

Setup for Failover Clustering and Microsoft Cluster Service

Update 2

11 APR 2019

VMware vSphere 6.7

VMware ESXi 6.7

vCenter Server 6.7

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

<https://docs.vmware.com/>

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

Copyright © 2006-2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Setup for Failover Clustering and Microsoft Cluster Service	5
1 Getting Started with MSCS	7
Clustering Configuration Overview	7
Clustering MSCS Virtual Machines on a Single Host	8
Clustering Virtual Machines Across Physical Hosts	9
Clustering Physical Machines with Virtual Machines	10
Hardware and Software Requirements for Clustering	11
Supported Shared Storage Configurations	11
PSP_RR Support for MSCS	12
iSCSI Support for MSCS	12
FCoE Support for MSCS	13
vMotion support for MSCS	14
VVol Support for MSCS	15
vSphere MSCS Setup Limitations	15
MSCS and Booting from a SAN	16
Set up CCR and DAG Groups	16
Setting up AlwaysOn Availability Groups with SQL Server 2012	16
Virtual Networking Best Practices for MSCS	17
2 Cluster Virtual Machines on One Physical Host	18
Create the First Node for Clusters on One Physical Host	18
Create Additional Nodes for Clusters on One Physical Host	19
Add Hard Disks to the First Node for Clusters on One Physical Host	20
Add Hard Disks to Additional Nodes for Clusters on One Physical Host	21
3 Cluster Virtual Machines Across Physical Hosts	22
Create the First Node for MSCS Clusters Across Physical Hosts	22
Create Additional Nodes for Clusters Across Physical Hosts	23
Add Hard Disks to the First Node for Clusters Across Physical Hosts	24
Add Hard Disks to the First Node for Clusters Across Physical Hosts with VVol	26
Add Hard Disks to Additional Nodes for Clusters Across Physical Hosts	27
4 Cluster Physical and Virtual Machines	29
Create the First Node for a Cluster of Physical and Virtual Machines	29
Create the Second Node for a Cluster of Physical and Virtual Machines	30
Add Hard Disks to the Second Node for a Cluster of Physical and Virtual Machines	31
Install Microsoft Cluster Service	32

Create Additional Physical-Virtual Pairs 32

5 Use MSCS in an vSphere HA and vSphere DRS Environment 33

Enable vSphere HA and vSphere DRS in a Cluster (MSCS) 33

Create VM-VM Affinity Rules for MSCS Virtual Machines 33

Enable Strict Enforcement of Affinity Rules (MSCS) 34

Set DRS Automation Level for MSCS Virtual Machines 34

Using vSphere DRS Groups and VM-Host Affinity Rules with MSCS Virtual Machines 35

Create a Virtual Machine DRS Group (MSCS) 36

Create a Host DRS Group (MSCS) 36

Set up VM-Host Affinity Rules for DRS Groups (MSCS) 37

6 vSphere MSCS Setup Checklist 38

About 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, Windows Server 2012 and above releases. 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 and above releases.

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.

Note All instances of "MSCS" in this document and elsewhere is intended to refer to all Windows OS and Microsoft application clustering configurations which utilize one or more shared disks among the clustered nodes, namely:

- The Microsoft Cluster Service (MSCS)
 - The Windows Server Failover Clustering (WSFC)
 - The SQL AlwaysOn Failover Cluster Instance (FCI)
-

vSphere Client and vSphere Web Client

Instructions in this guide reflect the vSphere Client (an HTML5-based GUI). You can also use the instructions to perform the tasks by using the vSphere Web Client (a Flex-based GUI).

Tasks for which the workflow differs significantly between the vSphere Client and the vSphere Web Client have duplicate procedures that provide steps according to the respective client interface. The procedures that relate to the vSphere Web Client, contain vSphere Web Client in the title.

Note In vSphere 6.7 Update 1, almost all of the vSphere Web Client functionality is implemented in the vSphere Client. For an up-to-date list of any remaining unsupported functionality, see [Functionality Updates for the vSphere Client](#).

Getting Started with MSCS

1

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.

This chapter includes the following topics:

- [Clustering Configuration Overview](#)
- [Hardware and Software Requirements for Clustering](#)
- [Supported Shared Storage Configurations](#)
- [PSP_RR Support for MSCS](#)
- [iSCSI Support for MSCS](#)
- [FCoE Support for MSCS](#)
- [vMotion support for MSCS](#)
- [VVol Support for MSCS](#)
- [vSphere MSCS Setup Limitations](#)
- [MSCS and Booting from a SAN](#)
- [Setting up Clustered Continuous Replication or Database Availability Groups with Exchange](#)
- [Setting up AlwaysOn Availability Groups with SQL Server 2012](#)
- [Virtual Networking Best Practices for MSCS](#)

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, FCoE or iSCSI. A quorum disk must have a homogenous set of disks. This means that if the configuration is done with FC SAN, then all of the cluster disks should be FC SAN only. Mixed mode is not supported.
- A private heartbeat network between nodes.

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

In ESXi 6.7, MSCS pass-through support for VVols (Virtual Volumes) permits the shared disk to be on a VVol storage that supports SCSI Persistent Reservations for VVols.

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 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.

Note Windows Server 2008 R2 and above releases support up to five nodes (virtual machines). Windows Server 2003 SP2 systems support two nodes.

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 1-1. 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.

Note A quorum disk can be configured with iSCSI, FC SAN or FCoE. A quorum disk must have a homogenous set of disks. This means that if the configuration is done with FC SAN, then all of the cluster disks should be FC SAN only. Mixed mode is not supported.

Figure 1-2. Virtual Machines Clustered Across Hosts




Note Windows Server 2008 SP2 and above systems support up to five nodes (virtual machines). Windows Server 2003 SP1 and SP2 systems support two nodes (virtual machines). For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#).

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 1-3. 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 1-4. 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 1-1. Clustering Requirements

Component	Requirement
Virtual SCSI adapter	LSI Logic Parallel for Windows Server 2003. LSI Logic SAS for Windows Server 2008 SP2 and above. VMware Paravirtual for Windows Server 2008 SP2 and above.
Operating system	Windows Server 2003 SP1 and SP2, Windows Server 2008 SP2 above releases. For supported guest operating systems see Table 6-2. Other Clustering Requirements and Recommendations .
Virtual NIC	VMXNET3 is required. You must activate in-guest Windows Receive Side Scaling (RSS) on the virtual NIC. Note Drivers, improvements and updates to the VMXNET3 virtual NIC are delivered through VMware Tools. Upgrade to the latest version of VMware Tools on any VM using the VMXNET3 virtual NIC.
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.
Number of nodes	Windows Server 2003 SP1 and SP2 : two-node clustering Windows Server 2008 SP2 and above: up to five-node clustering For supported guest operating systems see Table 6-2. Other Clustering Requirements and Recommendations .
NTP server	Synchronize domain controllers and cluster nodes with a common NTP server, and deactivate 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 1-2. 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	No	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.

PSP_RR Support for MSCS

ESXi 6.0 supports PSP_RR for MSCS.

- ESXi 6.0 supports PSP_RR for Windows Server 2008 SP2 and above releases. Windows Server 2003 is not supported.
- PSPs configured in mixed mode is supported. In a 2 node cluster one ESXi host can be configured to use PSP_FIXED and the other ESXi host can use PSP_RR.
- Shared disk quorum or data must be provisioned to guest in PassThrough RDM mode only.
- All hosts must be running ESXi 6.0.
- Mixed mode configurations of ESXi 6.0 with previous ESXi releases are not supported.
- Rolling upgrades of cluster Hosts from previous versions of ESXi to ESXi 6.0 builds are not supported.

iSCSI Support for MSCS

ESXi 6.0 supports iSCSI storage and up to 5 node MSCS cluster using Qlogic, Emulex and Broadcom adapters.

- ESXi 6.0 supports iSCSI for Windows Server 2008 SP2 and above releases. Windows Server 2003 is not supported.
- Cluster-across-box (CAB) and cluster-in-a-box (CIB) are supported. A mixture of CAB and CIB is not supported.
- No qualification is needed for SWiSCSI initiator in a guest operating system.

- N+1 cluster configuration comprising of a cluster between "N" virtual machines on separate ESXi hosts and one physical machine running Windows natively is supported.
- All hosts must be running ESXi 6.0.
- Mixed cluster nodes running FC or FCOE and iSCSI are not supported.
- Mixed mode of iSCSI config is supported. For example, Node A on ESXi with iSCSI software initiator and Node B on ESXi with Qlogic, Emulex or Broadcom hardware adapter.
- Mixed mode configurations of ESXi 6.0 with previous ESXi releases are not supported.
- Rolling upgrades of cluster Hosts from previous versions of ESXi to ESXi 6.0 builds are not supported.

FCoE Support for MSCS

ESXi 6.0 supports FCoE storage and up to 5 node MSCS clusters using Cisco FNIC and Emulex FCoE adapters.

- ESXi 6.0 supports FCoE for Windows Server 2008 SP2 and above releases. Windows Server 2003 is not supported.
- Cluster-across-box (CAB) and cluster-in-a-box (CIB) are supported. A mixture of CAB and CIB is not supported.
- CAB configurations are supported with some cluster nodes on physical hosts. In a CAB configuration, a max of one virtual machine in a host can see a LUN.
- In a CIB configuration all virtual machines must be on the same host.
- No qualification is needed for SWiSCSI and FCoE initiators in a guest operating system.
- N+1 cluster configuration, in which one ESXi host has virtual machines which are secondary nodes and one primary node is a physical box are supported.
- Standard affinity and anti-affinity rules apply for MSCS virtual machines.
- All hosts must be running ESXi 6.0.
- All hosts must be running FCoE initiators. Mixed cluster nodes running FC and FCoE are not supported.
- Mixed mode FCoE configuration is supported. For example, Node A on ESXi with an FCoE software adapter intel based card and Node B on ESXi with an Emulex or Cisco FCoE hardware adapter.
- Mixed mode configurations of ESXi 6.0 with previous ESXi releases are not supported.
- Rolling upgrades of cluster hosts from previous versions of ESXi to ESXi 6.0 builds are not supported.

vMotion support for MSCS

vSphere 6.0 adds support for vMotion of MSCS clustered virtual machines.

Pre-requisites for vMotion support:

- vMotion is supported only for a cluster of virtual machines across physical hosts (CAB) with pass-through RDMS.
- The vMotion network must be a 10Gbps Ethernet link. 1Gbps Ethernet link for vMotion of MSCS virtual machines is not supported.
- vMotion is supported for Windows Server 2008 SP2 and above releases. Windows Server 2003 is not supported.
- The MSCS cluster heartbeat time-out must be modified to allow 10 missed heartbeats.
- The virtual hardware version for the MSCS virtual machine must be version 11 and later .

Modifying the MSCS heartbeat time-out:

Failover cluster nodes use the network to send heartbeat packets to other nodes of the cluster. If a node does not receive a response from another node for a specified period of time, the cluster removes the node from cluster membership. By default, a guest cluster node is considered down if it does not respond within 5 seconds. Other nodes that are members of the cluster will take over any clustered roles that were running on the removed node.

An MSCS virtual machine can stall for a few seconds during vMotion. If the stall time exceeds the heartbeat time-out interval, then the guest cluster considers the node down and this can lead to unnecessary failover. To allow leeway and make the guest cluster more tolerant, the heartbeat time-out interval needs to be modified to allow 10 missed heartbeats. The property that controls the number of allowed heart misses is **SameSubnetThreshold**. You will need to modify this from its default value to 10. From any one of the participating MSCS cluster nodes run the following command: `cluster <cluster-name> /prop SameSubnetThreshold=10:DWORD`.

You can also adjust other properties to control the workload tolerance for failover. Adjusting delay controls how often heartbeats are sent between the clustered node. The default setting is 1 second and the maximum setting is 2 seconds. Set the **SameSubnetDelay** value to 1. Threshold controls how many consecutive heartbeats can be missed before the node considers its partner to be unavailable and triggers the failover process. The default threshold is 5 heartbeats and the maximum is 120 heartbeats. It is the combination of delay and threshold that determines the total elapsed time during which clustered Windows nodes can lose communication before triggering a failover. When the clustered nodes are in different subnets, they are called **CrossSubnetDelay** and **CrossSubnetThreshold**. Set the **CrossSubnetDelay** value to 2 and the **CrossSubnetThreshold** value to 10.

VVol Support for MSCS

ESXi 6.7 supports VVol (Virtual Volume) storage and up to 5 node MSCS clusters. The storage array must support SCSI persistent operations at the subsidiary LUN level.

- ESXi 6.7 supports VVol Storage for Windows Server 2008 SP2 and above releases.
- All hosts must be running ESXi 6.7 or above.
- Only Cluster-across-box (CAB) is supported.
- MSCS on VVols can work with any type of disk, "Thin" as well as "Thick"-provisioned disks.
- This feature enables customers to move away from using Pass-through RDM (physical compatibility mode).
- MSCS on VVols supports HA, DRS and vMotion.
- The underlying transport protocol can be FC, iSCSI or FCOE.
- Cluster-in-a-box (CIB) and a mixture of CAB and CIB is not supported.
- N+1 cluster configuration, in which one ESXi host has virtual machines which are secondary nodes and one primary node is a physical box is not supported.

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 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 on a single host (CIB).
- N-Port ID Virtualization (NPIV)
- 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.
- Suspend or resume of more than one MSCS node in an ESXi host with a five-node cluster in a box configuration is not supported. This I/O intensive operation is disruptive of the timing sensitive MSCS clustering software.
- Storage spaces are not supported with Failover clustering on Windows 2012 and above.

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

You can set up Clustered Continuous Replication (CCR) with Exchange 2007 or Database Availability Groups (DAG) with Exchange 2010 and higher in your vSphere environment. For supported versions of Exchange, see VMware knowledge base article 1037959.

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](#).

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

Setting up AlwaysOn Availability Groups with SQL Server 2012

You can set up AlwaysOn Availability Groups (AAG) with SQL Server 2012 in your vSphere environment.

vSphere 6.0 supports the following deployments of AAG:

- Using Availability Groups (AG) for high availability and disaster recovery solution (Non-shared disk configuration).
- Using Failover Cluster Instance (FCI) for high availability, and Availability Groups (AG) for disaster recovery solution (Shared Disk Configuration).

When working in a vSphere environment:

- Use virtual machines instead of physical machines as the cluster components.
- If the boot disks of the AAG virtual machines are on a SAN, see [MSCS and Booting from a SAN](#).

For more information, see Microsoft's documentation for AAG on the Microsoft Web site.

Virtual Networking Best Practices for MSCS

Some cluster workloads are more sensitive to network latency than others. To configure the network for MSCS, start with a thorough understanding of your workload network requirements.

The following guidelines apply to provisioning the network for a MSCS VM:

- Traffic types must be separated to keep like traffic contained to designated networks. vSphere can use separate interfaces for management, vSphere vMotion, and network-based storage traffic. Additional interfaces can be used for VM traffic such as MSCS traffic. Within VMs, different interfaces can be used to keep certain traffic separated.
- Use 802.1q VLAN tagging and virtual switch port groups to logically separate traffic. In the case of different logical traffic, include two or more NIC cards in a team to increase the network capacity and provide failover capability. If you configure the teaming and failover policy on a virtual switch, the policy is propagated to all port groups in the switch. If you configure the policy on the port group, it overrides the policy inherited from switch. Using different types of network adapters eliminates connectivity across all network adapters if there is an issue with the NIC driver.
- Use separate physical interfaces and dedicated port groups or virtual switches to physically separate traffic.
- If you are using iSCSI, the network adapters must be dedicated to either network communication or iSCSI, but not both.
- Enable jumbo frames on the virtual switches where you have enabled vSphere vMotion traffic or iSCSI traffic. Make sure that jumbo frames are also enabled on your physical network infrastructure before making this configuration on the virtual switches.
- Use multi-nic vMotion in cases where a vMotion operation is initiated for a large memory VM and its large memory size is intensively utilized. With multi-nic vMotion, every vMotion operation utilizes multiple port links and this speeds up the vMotion process.
- Use the VMXNET3 paravirtualized NIC. VMXNET 3 is the latest generation of paravirtualized NICs designed for performance. It offers several advanced features including multi-queue support, Receive Side Scaling, IPv4/IPv6 offloads, and MSI/MSI-X interrupt delivery.
- Follow the guidelines on guest operating system networking considerations and hardware networking considerations in the Performance Best Practices for VMware vSphere 6.0 guide.

Cluster Virtual Machines on One Physical Host

2

You can create an MSCS cluster with up to five nodes on a single ESXi host.

Note Windows Server 2008 SP2 and above systems support up to five nodes (virtual machines). For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#). Windows Server 2003 SP1 and SP2 systems support two nodes.

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

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.

Page	Action
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules.
Ready to Complete	Review your selections.

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

Note Do not add shared cluster disks at this time.

- In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
- Click the **New device** drop-down menu, select **Network**, and click **Add**.
- In the **New Virtual Machine - Edit Settings** dialog box, expand **Network adapter**. 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 **OK**.
- Install a Windows Server operating system on the virtual machine.

Create Additional Nodes for Clusters on One Physical Host

Create a template from the first virtual machine and deploy the second node from that template. You can have up to five nodes on Windows Server 2008.

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

Page	Action
Name and folder	Enter a name (for example, Node2_Template) and select a location.
Compute Resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .

Page	Action
Storage	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Ready to Complete	Click Finish to create the virtual machine template.

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

Page	Action
Name and folder	Enter a name (for example, Node2) and select a location.
Compute resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Clone options	Select Customize the operating system .

- 5 Select a new guest operating system from the list.
 - a Click the **Create a new specification** button to add a new guest operating system. Proceed through the **New VM Guest Customization Spec** wizard.
 - b Click **Finish** to exit the 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, select the newly created virtual machine, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, select **New Hard Disk**, and click **Add**.
- 3 Select the disk size.
- 4 Under Disk Provisioning, select **Thick Provision**.
You can also use a mapped SAN LUN set to virtual compatibility mode.
- 5 Expand the **New Hard Disk**. 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.

- 6 Click **OK**.

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

- 7 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **Change Type** drop-down menu.
- 8 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003 SP1 and SP2	LSI Logic Parallel
Windows Server 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#).

- 9 Click **OK**.
- 10 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **SCSI Bus Sharing** drop-down menu. Set SCSI Bus Sharing to **Virtual** and click **OK**.

Add Hard Disks to Additional Nodes 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, select the newly created virtual machine, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, select **Existing Hard Disk**, and click **Add**.
- 3 Select the same virtual device node you chose for the first virtual machine's shared storage disks (for example, **SCSI (1:0)**), and click **OK**.

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.

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

Cluster Virtual Machines Across Physical Hosts

3

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

A cluster across physical hosts requires specific hardware and software.

- To configure the network for MSCS, start with a thorough understanding of your workload network requirements. For more information, see [Virtual Networking Best Practices for MSCS](#).
- Supported shared storage configuration. For more information, see [Supported Shared Storage Configurations](#).
- 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 vCenter Server system.
Use the user name and password of the user who will have administrator permissions on the virtual machine.
- 2 In the vSphere Client, right-click the host and select **New Virtual Machine**.

- Proceed through the wizard to create the virtual machine.

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules.
Ready to Complete	Review your selections.

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

Note Do not add shared cluster disks at this time.

- In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
- Click the **New device** drop-down menu, select **Network**, and click **Add**.
- In the **New Virtual Machine - Edit Settings** dialog box, expand **Network adapter**. 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 **OK**.
- Install a Windows Server operating system on the virtual machine.

Create Additional Nodes for Clusters Across Physical Hosts

To create additional nodes in a cluster of virtual machines across physical hosts, you create a template of the first virtual machine and use it to deploy additional virtual machines onto another 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, right-click the first virtual machine you created and select **Clone > Clone to Template**.
- 2 Proceed through the wizard to create the virtual machine template.

Page	Action
Name and folder	Enter a name (for example, Node2_Template) and select a location.
Compute Resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Storage	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Ready to Complete	Click Finish to create the virtual machine template.

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

Page	Action
Name and folder	Enter a name (for example, Node2) and select a location.
Compute resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Clone options	Select Customize the operating system .

- 5 Select a new guest operating system from the list.
 - a Click the **Create a new specification** button to add a new guest operating system. Proceed through the **New VM Guest Customization Spec** wizard.
 - b Click **Finish** to exit the 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.
- If you are using VVol storage for shared disk, please see *Add Hard Disks to the First Node for Clusters Across Physical Hosts with VVol*.

Note Use RDMs in physical compatibility mode.

Procedure

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

2 Click the **New device** drop-down menu, select **RDM Disk**, and click **Add**.

3 Select an unformatted LUN.

4 Select a datastore.

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

5 Select **Physical** as the compatibility mode.

6 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.

7 Click **OK** to complete creating the disk.

The wizard creates a new hard disk.

8 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **Change Type** drop-down menu.

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

Operating System	Type of Controller
Windows Server 2003 SP1 and SP2	LSI Logic Parallel
Windows Server 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#).

10 Click **OK**.

11 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **SCSI Bus Sharing** drop-down menu. Set SCSI Bus Sharing to **Physical** and click **OK**.

Results

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 First Node for Clusters Across Physical Hosts with VVol

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

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.

Procedure

- 1 In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, select **New Hard Disk**, and click **Add**.
- 3 Select the disk size.
- 4 Under Disk Provisioning, select either Thick or Thin Provision.
- 5 Expand the New Hard Disk.
- 6 From the Virtual Device Node drop-down menu, select a new SCSI controller (for example, select **SCSI (1:0)**).

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

- 7 Click **OK**.

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

- 8 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **Change Type** drop-down menu.

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

Operating System	Type of Controller
Windows Server 2003 SP1 and SP2	LSI Logic Parallel
Windows Server 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#).

- 10 Click **OK**.
- 11 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **SCSI Bus Sharing** drop-down menu. Set SCSI Bus Sharing to **Physical** and click **OK**.

Results

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 Additional Nodes 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, select the newly created virtual machine, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, select **Existing Hard Disk**, and click **Add**.
- 3 In Disk File Path, browse to the location of the quorum disk specified for the first node.
- 4 Select **Physical** as the compatibility mode and click **Next**.

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

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.

- 6 Click **OK**.

The wizard creates a new hard disk.

- 7 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **Change Type** drop-down menu.
- 8 Select the appropriate type of controller, depending on your operating system.

Operating System	Type of Controller
Windows Server 2003 SP1 and SP2	LSI Logic Parallel
Windows Server 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 6-2. Other Clustering Requirements and Recommendations](#).

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

Cluster Physical and Virtual Machines

4

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.

- To configure the network for MSCS, start with a thorough understanding of your workload network requirements. For more information, see [Virtual Networking Best Practices for MSCS](#).
- 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.
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](#) is also visible from the virtual machine.

Procedure

- 1 Open the vSphere Client and connect to the vCenter Server system.

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

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

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules.
Ready to Complete	Review your selections.

- 4 Click **Finish** to complete creating the virtual machine.
- 5 In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
- 6 Click the **New device** drop-down menu, select **Network**, and click **Add**.
- 7 In the **New Virtual Machine - Edit Settings** dialog box, expand **Network adapter**. 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.
- 8 Click **Finish** to complete creating the device.
 - 9 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, select the newly created virtual machine, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, select **RDM Disk**, and click **Add**.
- 3 Select the LUN that is used by the physical machine.
- 4 Select the datastore, which is also the location of the boot disk.
- 5 Select **Physical** as the compatibility mode.
- 6 Expand the **New Hard Disk**. 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.

- 7 Click **OK**.
The wizard creates a new hard disk.
- 8 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **Change Type** drop-down menu.
- 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 and VMware Paravirtual

- 10 Click **OK**.
- 11 In the **New Virtual Machine - Edit Settings** dialog box, expand **SCSI controller** and select the **SCSI Bus Sharing** drop-down menu. 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 deactivate the storage validation heuristics.
- 2 Select the **Advanced** (minimum) configuration option, and click **OK**.

Results

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

5

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. vSphere includes vMotion support for MSCS cluster virtual machines.

Enable vSphere HA and vSphere DRS in a 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 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Services**, select **vSphere DRS** and click **Edit**.
- 4 Select the **Turn ON vSphere DRS** and **Turn on vSphere HA** check boxes.
- 5 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.

Note vMotion is supported only for a cluster of virtual machines across physical hosts with pass-through RDMs. For a cluster of virtual machines on one physical host and a cluster of virtual machines across physical hosts with non-pass-through RDMs, vMotion is not supported.

Procedure

- 1 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Configuration**, click **VM/Host Rules**.
- 4 Click **Add**.
- 5 In the **Create VM/Host Rule** dialog box, type a name for the rule.
- 6 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**.
- 7 Click **Add**.
- 8 Select the two virtual machines to which the rule applies and click **OK**.
- 9 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 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Services**, select **vSphere DRS** and click **Edit**.
- 4 Expand **Advanced Options** and click **Add**.
- 5 In the Option column, type `ForceAffinePoweron`.
- 6 In the Value column, type `1`.
- 7 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.

Procedure

- 1 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Services**, select **vSphere DRS** and click **Edit**.
- 4 Expand DRS Automation, under Virtual Machine Automation select the **Enable individual virtual machine automation levels** check box and click **OK**.
- 5 Under **Configuration**, select **VM Overrides** and click **Add**.
- 6 Click the **+** button, select the MSCS virtual machines in the cluster and click **OK**.
- 7 Click the Automation level drop-down menu, and select **Partially Automated**.
- 8 Click **OK**.

Note VMware recommends partially automated mode for MSCS virtual machines, but there is no technical restriction which prevents the setting for MSCS virtual machines to be fully automated. If the fully automated setting is used please tune the migration threshold to suit the workload running on the MSCS virtual machines.

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

You can 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 rule establishes an 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, all MSCS virtual machines must be in the same virtual machine DRS group and all hosts must be in the same host DRS group. The virtual machine and host DRS groups must be linked 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 both, a cluster of virtual machines on one physical host(CIB) and a cluster of virtual machines across physical hosts(CAB), create one virtual machine DRS group that contains all MSCS virtual machines. For example, VMGroup_1 contains MSCS_VM_1, MSCS_VM_2 ... MSCS_VM_5.

Procedure

- 1 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Configuration**, select **VM/Host Groups** and click **Add**.
- 4 In the **Create VM/Host Group** dialog box, type a name for the group.
- 5 Select **VM Group** from the **Type** drop down box and click **Add**.
- 6 Click the check box next to a virtual machine to add it. Continue this process until all desired virtual machines have been added.
 - 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 all MSCS virtual machines to one group.
- 7 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 both the ESXi hosts. For example, HostGroup_1 contains ESXi_HOST_1, ESXi_HOST_2.

For a cluster of virtual machines across physical hosts, create one host DRS group that contains all the ESXi hosts. For example, HostGroup_1 contains ESXi_HOST_1 running MSCS_VM_1, ESXi_HOST_2 running MSCS_VM_2 ESXi_HOST_5 running MSCS_VM_5 and ESXi_HOST_6 as a standby host.

Procedure

- 1 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Configuration**, select **VM/Host Groups** and click **Add**.
- 4 In the **Create VM/Host Group** dialog box, type a name for the group (for example, **HostGroup_1**).

- 5 Select **Host Group** from the **Type** drop down box and click **Add**.
- 6 Click the check box next to a host to add it. Continue this process until all desired hosts have been added.
- 7 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 \(MSCS\)](#).

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

Procedure

- 1 Browse to the cluster in the vSphere Client.
- 2 Click **Configure**.
- 3 Under **Configuration**, select **VM/Host Rules** and click **Add**.
- 4 In the **Create VM/Host 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

6

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 6-1. 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. See VMware knowledge base article http://kb.vmware.com/kb/1016106 to mark the device as perennially reserved on RDM LUNs used for MSCS virtual machines. This configuration is not supported on Windows Server 2008 and higher. This is a valid configuration only on Windows Server 2003.
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.

Table 6-1. Requirements for Clustered Disks (continued)

Component	Single-Host Clustering	Multihost Clustering
Clustered disks, physical compatibility mode (Shared VVol storage)	Not Supported.	SCSI bus sharing mode must be set to physical.
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 6-2. 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 <code>Thick Provision</code> option are RDM files (both physical and virtual compatibility mode).
Windows	<p>Use:</p> <ul style="list-style-type: none"> ■ Windows Server 2003 SP1 (32 bit/64 bit) ■ Windows Server 2003 SP2 (32 bit/64 bit) ■ Windows Server 2003 R2 (32 bit/64 bit) ■ Windows Server 2003 R2 SP1 (32 bit/64 bit) ■ Windows Server 2003 R2 SP2 (32 bit/64 bit) ■ Windows Server 2008 SP1 (32 bit/64 bit) ■ Windows Server 2008 SP2 (32 bit/64 bit) ■ Windows Server 2008 R2 (64 bit) ■ Windows Server 2008 R2 SP1 (32 bit/64 bit) ■ Windows Server 2012 ■ Windows Server 2012 R2 <p>For Windows Server 2003 SP1 and SP2, use only two cluster nodes. For Windows Server 2008 SP2 and above, you can use up to five cluster nodes. Disk I/O timeout is 60 seconds or more (<code>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue</code>).</p> <p>Note If you recreate the cluster, this value might be reset to its default, so you must change it again.</p> <p>The cluster service must restart automatically on failure (first, second, and subsequent times).</p>

Table 6-2. Other Clustering Requirements and Recommendations (continued)

Component	Requirement
ESXi configuration	<p>Do not overcommit memory. Set the Memory Reservation (minimum memory) option to the same as the amount of memory assigned to the virtual machine.</p> <p>If you must overcommit memory, the swap file must be local, not on the SAN.</p> <p>ESXi 5.0 uses a different technique to determine if Raw Device Mapped (RDM) LUNs are used for MSCS cluster devices, by introducing a configuration flag to mark each device as "perennially reserved" that is participating in an MSCS cluster. For ESXi hosts hosting passive MSCS nodes with RDM LUNs, use the <code>esxcli storage core device setconfig -d <naa.id> --perennially-reserved=true</code>. See KB 1016106 for more information.</p>
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 6-1. Requirements for Clustered Disks](#) and [Table 6-2. Other Clustering Requirements and Recommendations](#) before you contact technical support.

Table 6-3. Information Required by Technical Support

File or Information	Description or Location
<code>vm-support tarball</code>	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
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).