9 In the Virtual Machine Properties dialog box, select the new SCSI controller and click **Change Type**.

The Change SCSI Controller Type dialog box appears.

- 10 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003	LSI Logic Parallel
Windows Server 2008	LSI Logic SAS

- 11 Click **OK**.

- 12 On the Virtual Machine Properties dialog box, set SCSI Bus Sharing to **Physical** and click **OK**.

The virtual machine is connected to a public network and a private network with two virtual switches, and is connected to the quorum disk on FC SAN and the virtual machine virtual disk on local or remote storage.

Add Hard Disks to the Second Node for Clusters Across Physical Hosts

To allow shared access to clustered services and data, point the quorum disk of the second node to the same location as the first node's quorum disk. Point shared storage disks to the same location as the first node's shared storage disks.

Prerequisites

Before you begin, obtain the following information:

- Which virtual device node is for the first virtual machine's shared storage disks (for example, SCSI (1:0)).
- The location of the quorum disk specified for the first node.

Procedure

- 1 In the vSphere Client inventory, select the second virtual machine that you created and select **Edit Settings**.

The Virtual Machine Properties dialog box appears.

- 2 Click **Add**, select **Hard Disk**, and click **Next**.

- 3 Select **Use an existing virtual disk** and click **Next**.

- 4 In Disk File Path, browse to the location of the quorum disk specified for the first node.

- 5 Select **Physical** as the compatibility mode and click **Next**.

A SCSI controller is created when the virtual hard disk is created.

- 6 Select the same virtual device node you chose for the first virtual machine's shared storage disks (for example, **SCSI (1:0)**), and click **Next**.

NOTE The location of the virtual device node for this virtual machine's shared storage must match the corresponding virtual device node for the first virtual machine.

- 7 Click **Finish**.
The wizard creates a new hard disk and a new SCSI controller.
- 8 In the Virtual Machine Properties dialog box, select the new SCSI controller and click **Change Type**.
The Change SCSI Controller Type dialog box appears.
- 9 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003	LSI Logic Parallel
Windows Server 2008	LSI Logic SAS

- 10 Click **OK**.
- 11 Set SCSI Bus Sharing to **Physical** and click **OK**.

Cluster Physical and Virtual Machines

You can create an MSCS cluster in which each physical machine has a corresponding virtual machine. This type of configuration is known as a standby host cluster.

A standby host cluster has specific hardware and software requirements.

- Use ESXi hosts that have the following:
 - Two physical network adapters dedicated to the MSCS cluster and to the public and private networks.
 - One physical network adapter dedicated to the VMkernel.
- Use RDMS in physical compatibility mode (pass-through RDM). You cannot use virtual disks or RDMS in virtual compatibility mode (non-pass-through RDM) for shared storage.
- Use the STORport Miniport driver for the Fibre Channel (FC) HBA (QLogic or Emulex) in the physical Windows machine.
- Do not run multipathing software in the physical or virtual machines.
- Use only a single physical path from the host to the storage arrays in standby host configurations.

Create the First Node for a Cluster of Physical and Virtual Machines

The first node in a standby host setup is a physical machine.

For information about setting up a physical machine that participates in an MSCS cluster, see the Microsoft Cluster Service documentation.

Procedure

- ◆ Set up the physical machine using the settings listed in the table.

Component	Requirement
Windows Cluster Administrator application	Advanced minimum configuration if Windows 2003 is used.
Network adapters	At least two.

Component	Requirement
Storage	Access to the same storage on a SAN as the ESXi host on which the corresponding virtual machine will run.
Operating system	Installed on each physical machine.

Create the Second Node for a Cluster of Physical and Virtual Machines

To create the second node, you set up a virtual machine for clustering across physical machines.

You configure the virtual network adapters to handle virtual machine traffic for the cluster: a private network connection for the private heartbeat and a public network connection.

Prerequisites

Before you begin, make sure that the shared storage that is visible from the physical machine that you configured in [“Create the First Node for a Cluster of Physical and Virtual Machines,”](#) on page 17 is also visible from the virtual machine.

Procedure

- 1 Open the vSphere Client and connect to the ESXi host.

Use the user name and password of the user who will have administrator permissions the virtual machine.

- 2 In the vSphere Client inventory, right-click the host and select **New Virtual Machine**.
- 3 Proceed through the wizard to create the virtual machine.

Page	Action
Configuration	Select Typical .
Name and Location	Enter a name and select a location.
Datastore	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Guest Operating System	Select the Windows Server operating system that you intend to install.
Network Connections	Select the number of NICs to connect (1), the adapter type, and the network label. The network label refers to the network to be used for virtual machine traffic for the cluster, public or private.
Create Disk	Enter the appropriate value in Virtual disk size . Select Thick Provision to create a disk in <code>eagerzeroedthick</code> format.

- 4 Click **Finish** to complete creating the virtual machine.
- 5 In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.
The Virtual Machine Properties dialog box appears.
- 6 On the Hardware tab of the Virtual Machine Properties dialog box, click **Add**.
- 7 From the list of devices, select **Ethernet Adapter** and click **Next**.
- 8 Select the adapter type and network label.
 - If you selected a private network for the first network adapter, you must select a public network for this network adapter.
 - If you selected a public network for the first network adapter, you must select a private network adapter.
- 9 Click **Next**, and click **Finish** to complete creating the device.
- 10 Install a Windows Server operating system on the virtual machine.

Add Hard Disks to the Second Node for a Cluster of Physical and Virtual Machines

When you add hard disks to the second node, you set up the disks to point to the quorum disk and shared storage disks, if any, for the first node. The setup allows shared access to clustered services and data.

Procedure

- 1 In the vSphere Client inventory, select the newly created virtual machine and select **Edit Settings**.

The Virtual Machine Properties dialog box appears.

- 2 Click **Add**, select **Hard Disk**, and click **Next**.
- 3 Select **Raw Device Mappings** and click **Next**.
- 4 Select the LUN that is used by the physical machine.
- 5 Select the datastore, which is also the location of the boot disk, and click **Next**.

- 6 Select **Physical** as the compatibility mode and click **Next**.

A SCSI controller is created when the virtual hard disk is created.

- 7 From the **Virtual Device Node** drop-down menu, select a new SCSI controller (for example, **SCSI (1:0)**).

NOTE You must select a new virtual device node. You cannot use SCSI 0.

- 8 Click **Next**, and click **Finish**.

The wizard creates a new hard disk and a new SCSI controller.

- 9 In the Virtual Machine Properties dialog box, select the new SCSI controller and click **Change Type**.

The Change SCSI Controller Type dialog box appears.

- 10 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003	LSI Logic Parallel
Windows Server 2008	LSI Logic SAS

- 11 Click **OK**.

- 12 On the Virtual Machine Properties dialog box, set SCSI Bus Sharing to **Virtual** and click **OK**.

Install Microsoft Cluster Service

For Windows Server 2003 operating systems only, after you set up the first and second nodes, you must configure Microsoft Cluster Service.

See Microsoft's documentation for creating and configuring server clusters on the Microsoft Web site.

In complex storage solutions, such as an FC switched fabric, a storage unit might have a different identity (target ID or raw disk ID) on each computer in the cluster. Although this is a valid storage configuration, it causes a problem when you add a node to the cluster.

The following procedure allows you to avoid target identity problems when using clustering with Windows 2003.

Procedure

- 1 Within the Microsoft Cluster Administrator utility, on the Select Computer page, click **Advanced** to disable the storage validation heuristics.

- 2 Select the **Advanced** (minimum) configuration option, and click **OK**.

Microsoft Cluster Service should operate normally in the virtual machine after it is installed.

Create Additional Physical-Virtual Pairs

If you have additional physical machines, you can create additional clusters for each.

Procedure

- 1 On the ESXi host, set up an additional virtual machine for the physical machine.
- 2 Cluster the physical machine with the new virtual machine.

Use MSCS in an vSphere HA and vSphere DRS Environment

When you use MSCS in a vSphere High Availability (vSphere HA) or vSphere Distributed Resource Scheduler (DRS) environment, you must configure your hosts and virtual machines to use certain settings. All hosts that run MSCS virtual machines must be managed by a vCenter Server system.

Enable vSphere HA and vSphere DRS in a vCenter Server Cluster (MSCS)

All hosts that are running MSCS virtual machines can be part of a vCenter Server cluster with both vSphere HA and vSphere DRS enabled. You can enable vSphere HA and vSphere DRS in the Cluster Settings dialog box.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the Cluster Settings dialog box, select **Cluster Features**.
- 3 Select the **Turn on vSphere HA** and **Turn on vSphere DRS** check boxes.
- 4 Click **OK**.

Create VM-VM Affinity Rules for MSCS Virtual Machines

For MSCS virtual machines in a cluster, you must create VM-VM affinity or anti-affinity rules. VM-VM affinity rules specify which virtual machines should be kept together on the same host (for example, a cluster of MSCS virtual machines on one physical host). VM-VM anti-affinity rules specify which virtual machines should be kept apart on different physical hosts (for example, a cluster of MSCS virtual machines across physical hosts).

For a cluster of virtual machines on one physical host, use affinity rules. For a cluster of virtual machines across physical hosts, use anti-affinity rules.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the Cluster Settings dialog box under vSphere DRS, select **Rules**.
- 3 Click **Add**.
- 4 In the Rule dialog box, type a name for the rule.
- 5 From the **Type** drop-down menu, select a rule.
 - For a cluster of virtual machines on one physical host, select **Keep Virtual Machines Together**.
 - For a cluster of virtual machines across physical hosts, select **Separate Virtual Machines**.
- 6 Click **Add**.
- 7 Select the two virtual machines to which the rule applies and click **OK**.

- 8 Click **OK**.

Enable Strict Enforcement of Affinity Rules (MSCS)

To ensure that affinity and anti-affinity rules are strictly applied, set an advanced option for vSphere DRS. Setting the advanced option `ForceAffinePoweron` to 1 will enable strict enforcement of the affinity and anti-affinity rules that you created.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the Cluster Settings dialog box, select **vSphere DRS**.
- 3 Click **Advanced Options**.
- 4 In the Option column, type **ForceAffinePoweron**.
- 5 In the Value column, type **1**.
- 6 Click **OK**.

Set DRS Automation Level for MSCS Virtual Machines

You must set the automation level of all virtual machines in an MSCS cluster to Partially Automated. When you set the vSphere DRS automation level for the virtual machine to Partially Automated, vCenter Server will perform initial placement of virtual machines when they are powered on and will provide migration recommendations for them.

NOTE Migration of MSCS clustered virtual machines is not recommended.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane under vSphere DRS, select **Virtual Machine Options**.
- 3 Select the **Enable individual virtual machine automation levels** check box.
- 4 Change the virtual machine automation level for each MSCS virtual machine in the cluster.
 - a In the Virtual Machine column, select the virtual machine.
 - b In the Automation Level column, select **Partially Automated** from the drop-down menu.
 - c Repeat [Step 4a](#) and [Step 4b](#) for each virtual machine.
- 5 Click **OK**.

Using vSphere DRS Groups and VM-Host Affinity Rules with MSCS Virtual Machines

You can use the vSphere Client to set up two types of DRS groups: virtual machine DRS groups, which contain at least one virtual machine, and host DRS groups, which contain at least one host. A VM-Host affinity rule establishes an affinity (or anti-affinity) relationship between a virtual machine DRS group and a host DRS group.

You must use VM-Host affinity rules because vSphere HA does not obey VM-VM affinity rules. This means that if a host fails, vSphere HA might separate clustered virtual machines that are meant to stay together, or vSphere HA might put clustered virtual machines that are meant to stay apart on the same host. You can avoid this problem by setting up DRS groups and using VM-Host affinity rules, which are obeyed by vSphere HA.

For a cluster of virtual machines on one physical host, all MSCS virtual machines must be in the same virtual machine DRS group, linked to the same host DRS group with the affinity rule "Must run on hosts in group."

For a cluster of virtual machines across physical hosts, each MSCS virtual machine must be in a different virtual machine DRS group, linked to a different host DRS group with the affinity rule "Must run on hosts in group."



CAUTION Limit the number of hosts to two when you define host DRS group rules for a cluster of virtual machines on one physical host. (This does not apply to clusters of virtual machines across physical hosts.) Since vSphere HA does not obey VM-VM affinity rules, virtual machines in the configuration could be spread across hosts during a vSphere HA recovery from host failure if more than two hosts are included in a host DRS group rule.

Create a Virtual Machine DRS Group (MSCS)

Before you can create a VM-Host affinity rule, you must create the host DRS group and the virtual machine DRS group that the rule applies to.

For a cluster of virtual machines on one physical host, create one virtual machine DRS group that contains all MSCS virtual machines. For example, VMGroup_1 contains MSCS_VM_1 and MSCS_VM_2.

For a cluster of virtual machines across physical hosts, create one virtual machines DRS group for each MSCS virtual machine. For example, VMGroup_1 contains MSCS_VM_1 and VMGroup_2 contains MSCS_VM_2.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the cluster Settings dialog box under **vSphere DRS**, select **DRS Groups Manager**.
- 3 In the Virtual Machines DRS Groups section, click **Add**.
- 4 In the DRS Group dialog box, type a name for the group (for example, VMGroup_1).
- 5 In the left pane, select the virtual machine and click >> to add it to the group.
 - For a cluster of virtual machines on one physical host, add all MSCS virtual machines to one group.
 - For a cluster of virtual machines across physical hosts, add one MSCS virtual machine per group.
- 6 Click **OK**.

Create a Host DRS Group (MSCS)

Before you can create a VM-Host affinity rule, you must create the host DRS group and the virtual machine DRS group that the rule applies to.

For a cluster of virtual machines on one physical host, create one host DRS group that contains the ESXi hosts. For example, HostGroup_1 contains ESXi_HOST_1 and ESXi_HOST_2.

For a cluster of virtual machines across physical hosts, create groups with sets of hosts that do not overlap. This ensures that virtual machines placed in different host groups do not ever run on the same host simultaneously.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the cluster Settings dialog box under **vSphere DRS**, select **DRS Groups Manager**.
- 3 In the Host DRS Groups section, click **Add**.
- 4 In the DRS Group dialog box, type a name for the group (for example, **HostGroup_1**).
- 5 In the left pane, select a host and click >> to add it to the group.
- 6 Click **OK**.

Set up VM-Host Affinity Rules for DRS Groups (MSCS)

Create VM-Host affinity rules to specify whether the members of a selected virtual machine DRS group can run on the members of a specific host DRS group.

Prerequisites

Create virtual machine DRS groups that contain one or more MSCS virtual machines as described in [“Create a Virtual Machine DRS Group \(MSCS\)”](#) on page 22.

Create host DRS groups that contain one or more ESXi hosts, as described in [“Create a Host DRS Group \(MSCS\)”](#) on page 22.

Procedure

- 1 In the vSphere Client, right-click the cluster in the inventory and select **Edit Settings**.
- 2 In the left pane of the Cluster Settings dialog box under vSphere DRS, select **Rules**.
- 3 Click **Add**.
- 4 In the Rule dialog box, type a name for the rule.
- 5 From the **Type** menu, select **Virtual Machines to Hosts**.
- 6 Select the virtual machine DRS group and the host DRS group to which the rule applies.
For example, select VMGroup_1 and HostGroup_1.
- 7 Select **Must run on hosts in group**.
- 8 Click **OK**.

vSphere MSCS Setup Checklist

When you set up MSCS on ESXi, see the checklists to configure your environment according to the requirements. You can also use the checklists to verify that your setup meets the requirements if you need technical support.

Requirements for Clustered Disks

Each type of clustered disk has its own requirements, depending on whether it is in a single-host cluster or multihost cluster.

Table 3. Requirements for Clustered Disks

Component	Single-Host Clustering	Multihost Clustering
Clustered virtual disk (.vmdk)	SCSI bus sharing mode must be set to virtual.	Not supported.
Clustered disks, virtual compatibility mode (non-pass-through RDM)	Device type must be set to virtual compatibility mode. SCSI bus sharing mode must be set to virtual mode. A single, shared RDM mapping file for each clustered disk is required.	Device type must be set to virtual compatibility mode for cluster across boxes, but not for standby host clustering or cluster across boxes on Windows Sever 2008. SCSI bus sharing mode must be set to physical. Requires a single, shared RDM mapping file for each clustered disk.

Table 3. Requirements for Clustered Disks (Continued)

Component	Single-Host Clustering	Multihost Clustering
Clustered disks, physical compatibility mode (pass-through RDM)	Not supported.	Device type must be set to Physical compatibility mode during hard disk creation. SCSI bus sharing mode must be set to physical (the default). A single, shared RDM mapping file for each clustered disk is required.
All types	All clustered nodes must use the same target ID (on the virtual SCSI adapter) for the same clustered disk. A separate virtual adapter must be used for clustered disks.	

Other Requirements and Recommendations

The following table lists the components in your environment that have requirements for options or settings.

Table 4. Other Clustering Requirements and Recommendations

Component	Requirement
Disk	If you place the boot disk on a virtual disk, select Thick Provision during disk provisioning. The only disks that you should not create with the Thick Provision option are RDM files (both physical and virtual compatibility mode).
Windows	Use Windows Server 2003 SP2 (32 bit), Windows Server 2003 (64 bit) SP2, or Windows Server 2008 R2. Only two cluster nodes. Disk I/O timeout is 60 seconds or more (HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeoutValue). NOTE If you recreate the cluster, this value might be reset to its default, so you must change it again. The cluster service must restart automatically on failure (first, second, and subsequent times).
ESXi configuration	Do not overcommit memory. Set the Memory Reservation (minimum memory) option to the same as the amount of memory assigned to the virtual machine. If you must overcommit memory, the swap file must be local, not on the SAN.
Multipathing	Contact your multipathing software vendor for information and support of non-VMware multipathing software in vSphere.

Required Information for Technical Support

The following table lists the files and settings to collect if you require technical support. Technical support uses these files and settings to analyze clustering issues.

NOTE Verify that your setup complies with the checklists in [Table 3](#) and [Table 4](#) before you contact technical support.

Table 5. Information Required by Technical Support

File or Information	Description or Location
vm-support tarball	Contains the vmkernel log, virtual machine configuration files and logs, and so on.
Application and system event logs of all virtual machines with the problem	
Cluster log of all virtual machines with the problem	%ClusterLog%, which is usually set to %SystemRoot%\cluster\cluster.log.
Disk I/O timeout	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeoutValue

Table 5. Information Required by Technical Support (Continued)

File or Information	Description or Location
vSphere Client display names and Windows NETBIOS names of the virtual machines experiencing the problem	
Date and time that the problem occurred	
SAN configuration of the ESXi system	Details about LUNs, paths, and adapters.
(Optional) Memory dump of the virtual machine	Required if a guest virtual machine fails (a blue screen appears with an error message).

Index

A

- affinity rules
 - DRS groups and MSCS **23**
 - enforcing and MSCS **21**
 - MSCS **20, 21**
- anti-affinity rules, MSCS **20**
- automation level, MSCS **21**

B

- boot from SAN, MSCS **9**

C

- cluster across boxes, MSCS **6, 20**
- cluster in a box, MSCS **10, 20**
- clustering
 - MSCS physical hosts and virtual machines **7**
 - MSCS virtual machines on one host **10**
 - physical and virtual machines **17**
 - virtual machines across hosts **13**
- compatibility mode
 - physical **13, 15–17**
 - virtual **13, 17**

D

- Database Availability Group (DAG), MSCS **10**
- disks
 - adding to nodes **12, 13, 15**
 - formatting **10**
 - formatting MSCS **13**
 - quorum **12, 15, 19**
 - shared **12, 15, 19**
- DRS groups
 - host **22**
 - MSCS **21, 22**
 - virtual machine **22**

E

- eagerzeroedthick **10, 13**
- ESXi **17**
- Ethernet adapter **10**
- Ethernet adapters and MSCS **13**

F

- Fault Tolerance (FT), MSCS **8, 9**
- Fault Tolerance (FT) and MSCS **23**
- Fibre Channel (FC) SAN **12**
- Fibre Channel (FC) SAN and MSCS **13**

- format
 - disks **10**
 - disks MSCS **13**
 - eagerzeroedthick **10, 13**

H

- hardware requirements, MSCS **8**
- high availability and MSCS, See vSphere HA
- hosts, standby **17**

I

- iSCSI and MSCS **9**
- iSCSI SAN **9**

L

- LSI Logic Parallel **12, 15, 16, 19**
- LSI Logic SAS **12, 15, 16, 19**

M

- Microsoft Cluster Service (MSCS), installing **19**
- MSCS
 - boot from SAN **9**
 - cluster across boxes **6**
 - cluster in a box **6**
 - cluster virtual machines on one host **6**
 - Clustered Continuous Replication (CCR) **10**
 - clustering configurations **6**
 - clustering limitations **9**
 - clustering physical and virtual machines **7**
 - clustering requirements **8**
 - Database Availability Group (DAG) **10**
 - Database Availability Groups and Exchange 2010 **10**
 - disk format **8**
 - DRS automation level **21**
 - DRS groups **22**
 - Fault Tolerance (FT) **8, 9**
 - Fibre Channel (FC) SAN **6, 9**
 - getting started **5**
 - guest operating system requirements **8**
 - hardware requirements **8**
 - hardware version 7 **9**
 - iSCSI **9**
 - LSI Logic Parallel **8**
 - LSI Logic SAS **8**

- Microsoft Exchange and CCR, *See* Clustered Continuous Replication (CCR)
 - Microsoft Exchange and Database Availability Groups **10**
 - multipathing **9**
 - multiple host clustering **6**
 - N-Port ID Virtualization (NPIV) **9**
 - native multipathing (NMP) **9**
 - NFS **9**
 - non-pass-through RDM **9**
 - NTP server **8**
 - pass-through RDM **9**
 - physical compatibility mode **9**
 - requirements **8**
 - SAN **6, 10**
 - shared storage configurations **9**
 - single host clustering **6**
 - standby host **7**
 - types of applications **6**
 - virtual compatibility mode **9**
 - virtual machines across hosts **6**
 - virtual NICs **8**
 - virtual SCSI adapter **8**
 - VMkernel **10, 13**
 - vMotion **9**
 - MSCS nodes, creating second **14, 18**
 - multipathing, MSCS **9**
 - multipathing and MSCS **23**
- N**
- native multipathing (NMP), MSCS **9**
 - NFS, MSCS **9**
 - nodes
 - creating first **10, 17**
 - creating first MSCS **13**
 - creating second **11**
 - NTP server, MSCS **8**
- P**
- physical compatibility mode, MSCS **13**
 - physical-virtual pairs, creating **20**
- Q**
- quorum disk **12, 15, 19**
- R**
- RDM
 - non-pass-through **17**
 - pass-through **17**
 - RDM and MSCS **13, 14, 23**
 - rules, vSphere DRS **20**
- S**
- SAN, MSCS **6, 10**
 - SAN and MSCS **23**
 - SAN LUN **12, 15**
 - SAN, boot from **9**
 - SCSI bus sharing
 - physical **15, 16, 19**
 - virtual **12**
 - shared storage, disks **12, 15, 19**
 - standby host, MSCS **7**
 - storage
 - quorum disk **12, 15, 19**
 - shared **12, 15, 19**
- T**
- technical support checklist, MSCS **23**
- U**
- updated information **5**
- V**
- virtual compatibility mode, MSCS **13**
 - VMkernel **17**
 - VMkernel and MSCS **10, 13**
 - vmkfstools **11**
 - vmkfstools and MSCS **14**
 - vMotion, MSCS **9**
 - vSphere DRS, enabling and MSCS **20**
 - vSphere DRS and MSCS **20**
 - vSphere HA, enabling and MSCS **20**
 - vSphere HA and MSCS **20**
 - vSphere High Availability (HA), *See* vSphere HA