

Setup for Windows Server Failover Clustering

02 APR 2020

VMware vSphere 7.0

VMware ESXi 7.0

vCenter Server 7.0



vmware®

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

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

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

docfeedback@vmware.com

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

Copyright © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Setup for Windows Server Failover Clustering on VMware® vSphere®	5
1 Getting Started with WSFC on VMware® vSphere®	6
Clustering Configuration Overview	6
Hosting virtual nodes of a WSFC on a Single Host	7
Clustering Virtual Machines Across Physical Hosts	7
Clustering Physical Machines with Virtual Machines	8
Hardware and Software Requirements for WSFC on vSphere	9
Supported Shared Storage Configurations	10
Round Robin Path Selection Policy (PSP_RR) Support for disk resources used by WSFC	12
iSCSI Support for WSFC	12
FCoE Support for WSFC	12
vMotion support for WSFC	13
VMware vSphere® Virtual Volumes™ Support for WSFC	14
vSphere WSFC Setup Limitations	14
WSFC and Booting from a SAN	15
2 Cluster Virtual Machines Across Physical ESXi Hosts	16
Clustered VMDK support for WSFC	16
Enabling Clustered VMDK Support	16
Configuration Maximums for Clustered VMDK Support	17
Recommendations for using Clustered VMDKs with WSFC	17
Requirements for using Clustered VMDKs with WSFC	17
Limitations of Clustered VMDK support for WSFC	18
Create the First Node for WSFC Clusters Across Physical Hosts	18
Create Additional Nodes for a WSFC Across Physical Hosts	19
Add Hard Disks to the First VM, first node of a WSFC, for Clusters Across Physical Hosts with pRDMS	20
Add Hard Disks to the First Node of a WSFC Across Physical Hosts with vVols	21
Add Hard Disks to the First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS Datastores	22
Add Hard Disks to Additional Nodes for Clusters Across Physical Hosts	23
3 Use WSFC in an vSphere HA and vSphere DRS Environment	25
Enable vSphere HA and vSphere DRS in ESXi	25
Create VM-VM Affinity Rules for WSFC Virtual Machines	25
Set DRS Automation Level for WSFC Virtual Machines	26
Using vSphere DRS Groups and VM-Host Affinity Rules with WSFC Virtual Machines	27

	Create a Virtual Machine DRS Group (WSFC)	27
	Create a Host DRS Group (WSFC)	28
	Set up VM-Host Affinity Rules for DRS Groups (WSFC)	28
4	vSphere WSFC Setup Checklist	30
5	Cluster Virtual Machines on One Physical Host	33
	Create the First Node of a WSFC on One Physical ESXi Host	33
	Create Additional Nodes for Clusters on One Physical Host	34
	Add Hard Disks to the First Node for Clusters on One Physical Host	35
	Add Hard Disks to Additional Nodes for Clusters on One Physical Host	36
6	Cluster Physical and Virtual Machines	37
	Create the First Node for a Cluster of Physical and Virtual Machines	37
	Create the Second Node for a Cluster of Physical and Virtual Machines	38
	Add Hard Disks to the Second Node for a Cluster of Physical and Virtual Machines	38
	Install Microsoft Cluster Service	39
	Create Additional Physical-Virtual Pairs	39

About Setup for Windows Server Failover Clustering on VMware[®] vSphere[®]

Setup for Windows Server Failover Clustering describes the supported configurations for a WSFC with shared disk resources you can implement using virtual machines with Failover Clustering for Windows Server 2008 and above releases. You get step-by-step instructions for each configuration and a checklist of clustering requirements and recommendations.

Unless stated otherwise, the term Windows Server Failover Clustering (WSFC) applies to Failover Clustering with Windows Server 2008 and above releases.

Setup for Windows Server Failover Clustering covers ESXi and VMware[®] vCenter[®] Server.

Intended Audience

This information is for system administrators who are familiar with VMware technology and Failover Clustering for Windows.

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

Note All instances of "WSFC" 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)
 - Microsoft SQL Server Always On Failover Cluster Instances (FCI)
-

Note Other WSFC-based solutions not accessing shared storage such as SQL Server Always On Availability Groups or Exchange Database Availability Group, require no special storage configurations on the vSphere side using VMFS or NFS. This guide should not be used for such configurations.

Getting Started with WSFC on VMware vSphere

1

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

Note vSphere High Availability (vSphere HA) supports a Windows clustering solution. *vSphere Availability* describes vSphere HA functionality.

This chapter includes the following topics:

- [Clustering Configuration Overview](#)
- [Hardware and Software Requirements for WSFC on vSphere](#)
- [Supported Shared Storage Configurations](#)
- [Round Robin Path Selection Policy \(PSP_RR\) Support for disk resources used by WSFC](#)
- [iSCSI Support for WSFC](#)
- [FCoE Support for WSFC](#)
- [vMotion support for WSFC](#)
- [VMware vSphere® Virtual Volumes™ Support for WSFC](#)
- [vSphere WSFC Setup Limitations](#)
- [WSFC and Booting from a SAN](#)

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 WSFC clusters in several configurations, depending on your environment.

A typical clustering setup includes:

- Disks resources that are shared between nodes. A shared disk is required as a quorum disk. In a cluster of virtual machines across physical hosts, shared disks can be RDMS, vVol or VMFS VMDKs. If you are using RDMS or vVol as shared disks, they can accessed using Fibre Channel (FC) SAN,

FCoE or iSCSI. If you are using VMFS VMDKs as shared disks, only FC SAN connected storage devices are supported. A WSFC must have a homogenous set of disks. This means that if the configuration is done with disks attached using FC SAN, then all of the cluster disks should be FC SAN only. Mixed mode is not supported.

In ESXi 7.0, clustered VMDKs on VMFS datastores are supported in configurations when the VMs hosting the nodes of a cluster are located on different ESXi hosts, known as cluster-across-boxes (CAB) configuration.

Hosting virtual nodes of a WSFC on a Single Host

A cluster of WSFC virtual machines on a single host (also known as a cluster in a box or CIB) 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.

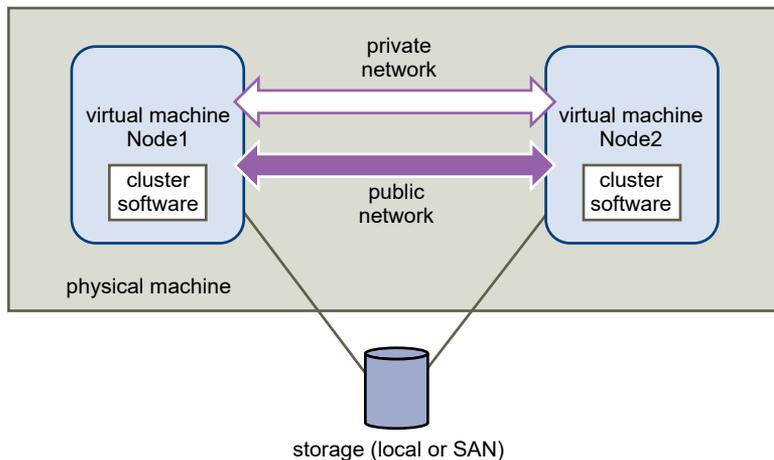
Note Do not host virtual nodes of a WSFC on the same ESXi host for a production deployment.

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

The following figure shows a cluster in a box setup.

- Two virtual machines on the same ESXi host run WSFC.
- 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 ESXi hosts (also known as a cluster across boxes or CAB) protects against software failures and hardware failures on the physical ESXi hosts by placing the cluster

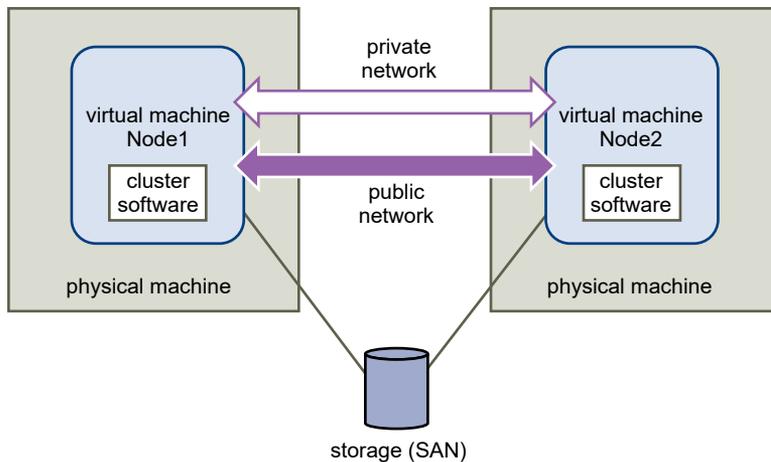
nodes on separate ESXi hosts. This configuration requires shared storage for the clustered disk resources.

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

- Two virtual machines on two different ESXi hosts run WSFC.
- 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.

Note In a cluster of virtual machines across physical hosts, shared disks can be physical mode or vVol with vSphere 6.7 or later. You can use VMFS VMDKs with vSphere 7.0 or later. If you are using RDMs or vVol as shared disks, they can be on Fibre Channel (FC) SAN, FCoE or iSCSI. If you are using VMFS VMDKs as shared disks, they must be on FC SAN only. A clustered 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). For supported guest operating systems see [Table 4-2. Other Clustering Requirements and Recommendations](#).

Note Private and Public network interconnect can share a single virtual NIC in a VM.

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

Clustering Physical Machines with Virtual Machines

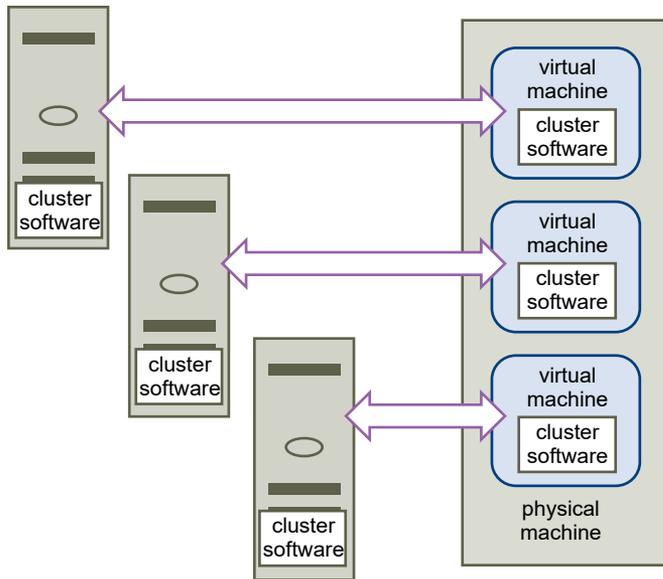
For a simple WSFC clustering solution with low hardware requirements with RDMs, you might choose to have one standby host.

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 ESXi host. Each virtual machine is running WSFC.

Note This solution is not applicable for vVol and clustered VMDKs.

Figure 1-3. Clustering Physical and Virtual Machines



Hardware and Software Requirements for WSFC on vSphere

All vSphere WSFC configurations require certain hardware and software components.

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

Table 1-1. Clustering Requirements

Component	Requirement
Virtual SCSI adapter	LSI Logic SAS for Windows Server 2008 SP2 and above. VMware Paravirtual for Windows Server 2008 SP2 and above. (recommended)
Operating system	Windows Server 2008 SP2 and above releases. For supported guest operating systems see Table 4-2. Other Clustering Requirements and Recommendations.
Virtual NIC	VMXNET3 is required. You must enable 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.

Table 1-1. Clustering Requirements (continued)

Component	Requirement
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	pRDM (physical compatibility mode) vVol disks can be thin or thick format VMFS VMDKs with vSphere 7.0 eagerzeroedthick format
Disk and networking setup	Add networking before disks.
Number of nodes	Windows Server 2008 SP2 and above: up to five-node clustering For supported guest operating systems see Table 4-2. Other Clustering Requirements and Recommendations .
NTP server	Synchronize Windows Active Directory domain controllers and WSFC nodes with a common NTP server, and disable host-based time synchronization. See KB 1189 for more information.

The following table lists requirements that apply to Clustered VMDK as a disk resource for a WSFC.

Table 1-2. Additional Requirements for Clustered VMDK WSFC

Component	Requirement
Windows Cluster Parameters	Set the Windows Cluster Parameter QuorumArbitrationTimeMax to 60
Physical disk /Storage Array LUN	<ul style="list-style-type: none"> ■ The physical disk that virtual disks are stored on must support ATS SCSI ■ The physical disk must also support SCSI-3 Persistent Reservations, specifically Write Exclusive All Registrants (WEAR) type reservations
VMFS Datastore	<ul style="list-style-type: none"> ■ VMFS version must be VMFS6 ■ Fibre Channel (FC) only ■ The datastore must have attribute ClusteredVMDK set. See Enabling Clustered VMDK Support.

Supported Shared Storage Configurations

Different WSFC 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.

Differences between Physical compatibility RDMs and Virtual compatibility RDMs

An RDM (Raw Device Mapped) is a special mapping file in a VMFS volume that manages metadata for its mapped device. The mapping file is presented to the management software as an ordinary disk file, available for file-system operations. To the virtual machine, the storage virtualization layer presents the mapped device as a virtual SCSI device. There are two types of compatibility modes for RDMs.

- A Physical compatibility RDM (pRDM), specifies minimal SCSI virtualization of the mapped device.
 - Virtual Machine Snapshots are not available when the RDM is used in physical compatibility mode.
 - RDMs in physical compatibility mode are required to allow SCSI commands be directly passed to a LUN to satisfy the requirements of SCSI-3 Persistent Reservations used by WSFC.
- A Virtual compatibility RDM (vRDM), specifies full virtualization of the mapped device.
 - VMkernel sends only READ and WRITE to the mapped device. The mapped device appears to the guest operating system the same as a virtual disk file in a VMFS volume.
 - If you are using a raw disk in virtual mode, you can realize the benefits of VMFS, such as advanced file locking for data protection and snapshots for streamlining development processes.
 - Virtual mode is more portable across storage hardware than physical mode, presenting the same behavior as a virtual disk file.
 - Supported for CIB configuration only.

Note Do not use cluster-in-a-box configuration for a production deployment.

For more information, see [KB 2147661](#).

Table 1-3. 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)
VMDK	No	Yes	No
pRDM (physical compatibility mode)	No	Yes (recommended)	Yes
vRDM (virtual compatibility mode)	Yes (not recommended)	No	No

Note The multi-writer feature must not be used for a clustered disk resource for WSFC.

Use of in guest options for storage such as iSCSI or SMB shares within guest operating systems configured with WSFC, in any configuration supported by Microsoft, is transparent to ESXi hosts and there is no need for explicit support statements from VMware.

Round Robin Path Selection Policy (PSP_RR) Support for disk resources used by WSFC

ESXi 7.0 supports PSP_RR for WSFC.

- ESXi 7.0 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 WSFC cluster in CAB configuration one ESXi host can be configured to use PSP_FIXED and the other ESXi host can use PSP_RR.
- All hosts must be running ESXi 6.0 or later.
- Mixed mode configurations of ESXi 6.0 or later with previous ESXi releases are not supported.

iSCSI Support for WSFC

ESXi 7.0 supports iSCSI storage and up to 5 node WSFC clusters.

- ESXi 7.0 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.
- Clustered VMDKs on VMFS datastore in CAB are not supported over iSCSI.
- No qualification is needed for software iSCSI 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 or later.
- Mixed cluster nodes running FC or FCOE and iSCSI are not supported.
- Mixed mode of iSCSI config is supported.

FCoE Support for WSFC

ESXi 7.0 supports FCoE storage and up to 5 node WSFC clusters.

- ESXi 7.0 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.
- Clustered VMDK is not supported on a datastore connected over FCoE.

- 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.
- Required DRS affinity (CIB) or Anti-affinity (CAB) rules for WSFC virtual machines.
- All hosts must be running FCoE initiators. Mixed cluster nodes running FC and FCoE are not supported.
- Mixed mode FCoE configuration is supported.

vMotion support for WSFC

vSphere 7.0 supports vMotion of a VM hosting a node of a WSFC.

Pre-requisites for vMotion support:

- vMotion is supported only for a cluster of virtual machines across physical hosts (CAB).
- The vMotion network must be a 10Gbps Ethernet link. 1Gbps Ethernet link for vMotion of WSFC virtual machines is not supported.
- vMotion is supported for Windows Server 2008 SP2 and above releases. Windows Server 2003 is not supported.
- The WSFC cluster heartbeat time-out must be modified at least to the values listed below:
 - `(get-cluster -name <cluster-name>).SameSubnetThreshold = 10`
 - `(get-cluster -name <cluster-name>).CrossSubnetThreshold = 20`
 - `(get-cluster -name <cluster-name>).RouteHistoryLength = 40`
- The virtual hardware version for the WSFC virtual machine must be version 11 and later.

Modifying the WSFC heartbeat time-out:

WSFC 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 in Windows 2008 R2, 2012, 2012 R2. Other nodes that are members of the cluster will take over any clustered roles that were running on the removed node.

An WSFC 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 at least 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 at least 10. From any one of the participating WSFC cluster nodes run the following command:

```
(get-cluster -name <cluster-name>).SameSubnetThreshold = 10
```

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 20.

Note Recommended values for WSFC hearbeats settings are now defaults in Windows Server 2016 and above.

VMware vSphere® Virtual Volumes™ Support for WSFC

ESXi 7.0 supports VMware vSphere® Virtual Volumes™ storage and up to 5 node WSFC clusters. The storage array must support SCSI3-PRs at the subsidiary LUN level.

- ESXi 7.0 supports vVols Storage for Windows Server 2008 SP2 and above releases.
- Only Cluster-across-box (CAB) is supported.
- WSFC 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 pRDM.
- WSFC on vVols supports vSphere HA, DRS and vMotion.
- The underlying transport protocol can be FC, SCSI 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 WSFC Setup Limitations

Before you set up WSFC in a virtual environment, 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 WSFC setups with this release of vSphere:

- Using VMDKs on NFS datastore as a shared disk resource for WSFC.
- Mixed environments, such as configurations where one cluster node is running a different version of ESXi than another cluster node.
- vSphere Fault Tolerance (FT).
- Migration with vMotion® of clustered virtual machines using vRDMs.

- N-Port ID Virtualization (NPIV)
- ESXi hosts or clusters that have memory overcommitment are not suitable for deploying WSFC virtual machines. Memory overcommitment can cause virtual machines to stall for short durations. This can be significantly disruptive as the WSFC clustering mechanism is time-sensitive and timing delays can cause the virtual machines to behave incorrectly.
- Suspend or resume of a VM, node of a WSFC.
- Storage spaces are not supported with Failover clustering on Windows 2012 and above.
- VM configurations changes:
 - Hot adding memory.
 - Hot adding CPU.
 - Increasing the size of a shared disk.
 - Pausing or resuming the virtual machine state. This I/O intensive operation is disruptive of the timing sensitive WSFC clustering software.
- Sharing disks between virtual machines without a clustering solution may lead to data corruptions

WSFC and Booting from a SAN

You can boot a VM from a SAN, node of a WSFC.

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 Failover Clustering for Windows Server 2008 guest operating systems.
- Test clustered configurations in different failover scenarios before you put them into production environments.

Cluster Virtual Machines Across Physical ESXi Hosts

2

You can create a WSFC cluster that consists of two or more virtual machines on two ESXi or more hosts. Use this method for production deployments.

A cluster across physical hosts requires specific hardware and software.

- Up to 5 nodes in a single WSFC is supported
- Supported shared storage configuration, such as RDM in physical compatibility mode, vVol, clustered VMDK. For more information, see [Supported Shared Storage Configurations](#).

Clustered VMDK support for WSFC

vSphere 7.0 introduces support for the use of VMDKs on clustered datastore as shared disk resources for a WSFC. Using VMDKs reduces the extra overhead to manage the virtual disks compared to pRDMs.

Microsoft Clustering Service uses SCSI-3 PRs commands to coordinate access to a clustered disk resource. These commands (PR-IN and PR-Out) are emulated at VSCSI layer on a datastore. The feature requires support from the datastore perspective. A datastore configured to host clustered VMDKs is referred to as a clustered VMDK datastore in this document.

Enabling Clustered VMDK Support

You can enable clustered VMDK support when you create a new VMFS6 datastore, or enable clustered VMDK on an existing VMFS6 datastore. Before enabling clustered VMDK support, ensure all hosts connected to the datastore are using ESXi 7.0 or later and are managed by vCenter 7.0 or later. All hosts connected to the datastore must be managed by the same vCenter.

Procedure

- 1 Navigate to any inventory object that is a valid parent object of a datastore, such as a host, a cluster, or a data center, and click the **Datastores** tab.
- 2 Right-click the parent object and select **Create a New Datastore**.
- 3 On the **New Datastore** page choose a device with the attribute **Clustered VMDK Supported** set to **Yes**. Create a VMFS6 version Datastore.
- 4 Once the Datastore is created go to the **Configure Datastore** tab.

- 5 Under **General** go to the **Datastore Capabilities** section. Click the **Enable** button for **Clustered VMDK**.
- 6 Once enabled, you can use the **Disable** button to disable clustered VMDK. Disable will not work if there are any WSFC VMs using clustered VMDKs on the datastore. All virtual machines with shared disks on the datastore must be powered off.

Configuration Maximums for Clustered VMDK Support

Configuration Maximums for Clustered VMDK support.

Configuration	Maximum
Max number of VMs in a single WSFC cluster	5
Max number of WSFC clusters (multi-cluster) running on the same set of ESXi hosts	3
Max number of clustered VMDKs per ESXi host	128

Recommendations for using Clustered VMDKs with WSFC

Follow these recommendations when using Clustered VMDKs with WSFC.

- 1 Do not present LUNs used for clustered VMDKs to ESXi hosts unless the host is configured with ESXi 7.0. This might cause slow boot times, hostd to become unresponsive and other issues. A host with a version lower than ESXi 7.0 cannot mount a clustered VMDK datastore. This is because the ESXi hosts on which WSFC VMs run must have physical SCSI-3 WEAR type reservation on the LUN.
- 2 Make sure that all VMs hosting nodes of WSFC are migrated off or powered off properly before removing it from a clustered VMDK datastore to ensure the resources, like heartbeat (HB) slots, are freed. If a VM fails or there is an APD on the clustered VMDK datastore during power-off, always power on and power off before removing the VM from a cluster.
- 3 Do not combine clustered and non-clustered VMDKs on the same clustered datastore. Although VMs that are using non-shared disks on a clustered datastore will continue to work normally and support all operations such as snapshot and clone.
- 4 Do not keep clustered VMDKs for different clusters on the same shared datastore. Use a different clustered datastore for different WSFC clusters.
- 5 Set vHardware (virtual compatibility) to vSphere 7.0 while using the clustered VMDK feature.

Requirements for using Clustered VMDKs with WSFC

You must follow these requirements when using Clustered VMDKs with WSFC.

- 1 You must set the Windows Cluster Parameter **QuorumArbitrationTimeMax** to 60.
- 2 The physical disk that virtual disks are stored on must support ATS SCSI commands.
- 3 The physical disk must support SCSI-3 Persistent Reservations, specifically Write Exclusive All Registrants (WEAR) type reservations.

- 4 VMDKs must be Eager Zeroed Thick Provisioned.

Limitations of Clustered VMDK support for WSFC

Clustered VMDK support is only possible with the following limitations:

- 1 Cluster in a Box (CIB) configuration is not supported.
- 2 Spanned VMFS datastores cannot store clustered VMDKs.
- 3 A physical LUN with a clustered datastore cannot have any additional VMFS datastores.
- 4 On-line Storage vMotion and snapshots are not supported.
- 5 Only disks with 512 sector size are supported.
- 6 VMFS5 and earlier are not supported for storing clustered VMDKs.
- 7 Only datastores connected through Fibre Channel (FC) are supported in this release.
- 8 Clustered VMDK datastore cannot be used as an ESXi installation disk.
- 9 Diagnostic partition is not supported on disks that are marked clustered.
- 10 SCSI-2 reservations are not supported on clustered VMDKs.
- 11 Hot expansion of a VMDK that is associated with a clustered VM is not supported.
- 12 A mix of clustered VMDKs with other types of clustered disks such as pRDMs or vVol in the same WSFC is not supported.
- 13 Moving more than one node of WSFC to the same ESXi host is not allowed and causes vMotion to fail. You must use anti-affinity DRS Rules to separate VMs and nodes of the cluster on different ESXi hosts.
- 14 Cloning and Fault Tolerance are not supported.

Create the First Node for WSFC Clusters Across Physical Hosts

To create the first node, you create and configure a virtual machine and install a guest operating system on 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 if required or applicable.
Ready to Complete	Review your selections.

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

Note Do not add shared cluster disks at this time.

- 5 Install a Windows Server operating system on the virtual machine.

Create Additional Nodes for a WSFC Across Physical Hosts

To create additional nodes in a cluster of virtual machines across physical hosts, you can create a template from the the first virtual machine and use it to deploy additional virtual machines or use the procedure outlined previously to manually create additional VMs 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.

Caution If you clone a virtual machine with clustered VMDKs, unmap all clustered VMDKs before cloning, and remap them after cloning is complete.

Caution Use Windows sysprep operations before cloning to a template.

Procedure

- 1 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 .
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.

- Right-click the virtual machine template and select **Deploy VM from this Template**.
- 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 .

- Select a new guest operating system from the list.
 - Click the **Create a new specification** button to add a new guest operating system. Proceed through the **New VM Guest Customization Spec** wizard.
 - Click **Finish** to exit the wizard.
- Click **Finish** to deploy the virtual machine.

Add Hard Disks to the First VM, first node of a WSFC, for Clusters Across Physical Hosts with pRDMs

In an WSFC cluster with shared disk, storage disks are shared between nodes.

Prerequisites

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

- If you are using physical compatibility mode RDMs, 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 vVols storage for shared disk, please see *Add Hard Disks to the First Node for Clusters Across Physical Hosts with VMware vSphere® Virtual Volumes™*.

- If you are using Clustered VMDKs on VMFS datastores, please see *Add Hard Disks to the First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS datastores*.

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, to add a storage controller select **VMware Paravirtual**, and click **Add**.

Note LSI SAS is supported, but VMware Paravirtual is recommended.

- 3 Click the **New device** drop-down menu, select **RDM Disk**, and click **Add**.
- 4 Select an unformatted LUN.
- 5 Select a datastore.

This datastore must be on shared storage because you need a single shared RDM file for each shared disk.

- 6 Select **Physical** as the compatibility mode.
- 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 **OK** to complete creating the disk.
The wizard creates a new hard disk.
- 9 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 disks are added to the VM, a node of a WSFC, and prepared to be clustered across additional nodes.

Add Hard Disks to the First Node of a WSFC Across Physical Hosts with vVols

In an WSFC cluster, storage disks are shared between nodes.

Prerequisites

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

- Use HW Version 13 or later. See [Virtual Volumes \(vVols\) now supports WSFC](#).
- Prepare vVol deployment. See *vSphere Storage* documentation for more information.

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 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 4-2. 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 disks based on vVols are added to the VM, a node of a WSFC, and prepared to be clustered across additional nodes.

Add Hard Disks to the First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS Datastores

Add hard disks to the first node for cluster across physical hosts with clustered VMDKs on VMFS datastores.

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, to add a storage controller select **VMware Paravirtual**, and click **Add**.

Note LSI SAS is supported, but VMware Paravirtual is recommended.

- 3 Click the **New device** drop-down menu, select **New Hard Disk**, and click **Add**.
- 4 Select the disk size.
- 5 Under Disk Provisioning, select **eagerzeroed thick**.
- 6 Expand the New Hard Disk.
- 7 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.

- 8 Click **OK**.

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

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

- 10 Select the appropriate type of controller.

- LSI Logic SAS
- VMware Paravirtual

For supported guest operating systems see [Table 4-2. Other Clustering Requirements and Recommendations](#).

- 11 Click **OK**.

- 12 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**.

Add Hard Disks to Additional Nodes for Clusters Across Physical Hosts

To allow shared access to disk resources, point to existing disks on the VM, the first node of a WSFC cluster. Use the same SCSI IDs while assigning disks to all additional nodes.

Prerequisites

Before you begin, obtain the following information:

- Obtain SCSI IDs for all virtual disks to be shared.
- Obtain disk file path on datastore for all shared disks.

Note The procedure for clustered VMDK, pRDMs and vVol is the same.

Procedure

- 1 In the vSphere Client, select the newly created virtual machine, a second node of the WSFC, right-click and select **Edit Settings**.
- 2 Click the **New device** drop-down menu, to add a storage controller select **VMware Paravirtual**, and click **Add**.

Note LSI SAS is supported, but VMware Paravirtual is recommended.

- 3 Click the **New device** drop-down menu, select **Existing Hard Disk**, and click **Add**.
- 4 In Disk File Path, browse to the location of the disk to be shared specified for the first node.
- 5 Select **Physical** as the compatibility mode and click **Next**.
- 6 Select the same SCSI ID you chose for the first virtual machine's shared storage disks (for example, **SCSI (1:0)**), and click **OK**.

Note The disk SCSI ID for this virtual machine's shared storage must match the corresponding SCSI ID for the first virtual machine.

- 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 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 4-2. Other Clustering Requirements and Recommendations](#).

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

Use WSFC in an vSphere HA and vSphere DRS Environment

3

When you use WSFC 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 WSFC virtual machines must be managed by a vCenter Server system.

VMware vSphere 7.0 provides support for vMotion of a VM, node of a WSFC.

Enable vSphere HA and vSphere DRS in ESXi

ESXi hosts that are running WSFC virtual machines can be part of a vSphere 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 WSFC Virtual Machines

For WSFC 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 WSFC 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 WSFC virtual machines across physical hosts).

For a cluster of virtual machines on one physical host, use affinity rules.

Note Do not use a cluster on one physical host for production deployment.

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

Note vMotion is supported for a virtual machines, nodes of a WSFC, across physical ESXi hosts.

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 ESXi host, select **Keep Virtual Machines Together**.
 - For a cluster of virtual machines across physical ESXi 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**.

Set DRS Automation Level for WSFC Virtual Machines

Set the automation level of all virtual machines in an WSFC 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 WSFC 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 WSFC virtual machines, but there is no technical restriction which prevents the setting for WSFC virtual machines to be fully automated. If the fully automated setting is used please tune WSFC advanced settings to prevent WSFC failover during vMotion.

Using vSphere DRS Groups and VM-Host Affinity Rules with WSFC 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 WSFC 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 WSFC 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 ESXi host. (This does not apply to clusters of virtual machines across physical ESXi 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 (WSFC)

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 WSFC virtual machines. For example, VMGroup_1 contains WSFC_VM_1, WSFC_VM_2 ... WSFC_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 WSFC virtual machines to one group.
 - For a cluster of virtual machines across physical hosts, add all WSFC virtual machines to one group.
- 7 Click **OK**.

Create a Host DRS Group (WSFC)

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 WSFC_VM_1, ESXi_HOST_2 running WSFC_VM_2 ESXi_HOST_5 running WSFC_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 (WSFC)

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 WSFC virtual machines as described in [Create a Virtual Machine DRS Group \(WSFC\)](#).

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

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 WSFC Setup Checklist

4

When you set up a WSFC on VMware vSphere, 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 4-1. Requirements for Clustered Disks

Component	Single-Host Clustering	Multihost Clustering
VMDK	SCSI bus sharing mode set to virtual is supported. (Not recommended.)	SCSI bus sharing mode set to physical is supported.
vRDM	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.	Not supported.
pRDM	Not supported.	RDM device type must be set to Physical compatibility mode during hard disk creation. SCSI bus sharing mode of vSCSI Controller hosting pRDM must be set to physical. A single, shared RDM mapping file for each clustered disk is required.
vVols storage	Not Supported.	SCSI bus sharing mode of vSCSI controller must be set to physical.
All types	All clustered nodes must use the same target SCSI ID (on the virtual SCSI adapter) for the same clustered disk. A separate virtual SCSI adapter must be used for clustered disks. Multiple virtual SCSI adapters for clustered disks (up to supported maximum) can be used to enhance performance. In this setup at least one controller should be dedicated to non-shared disk.	

Other Requirements and Recommendations

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

Table 4-2. Other Clustering Requirements and Recommendations

Component	Requirement
Disk	<p>Select Thick Provision during disk provisioning.</p> <p>Modify the .vmx file for the WSFC VMs to add these lines:</p> <ul style="list-style-type: none"> ■ <code>scsiX.returnNoConnectDuringAPD = "TRUE"</code> ■ <code>scsiX.returnBusyOnNoConnectStatus = "FALSE"</code> ■ Where X is the boot device SCSI bus controller ID number. By default, X is set to 0. <p>The only disks that you should not create with the Thick Provision option are RDM files (both physical and virtual compatibility mode).</p> <p>Disk I/O timeout is 60 seconds or more (HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeoutValue).</p>
Windows	<p>Use:</p> <ul style="list-style-type: none"> ■ 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 ■ Windows Server 2016 ■ Windows Server 2019 with vSphere 6.7 Update 3 <p>For Windows Server 2008 SP2 and above, you can use up to five cluster nodes.</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>
ESXi	<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>ESXi uses a different technique to determine if Raw Device Mapped (RDM) LUNs are used for WSFC cluster devices, by introducing a configuration flag to mark each device as "perennially reserved" that is participating in an WSFC cluster. For ESXi hosts hosting passive WSFC nodes with RDM LUNs, use the <code>esxcli</code> command to mark the device as perennially reserved:</p> <p><code>esxcli storage core device setconfig -d <naa.id> --perennially-reserved=true</code>. See KB 1016106 for more information.</p>
Multipathing	<p>Contact your multipathing software vendor for information and support of non-VMware multipathing software in vSphere.</p>

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 issues with a WSFC setup.

Note Verify that your setup complies with the checklists in [Table 4-1. Requirements for Clustered Disks](#) and [Table 4-2. Other Clustering Requirements and Recommendations](#) before you contact technical support.

Table 4-3. 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 or use Get-ClusterLog PowerShell Cmdlet
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).

Cluster Virtual Machines on One Physical Host

5

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

Note For production deployments, do not set up a WSFC on VMs residing 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 4-2. Other Clustering Requirements and Recommendations](#).

Create the First Node of a WSFC on One Physical ESXi Host

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

Page	Action
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.

- 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 or later.

Note For a clean template, use Windows Sysprep.

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 .
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.

- Right-click the virtual machine template and select **Deploy VM from this Template**.
- 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 .

Page	Action
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 WSFC cluster, storage disks are shared between nodes. You set up a quorum disk and at least one 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 2008 SP2 and above	LSI Logic SAS and VMware Paravirtual

For supported guest operating systems see [Table 4-2. 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, use **Add an Existing disk** on the second VM to point to the disks from the first VM that should be shared. Use the same SCSI IDs.

Prerequisites

Before you begin, obtain the following information:

- Obtain SCSI IDs for all virtual disks to be shared.

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 Physical and Virtual Machines

6

You can create an WSFC cluster in which each physical server 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.

- Network configuration should be consistent between a physical and virtual node of a WSFC. Ensure that all required VLANs are available to ESXi hosts, hosting a virtual node of a WSFC cluster.
- Use RDMs in physical compatibility mode. You cannot use virtual disks or RDMs in virtual compatibility mode 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 WSFC cluster, see the Failover Cluster for Windows 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.
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, a second node of a WSFC cluster, where the first node is a physical server.

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 ESXi host.

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 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 shared storage disks for the first node. The setup allows shared access to clustered disks.

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 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

After you set up the first and second nodes, you must configure Microsoft Failover Clustering.

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 might cause a problem when you add a node to the cluster.

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.