Cluster Volume Manager

Improved storage management with Symantec Storage Foundation 6.1 for Windows

Table of Contents

What’s improved with storage management using SFW?

Symantec Storage Foundation for Windows provides a comprehensive solution for storage management. It lets you optimize your storage through the use of dynamic disk groups and dynamic volumes.

Context

In a clustered environment, at any given point of time, a dynamic disk group (VMDG) created on the shared disks is imported on any one cluster node and remains deported on the other cluster nodes.

In this model, in the event of a failover or as part of a planned move, the unit of failover is a disk group. The complete disk group needs to be deported from an online cluster node and imported on the target node.

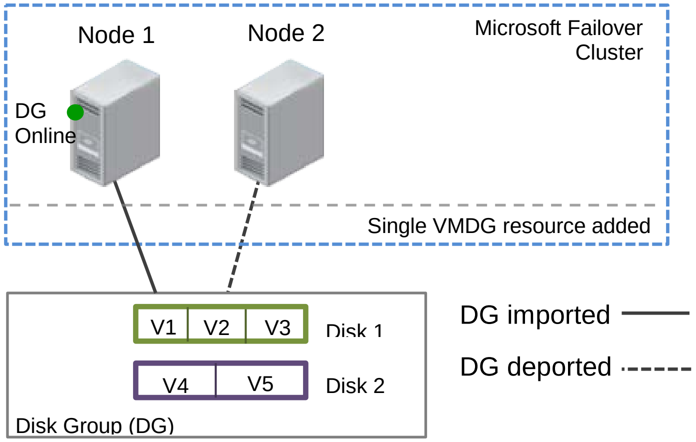

What’s new?

With the 6.1 release, SFW introduces support for volume-level failover. SFW now includes the Cluster Volume Manager (CVM) component.

You can use CVM to create a cluster shared disk group (CSDG) in a Microsoft Failover Cluster environment. A CSDG allows you to share volume configurations and enable failover capabilities at volume level.

Note: CVM is supported only in a Microsoft Hyper-V environment. It is not supported for a physical environment.

The cluster nodes can simultaneously access and manage the set of disks in a CSDG. The logical view of disk and volume configuration is available on all the nodes. Even though the disk group is shared, only one node in a given cluster is granted access to a volume. The same volume cannot be accessed from different nodes. This volume access mode is used to avoid data corruption.

When a volume is brought online on a node, then that node is given access to the volume, while the rest of the nodes cannot access the volume. To give another node access to the same volume, you must first take the volume offline on the node where it is online, and then bring it online on the other node.

In this model, in the event of a failover or as part of a planned move, the unit of failover is the volume. You do not need to deport the entire disk group from an online cluster node.

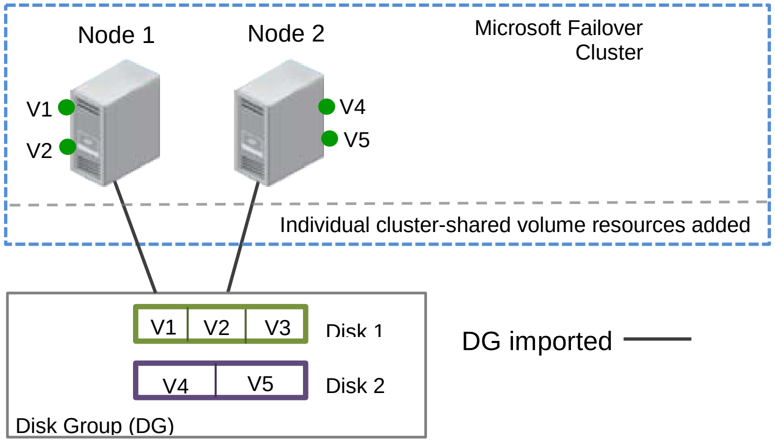
Node 1
Node 2
Microsoft Failover Cluster
V1
V2
V4
V5
Individual cluster-shared volume resources added
V1
V2
V3
Disk 1
V4
V5
Disk 2
Disk Group (DG)
DG imported

Cluster Volume Manager– Detailed insight

The Cluster Volume Manager (CVM) is based on a Master-and-Slave architecture pattern. The node from which the CSDG is configured takes the Master role and the subsequent nodes join the cluster as Slaves. When a slave tries to join a cluster, the master sends it a list of the disk IDs that it has imported, and the slave checks to see if it can access them all. If the slave cannot access one of the listed disks, it abandons its attempt to join the cluster. If it can access all of the listed disks, it joins the cluster and imports the same shared disk groups as the master. When a node leaves the cluster gracefully, it deports all its imported shared disk groups, but they stay imported on the remaining nodes.

In a CVM cluster, the applications on each cluster node can simultaneously read and write data. This simultaneous data access introduces the chances of data corruption. To overcome this issue, all the cluster operations are coordinated through the Master node. The operations are marked as complete only after the required changes are committed across all the Slaves.

Any operation performed through a single node is redirected through the Master node.

The following figure depicts a CVM cluster consisting of four nodes. The cluster nodes access shared disk groups and the objects configured within the disk groups.

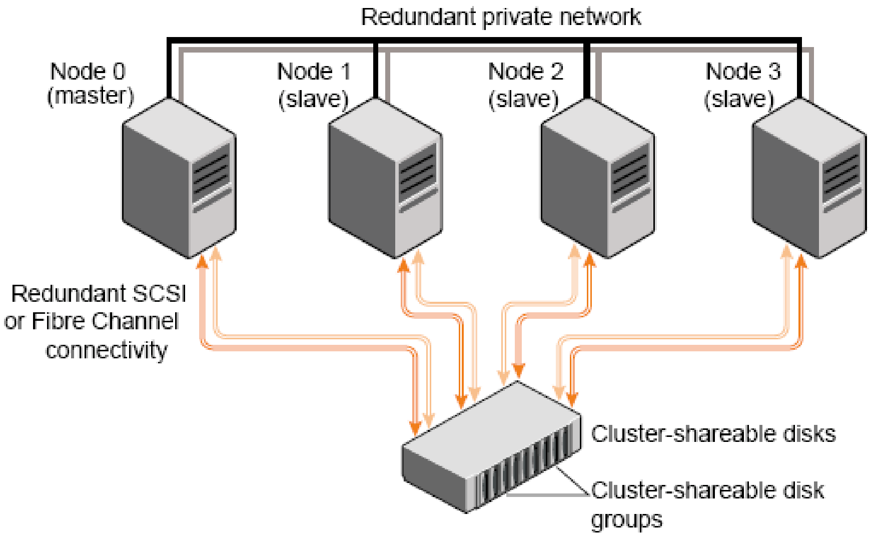
Redundant private network
Node 0 (master)
Node 1 (slave)
Node 2 (slave)
Node 3 (slave)
Redundant SCSI or Fibre Channel connectivity
Cluster-shareable disks
Cluster-shareable disk groups

In this example, node 0 is configured as the CVM master node and nodes 1, 2, and 3 are configured as CVM slave nodes. The nodes are fully connected by a private network and they are also separately connected to shared external storage through SCSI or fibre channel in a storage area network (SAN).

Each node has two independent paths to the disks, which are configured in one or more cluster-shared disk groups. Multiple paths provide resilience against failure of one of the paths, but they are not required for configuring a cluster. The disks may also be connected by single paths.

The private network allows the nodes to share information about system resources and about each other's state. Using the private network, any node can recognize which other nodes are currently active, which are joining or leaving the cluster, and which have failed. The private network requires at least two communication channels to provide redundancy against the failure of one of the channels. If only one channel were used, its failure would be indistinguishable from node failure—a condition known as network partitioning.

How CVM works

CVM functions with minimal disruption if one or more nodes lose connectivity to the shared storage. When CVM detects a loss of storage connectivity for an online disk group, it performs appropriate error handling for the situation. For example, CVM may redirect I/O over the network, detach a plex, or disable a volume for all disks, depending on the situation.

CVM follows the local detach policy to handle failures.

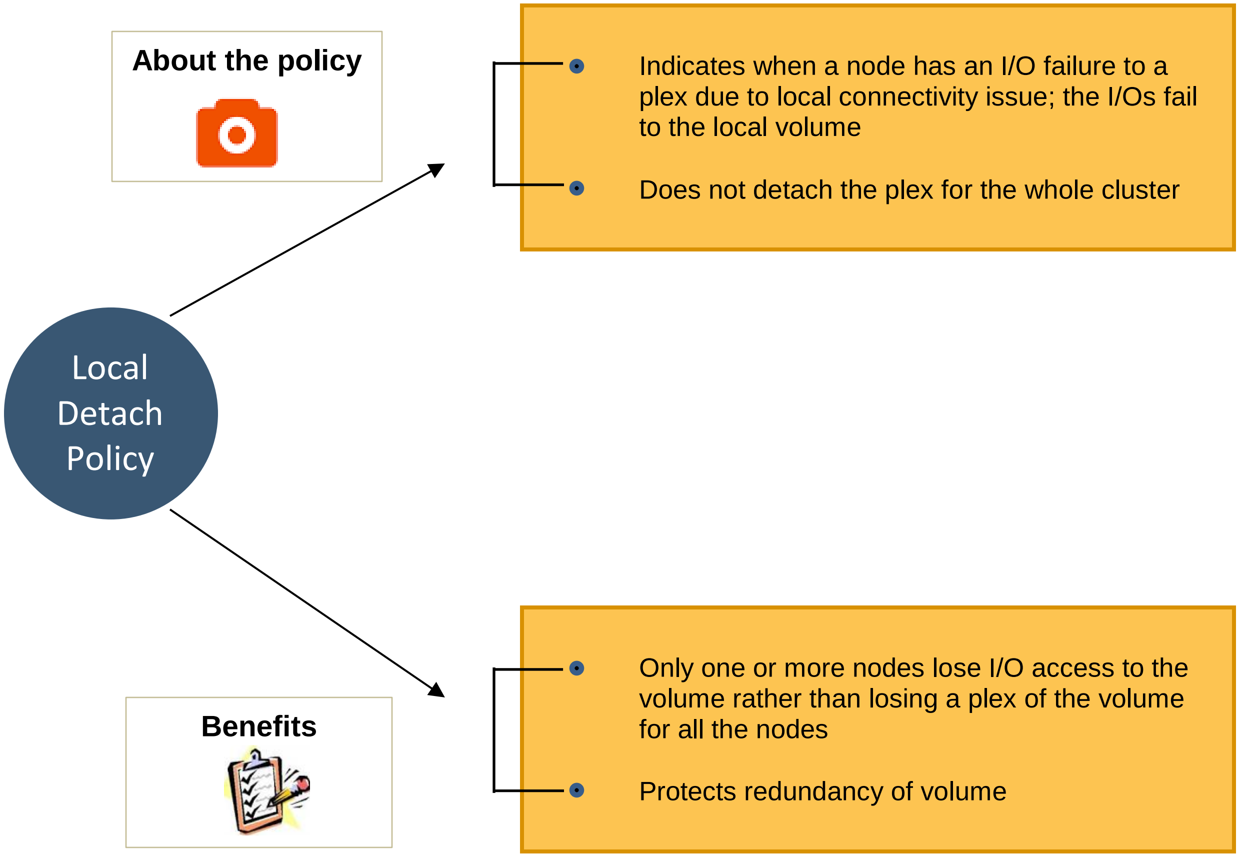

Local Detach Policy

About the policy

- Indicates when a node has an I/O failure to a plex due to local connectivity issue; the I/Os fail to the local volume
- Does not detach the plex for the whole cluster

Benefits

- Only one or more nodes lose I/O access to the volume rather than losing a plex of the volume for all the nodes
- Protects redundancy of volume

Types of storage connectivity issues and CVM behavior

The behavior of CVM in the event of a storage connectivity failure depends on the type of failure that has occurred. The type of storage connectivity failure depends on the scope of failure. CVM determines whether the failure affects all the nodes (global failure) or only some nodes (local failure). CVM also determines whether the failure affects one or more plexes of the volume. If the failure affects all the plexes, it is considered as a total failure. Otherwise, it is considered as a partial failure.

Based on the scope of failure, storage disconnect is classified into the following categories:

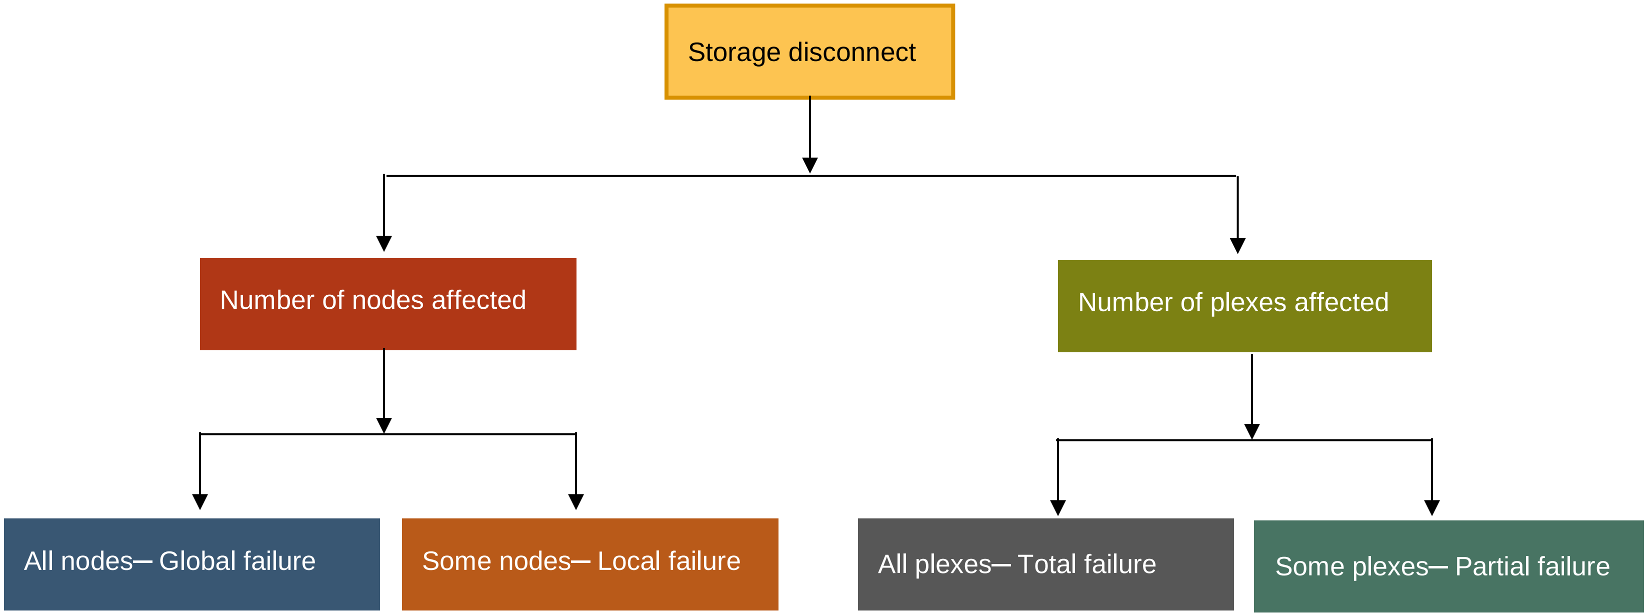
Storage disconnect
Number of nodes affected
Number of plexes affected
All nodes– Global failure
Some nodes– Local failure
All plexes– Total failure
Some plexes– Partial failure

CVM response:

Type of failure	CVM behavior
Global partial failure	Detaches the plex
Global total failure	Disables the volume
Local partial failure	Fails I/O to the volume from the nodes that cannot access the plex
Local total failure	Fails I/O to the volume from the nodes that cannot access the plex

Types of storage connectivity issues and CVM behavior

The following figures depict a storage connectivity failure and how CVM uses the local detach policy to handle the failure:

A Global partial failure

A global partial failure indicates that all the nodes in the cluster are affected, but not all the plexes in the volume.

Storage disconnect

The following figure depicts a global partial failure in which all the nodes in the cluster have lost access to Array B, which has plex B for the volume.

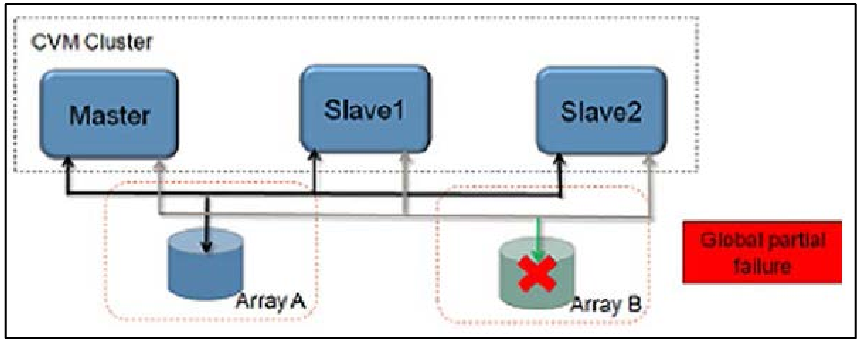

Global Partial Failure

CVM behavior

The following figure depicts how CVM handles the global partial failure using the local detach policy.

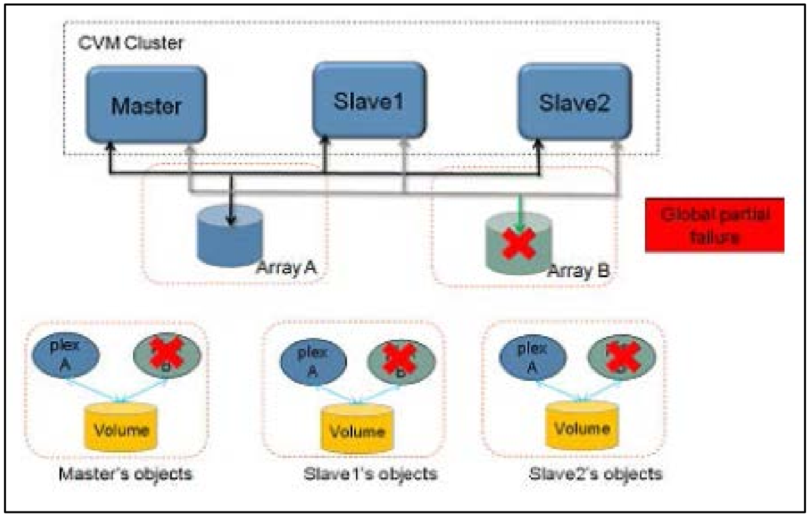

CVM detaches the plex B to maintain the consistency of the mirror. I/O continues to the other plexes in the volume. This reduces the redundancy of the volume.

Types of storage connectivity issues and CVM behavior

The following figures depict a storage connectivity failure and how CVM uses the local detach policy to handle the failure:

B Global total failure

A global total failure indicates that all nodes in a cluster and all the plexes in the volume are affected.

Storage disconnect type

The following figure depicts a global total failure in which all the nodes in the cluster have lost access to Array A.

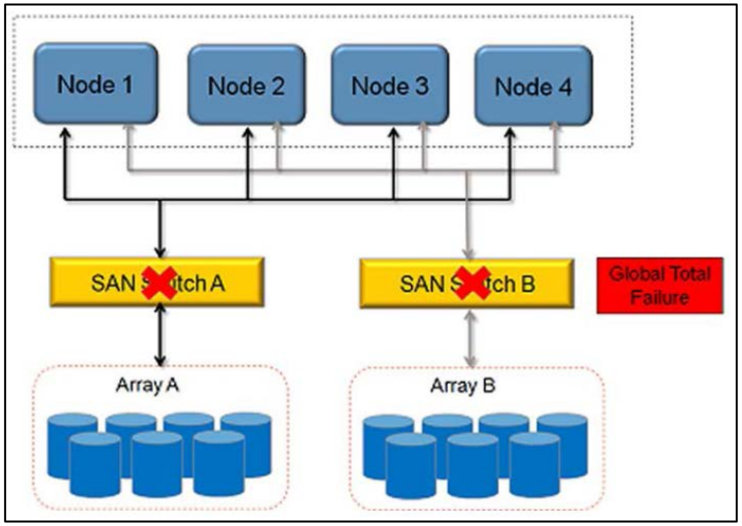

Global Total Failure

CVM response

The following figure depicts how CVM handles the global total failure using the local detach policy.

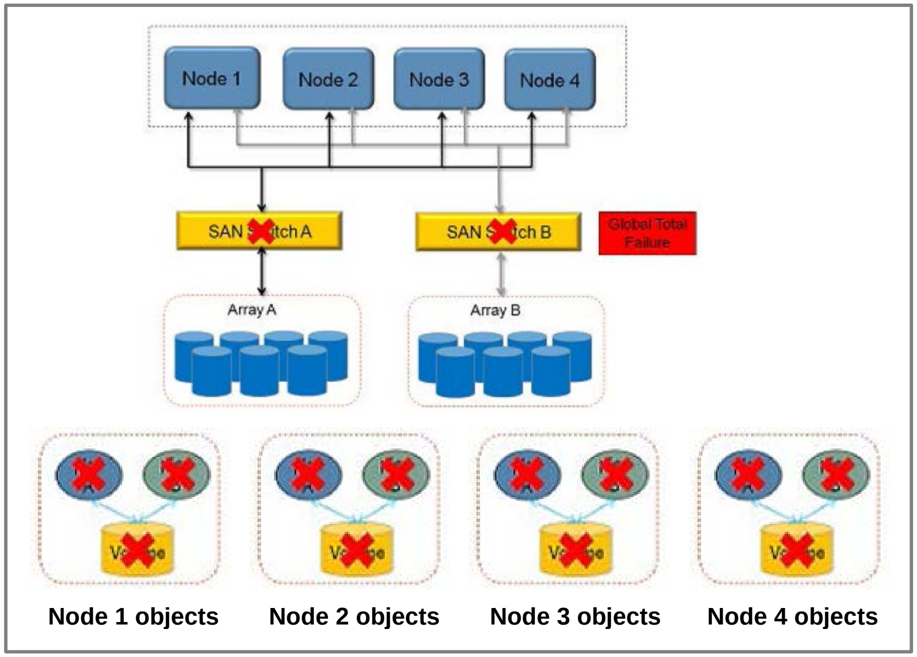

CVM disables the volume. Since no plexes are available, the volume is not available for any I/Os. If the failure occurs on all the nodes at the same time, no plexes are detached.

Types of storage connectivity issues and CVM behavior

The following figures depict a storage connectivity failure and how CVM uses the local detach policy to handle the failure:

C Local partial failure

A local partial failure indicates that some of the nodes in the cluster are affected, but none of the plexes in the volume are affected

Storage disconnect type

The following figure depicts the local partial failure in which one node from the cluster has lost access to Array B.

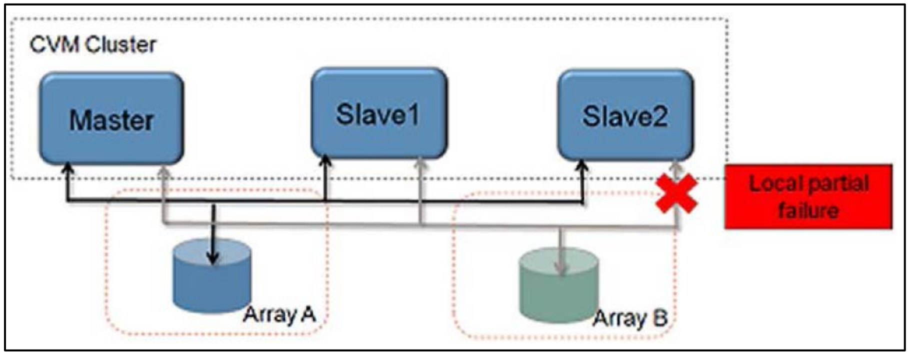

Local Partial Failure

CVM response

The following figure provides information about how CVM handles the global total failure using the local detach policy.

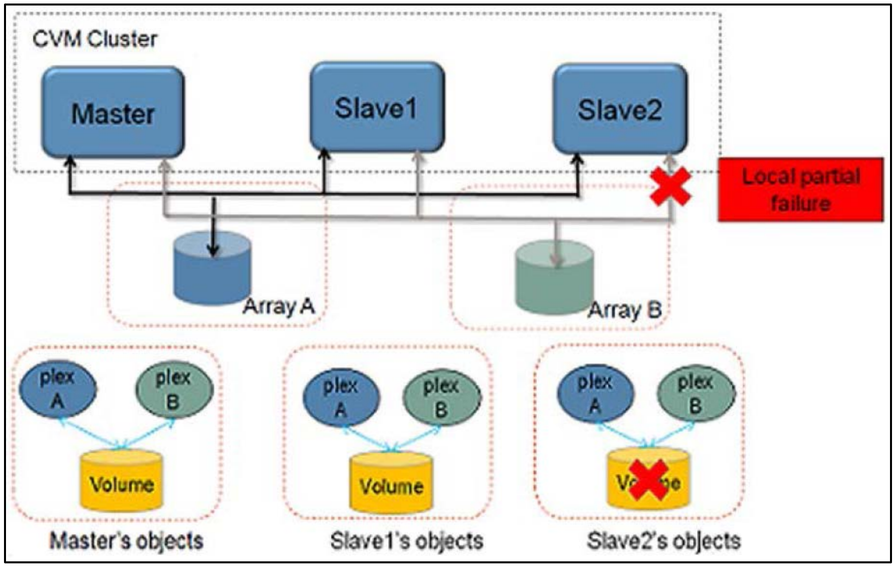

CVM fails to perform I/Os locally to the volume. The local detach policy indicates that CVM should ensure that a local connectivity error only affects the local node. When the local I/O to the volume fails, the applications need to be failed over to another node.

Types of storage connectivity issues and CVM behavior

The following figures depict a storage connectivity failure and how CVM uses the local detach policy to handle the failure:

D Local total failure

A local total failure indicates that all the plexes in the volume are affected, but not all the nodes in the cluster.

Storage disconnect type

The following figure depicts the local total failure in which one node from the cluster has lost access to all plexes.

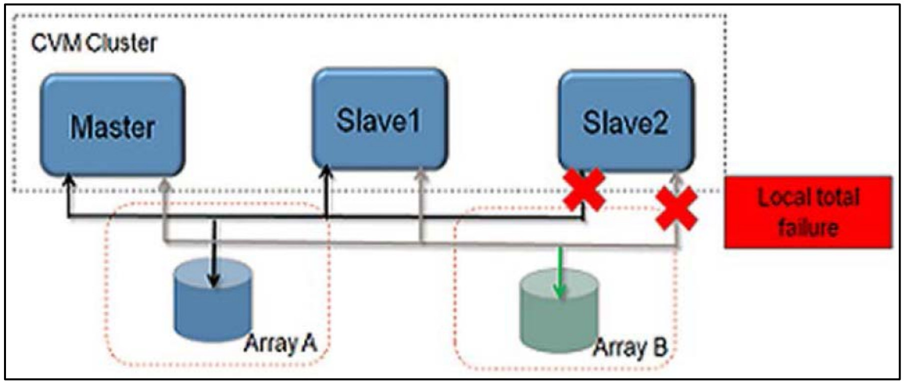

Local Total Failure

CVM response

The following figure provides information about how CVM handles the global total failure using the local detach policy.

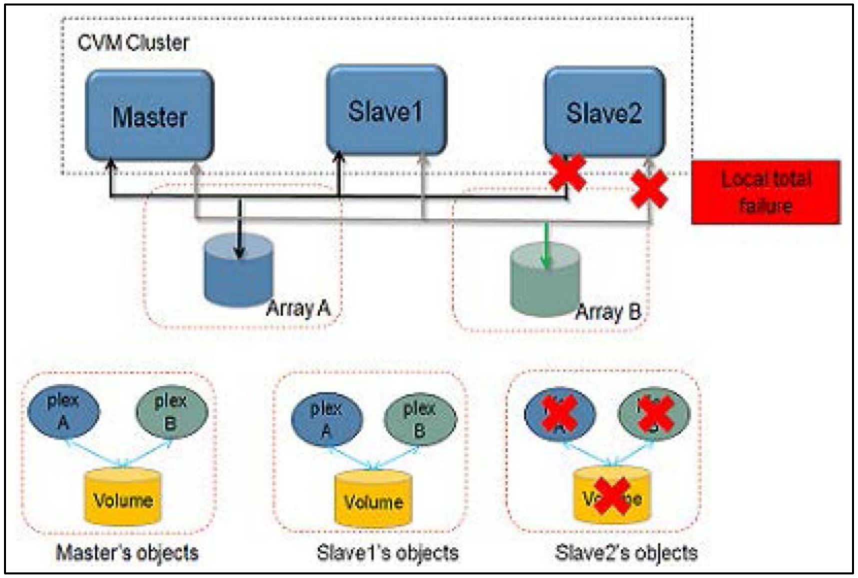

CVM fails to perform I/Os locally on the volume. I/O fails to the volume from the node or nodes that cannot access the plex. I/O continues to the volume from the other nodes.

Configuring a CVM cluster for Microsoft Failover Cluster

Configuring a CVM cluster involves the following tasks:

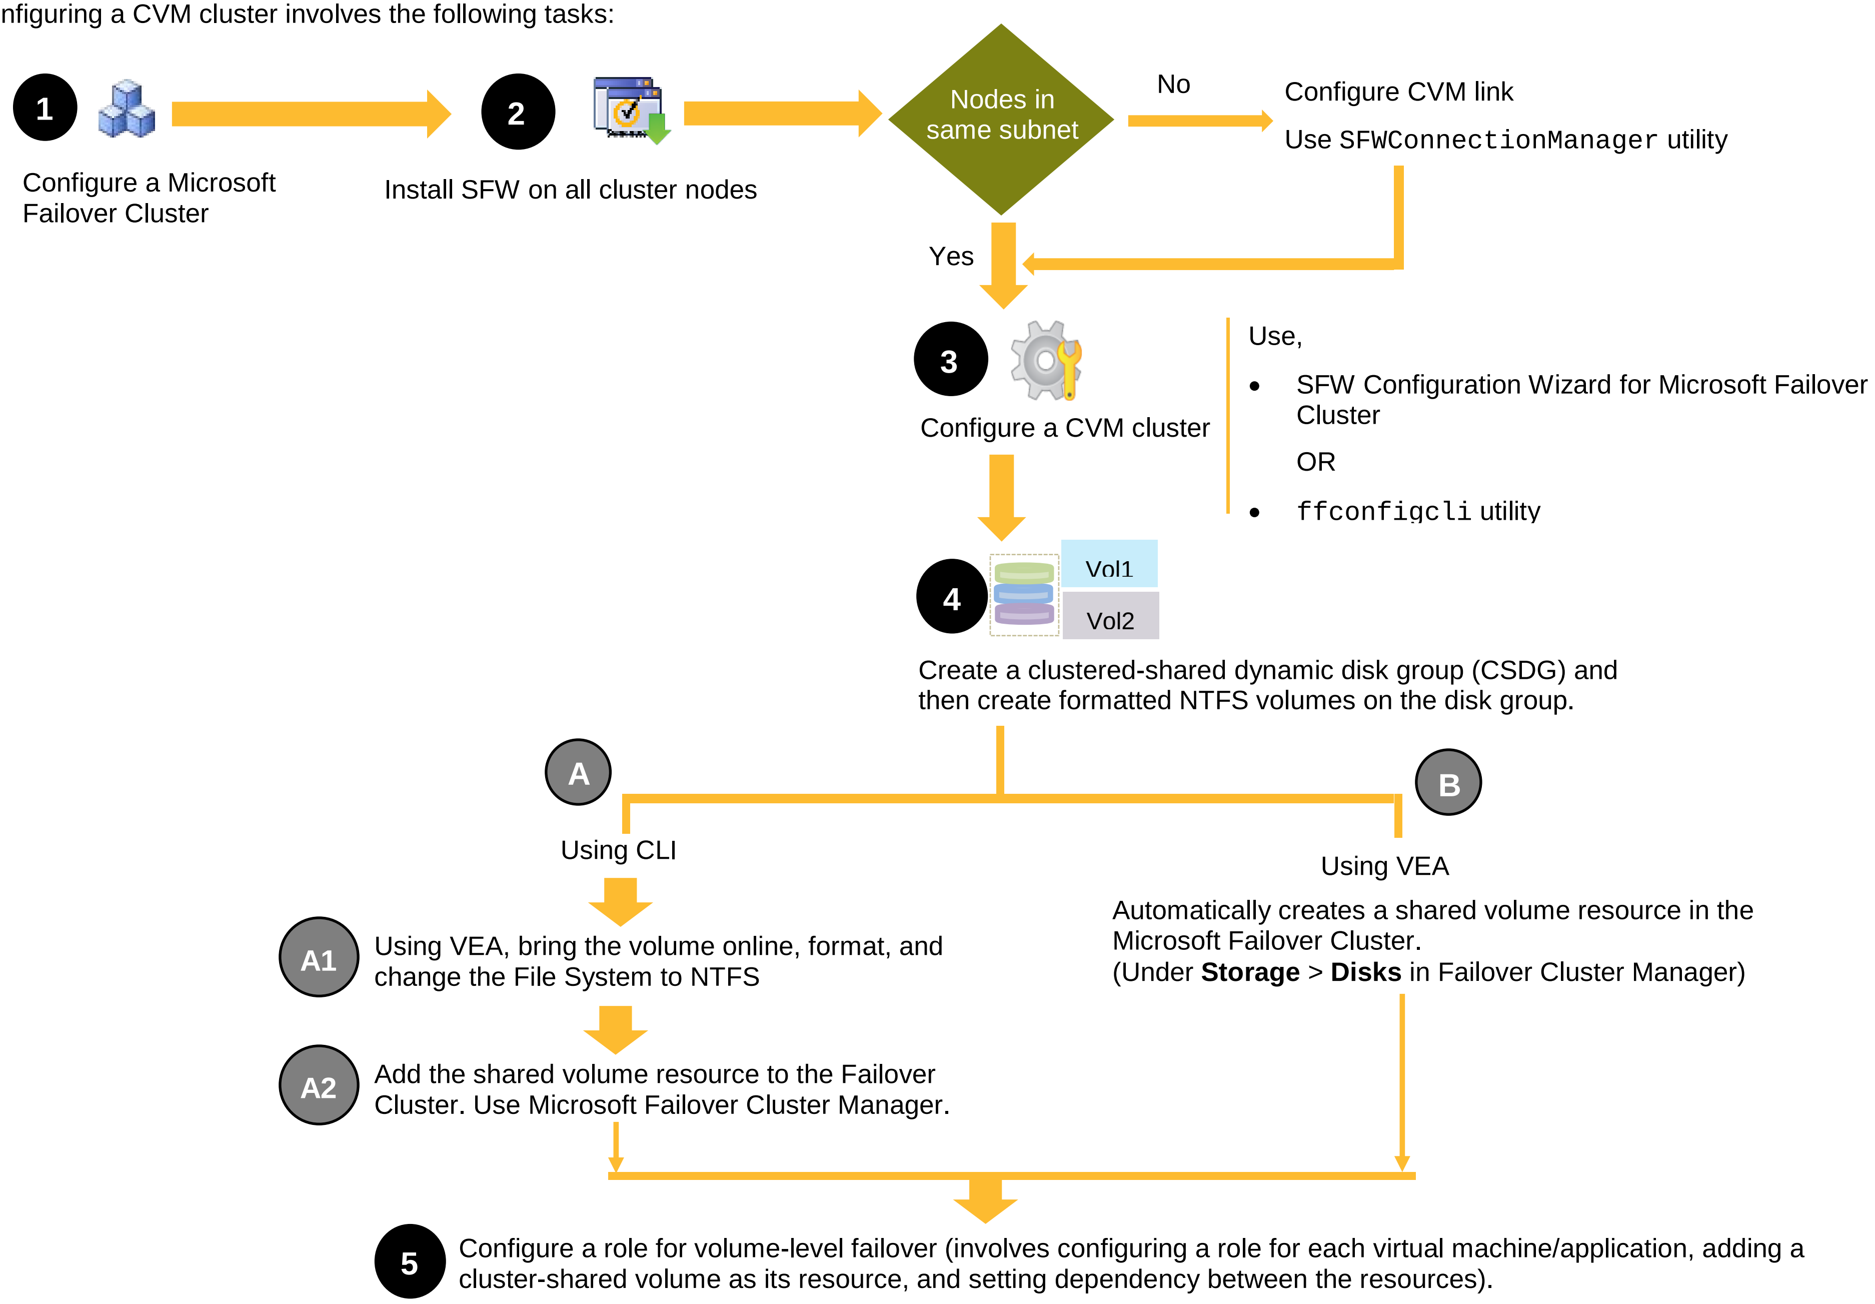

1 Configure a Microsoft Failover Cluster

2 Install SFW on all cluster nodes

Nodes in same subnet

No

Configure CVM link

Use `SFWConnectionManager` utility

Yes

3 Configure a CVM cluster

Use,

- SFW Configuration Wizard for Microsoft Failover Cluster

  OR

- `ffconfigcli` utility

4

Create a clustered-shared dynamic disk group (CSDG) and then create formatted NTFS volumes on the disk group.

A Using CLI

A1 Using VEA, bring the volume online, format, and change the File System to NTFS

A2 Add the shared volume resource to the Failover Cluster. Use Microsoft Failover Cluster Manager.

B Using VEA

Automatically creates a shared volume resource in the Microsoft Failover Cluster.
(Under **Storage** > **Disks** in Failover Cluster Manager)

5 Configure a role for volume-level failover (involves configuring a role for each virtual machine/application, adding a cluster-shared volume as its resource, and setting dependency between the resources).

Pre-configuration notes

Notes

- After you configure a Microsoft Failover Cluster, ensure that the Microsoft failover cluster service is running on all nodes.
- Ensure that you select the **Microsoft Failover Cluster** option while installing SFW.
- You can choose to configure a Microsoft Failover Cluster after installing SFW. If you choose to do so, then you need to manually register the Volume Manager Shared Volume resource.
- You must ensure that the selected drive letter for the new cluster-shared volume is available and not in use on any of the cluster nodes.
- If you make any changes to the cluster network configuration of the Microsoft failover cluster after configuring CVM, then you need to unconfigure the CVM cluster and reconfigure it using either the wizard or the ffconfigcli utility. This requires application downtime.
- You can add a node to a CVM cluster. To successfully add a node to the CVM cluster, the node must be added to the Microsoft failover cluster first.

For details about configuring Failover Cluster, refer to Microsoft documentation.

For details about installing SFW and/or manually registering the Volume Manager Shared Volume resource, refer to the *Storage Foundation and High Availability Solution Installation and Upgrade Guide*.

Configuring CVM links for multi-subnet cluster networks

To successfully create a CVM cluster, the participating cluster nodes must share a common cluster network. However, in case of a cross-subnet configuration, where you have nodes across different cluster networks, you need to use the `SFWConnectionManager` utility to create a CVM link to establish cluster and network communications across the nodes that will form a CVM cluster.

You must create a CVM link using `SFWConnectionManager` before you run the SFW Configuration Wizard for Microsoft Failover Cluster (or its CLI equivalent `ffconfigcli`).

Use the following commands to perform the required operation:

Command	Description
`SFWConnectionManager -createlink <ClusterNetworkName 1>,<ClusterNetworkName 2>,...<ClusterNetworkName n>`	Creates a CVM link across cluster networks **Note**: If a cluster network's name has spaces, then it must be specified within double quotes.
`SFWConnectionManager -deletelink <CVMLinkName>`	Deletes an existing CVM link across cluster networks In place of `<CVMLinkName>`, use the -all option at the end of the command to delete all the CVM links.
`SFWConnectionManager -displaylinks`	Displays information about the CVM links and their cluster networks

Note

If you make any changes to the existing CVM link of the CVM cluster using `SFWConnectionManager` after running the wizard (or `ffconfigcli`), then you need to unconfigure and re-configure CVM cluster for the changes to be reflected.

Configuring a CVM cluster using the configuration wizard

Use the SFW Configuration Wizard for Microsoft Failover Cluster to configure a CVM cluster over the underlying Microsoft Failover Clustering networks.

Note: If User Access Control (UAC) is enabled, run the program or commands in the “Run as administrator” mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

1 On a cluster node, launch Solutions Configuration Center (SCC). (From the Start menu list all the applications installed and double-click Solutions Configuration Center)

From the Solutions view in the SCC, click **Launch** next to **SFW Configuration Wizard for Microsoft Failover Cluster**.

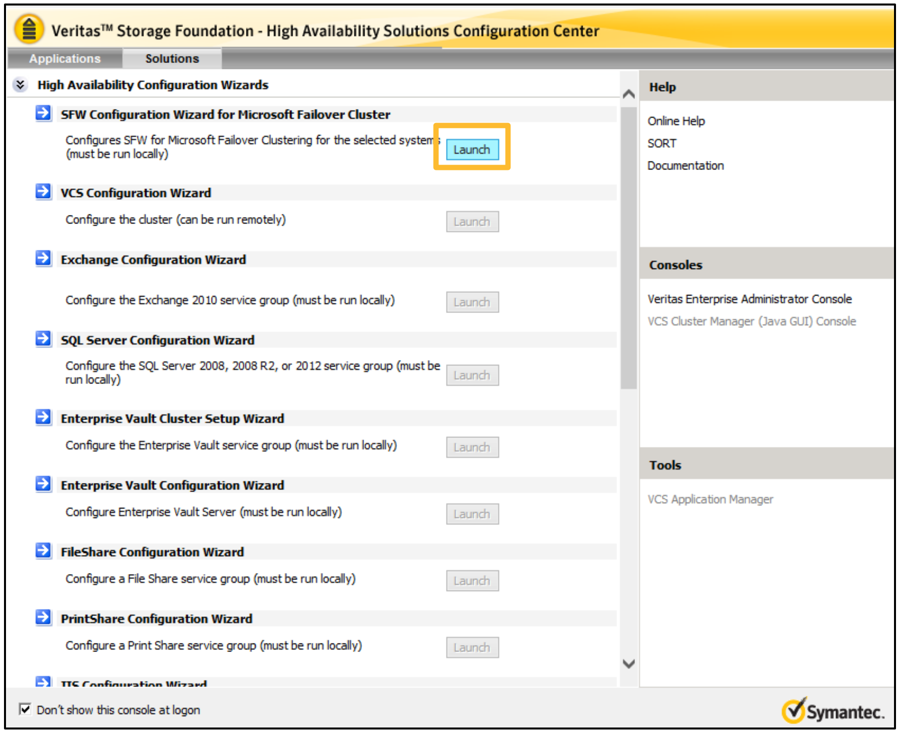

Veritas™ Storage Foundation - High Availability Solutions Configuration Center

Applications
Solutions
Help
Online Help
SORT
Documentation
Consoles
Veritas Enterprise Administrator Console
VCS Cluster Manager (Java GUI) Console
Tools
VCS Application Manager
Launch
☑ Don't show this console at logon
Symantec.

High Availability Configuration Wizards

SFW Configuration Wizard for Microsoft Failover Cluster
Configures SFW for Microsoft Failover Clustering for the selected systems (must be run locally)

VCS Configuration Wizard
Configure the cluster (can be run remotely)

Exchange Configuration Wizard
Configure the Exchange 2010 service group (must be run locally)

SQL Server Configuration Wizard
Configure the SQL Server 2008, 2008 R2, or 2012 service group (must be run locally)

Enterprise Vault Cluster Setup Wizard
Configure the Enterprise Vault service group (must be run locally)

Enterprise Vault Configuration Wizard
Configure Enterprise Vault Server (must be run locally)

FileShare Configuration Wizard
Configure a File Share service group (must be run locally)

PrintShare Configuration Wizard
Configure a Print Share service group (must be run locally)

IIS Configuration Wizard

2 On the SFW Configuration Wizard for Microsoft Failover Cluster panel, select the nodes to configure CVM support and then click **Configure**.

Note: While configuring CVM for the first time, even if one node is selected, the wizard configures all the nodes, including those that are not selected.

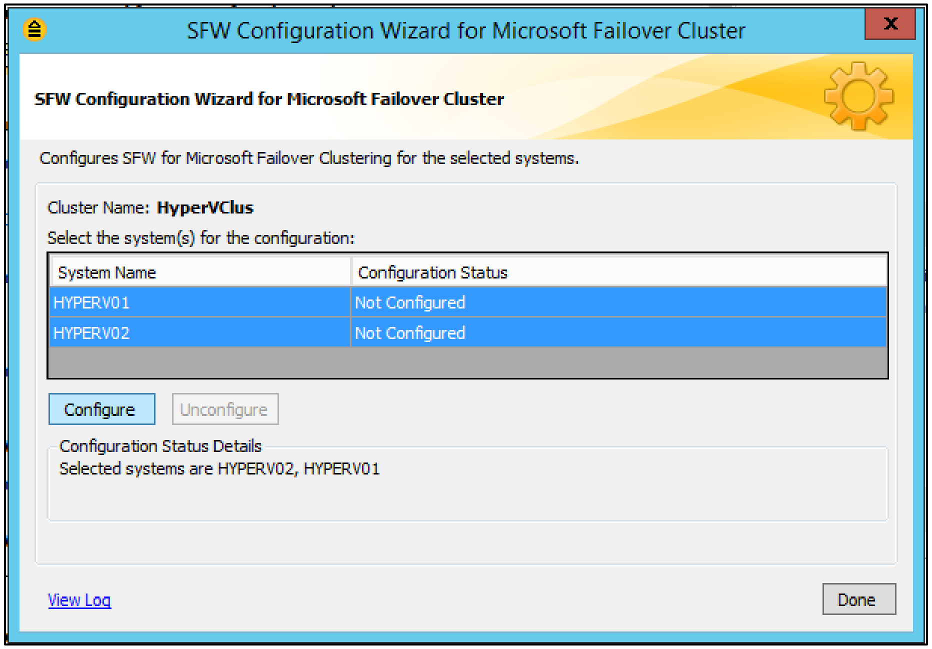

SFW Configuration Wizard for Microsoft Failover Cluster

SFW Configuration Wizard for Microsoft Failover Cluster

Configures SFW for Microsoft Failover Clustering for the selected systems.

Cluster Name: **HyperVClus**

Select the system(s) for the configuration:

System Name	Configuration Status
HYPERV01	Not Configured
HYPERV02	Not Configured

Configuration Status Details
Selected systems are HYPERV02, HYPERV01

The wizard configures the CVM cluster to communicate over the underlying Microsoft Failover Cluster. The configuration status for each selected node appears as “Configured”. Click **Done**.

Creating a CSDG using VEA

The following workflow represents the main tasks for creating a CSDG using VEA.

Note: Only the wizard panels that have considerable user actions are shown here. You must follow the wizard till the Finish page to complete the workflow.

1

On a cluster node, launch VEA.
In the tree-view, right-click **Disk Groups** and then select **New Dynamic Disk Group**.

2

On the New Dynamic Disk Group panel, perform the following steps:

- Specify a name for the disk group in the **Group name** text box.
- Select **Create cluster shared disk group**.
- Select the required disks from the Available disks box and click **Select**.
- Click **Next**.

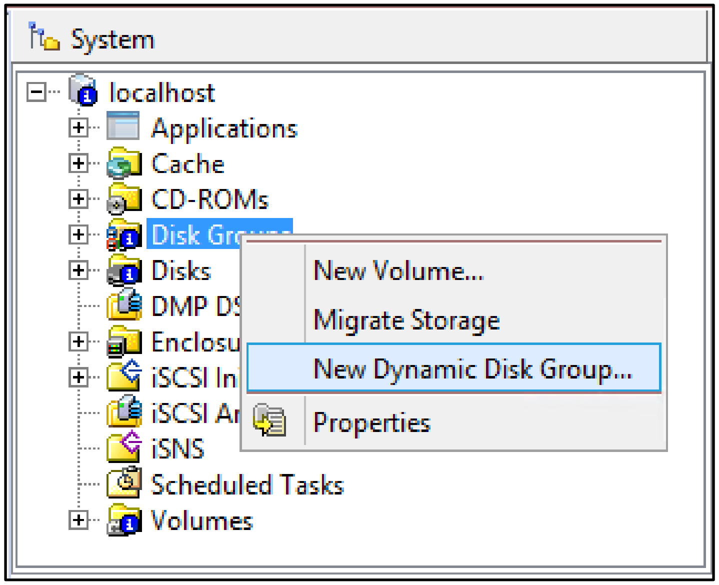

The New Dynamic Disk Group Wizard appears.

On the Welcome panel, click **Next**.

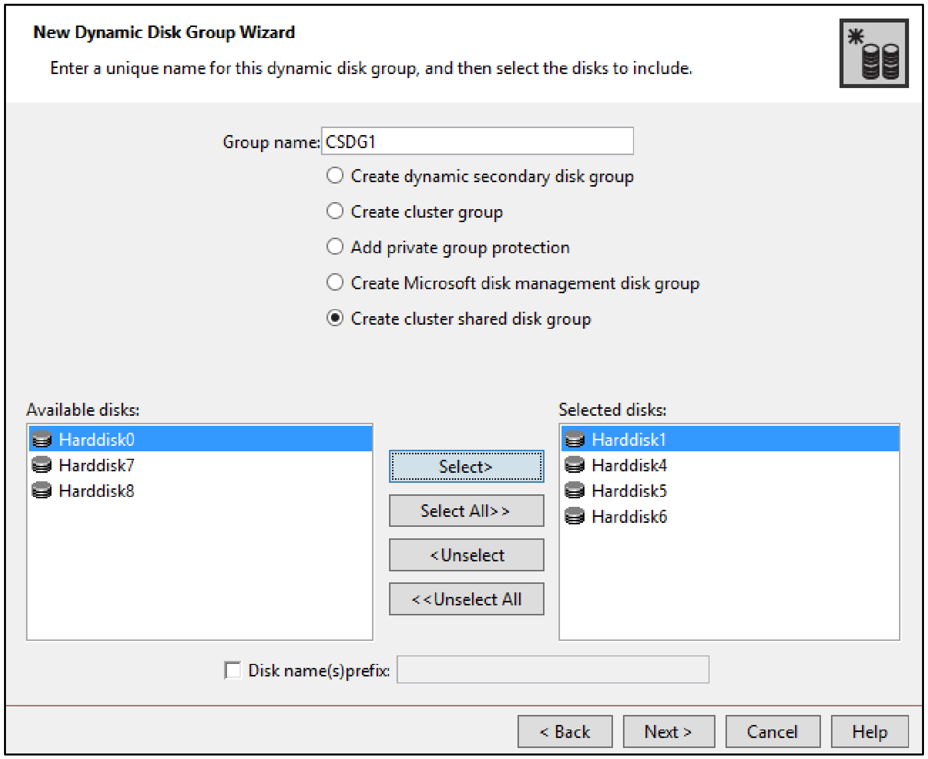

New Dynamic Disk Group Wizard

Enter a unique name for this dynamic disk group, and then select the disks to include.

Group name: CSDG1

○ Create dynamic secondary disk group

○ Create cluster group

○ Add private group protection

○ Create Microsoft disk management disk group

● Create cluster shared disk group

Available disks:

- Harddisk0
- Harddisk7
- Harddisk8

Selected disks:

- Harddisk1
- Harddisk4
- Harddisk5
- Harddisk6

□ Disk name(s)prefix:

Follow the wizard till the Finish page to complete the workflow.

Creating volumes on a CSDG using VEA

The following workflow represents the main tasks for creating volumes on a CSDG.

Note: Only the wizard panels that have considerable user actions are shown here. You must follow the wizard till the Finish page to complete the workflow.

1

On a cluster node, launch VEA.
In the tree-view, right-click the disk group you have created for the quorum and then select **New Volume**.

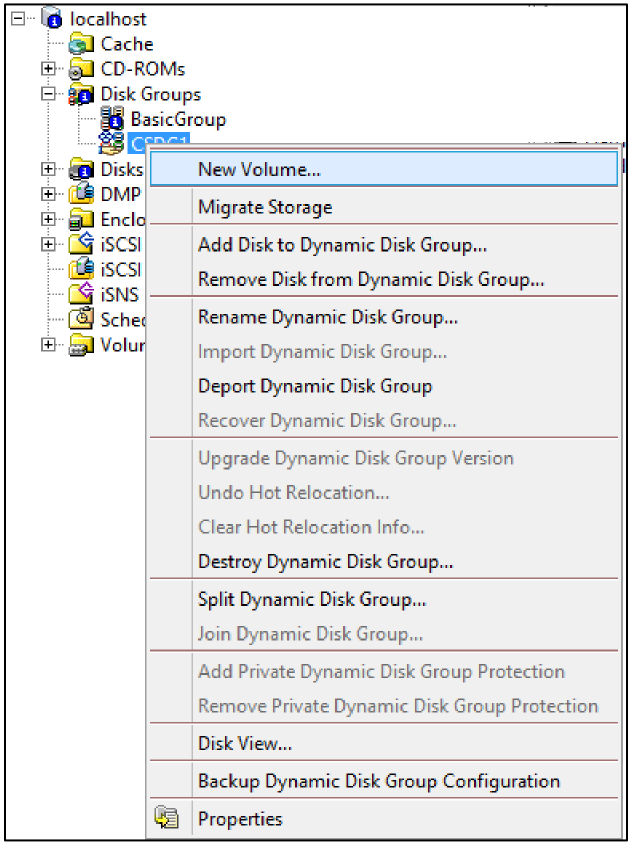
localhost
Cache
CD-ROMs
Disk Groups
BasicGroup
CSDG1
Disks
DMP
Enclo
iSCSI
iSCSI
iSNS
Sched
Volur
New Volume...
Migrate Storage
Add Disk to Dynamic Disk Group...
Remove Disk from Dynamic Disk Group...
Rename Dynamic Disk Group...
Import Dynamic Disk Group...
Deport Dynamic Disk Group
Recover Dynamic Disk Group...
Upgrade Dynamic Disk Group Version
Undo Hot Relocation...
Clear Hot Relocation Info...
Destroy Dynamic Disk Group...
Split Dynamic Disk Group...
Join Dynamic Disk Group...
Add Private Dynamic Disk Group Protection
Remove Private Dynamic Disk Group Protection
Disk View...
Backup Dynamic Disk Group Configuration
Properties

The New Volume Wizard appears.

On the Welcome panel, click **Next**.

2

On the Assign Disks for Volume panel, perform the following steps:

- Select the disk group from the **Group name** drop-down list.
- Select the appropriate site preference option.
- Select the required disks from the Available disks box and click **Select**.
- Click **Next**.

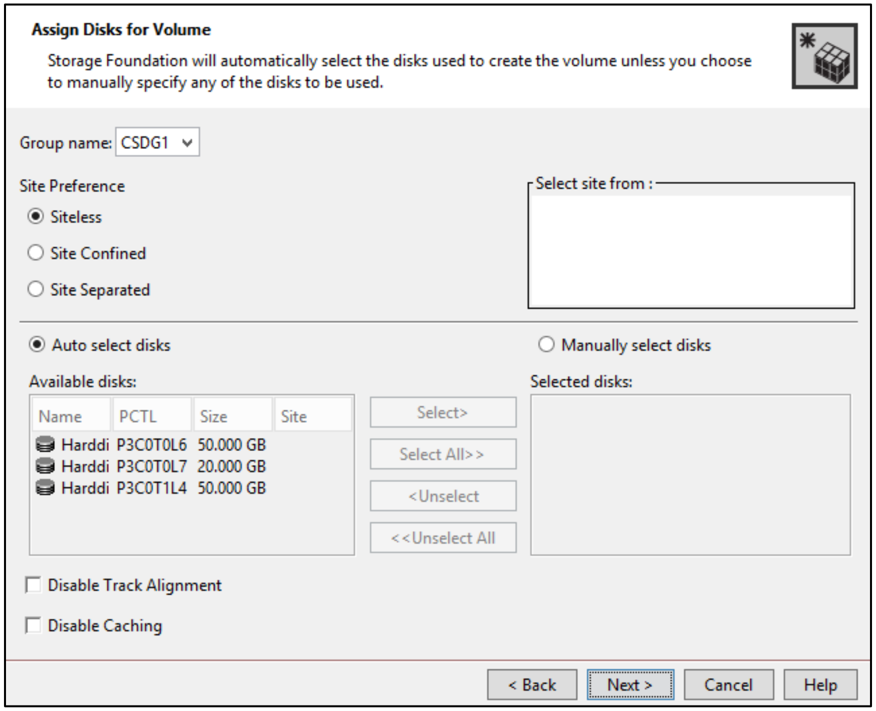

Assign Disks for Volume

Storage Foundation will automatically select the disks used to create the volume unless you choose to manually specify any of the disks to be used.

Group name: CSDG1

Site Preference

● Siteless

○ Site Confined

○ Site Separated

Select site from :

● Auto select disks ○ Manually select disks

Available disks:

Name	PCTL	Size	Site
Harddi	P3C0T0L6	50.000 GB	
Harddi	P3C0T0L7	20.000 GB	
Harddi	P3C0T1L4	50.000 GB	

Selected disks:

□ Disable Track Alignment

□ Disable Caching

(Contd…)

Creating volumes on a CSDG using VEA

-contd-

3

On the New Volume panel, perform the following steps:

1. Specify a name for the volume in the **Volume name** text box.
2. Specify the size for the volume to be created.
3. Select the layout for the volume to be created.
4. From the Mirror Info section, select **Mirrored** and then specify the number of disks on which you want to mirror the volume data.
5. Click **Next**.

4

On the Add Drive Letter and Path panel, select the drive letter and then click **Next**.

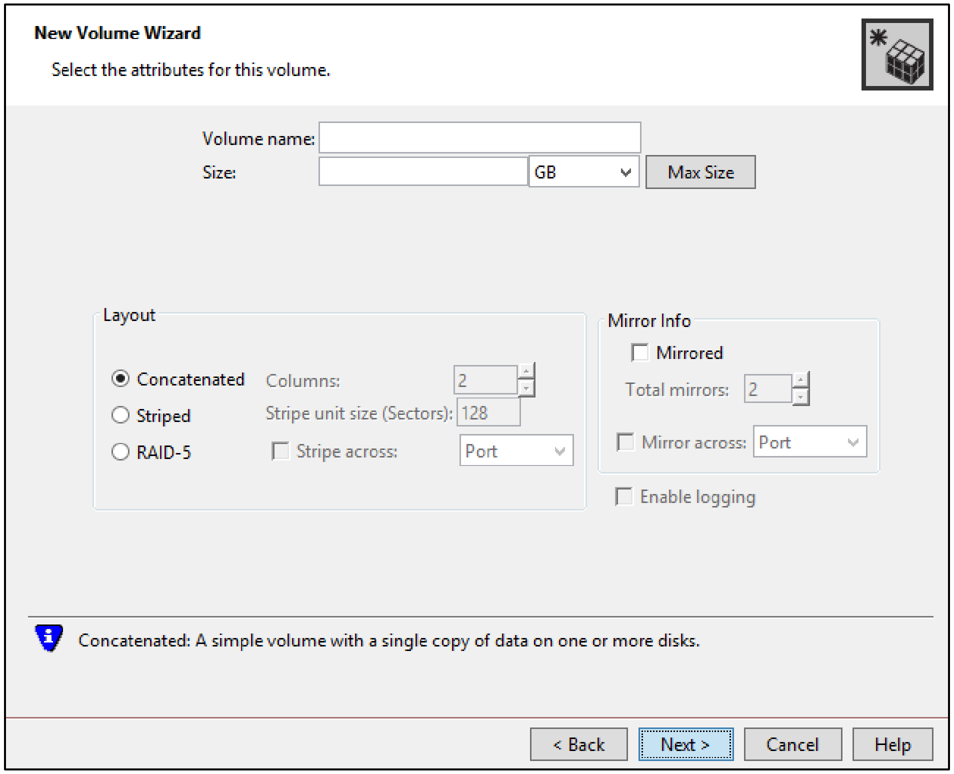

New Volume Wizard

Select the attributes for this volume.

Volume name:

Size: GB

Layout

● Concatenated Columns: 2

○ Striped Stripe unit size (Sectors): 128

○ RAID-5 □ Stripe across: Port

Mirror Info

□ Mirrored

Total mirrors: 2

□ Mirror across: Port

□ Enable logging

Concatenated: A simple volume with a single copy of data on one or more disks.

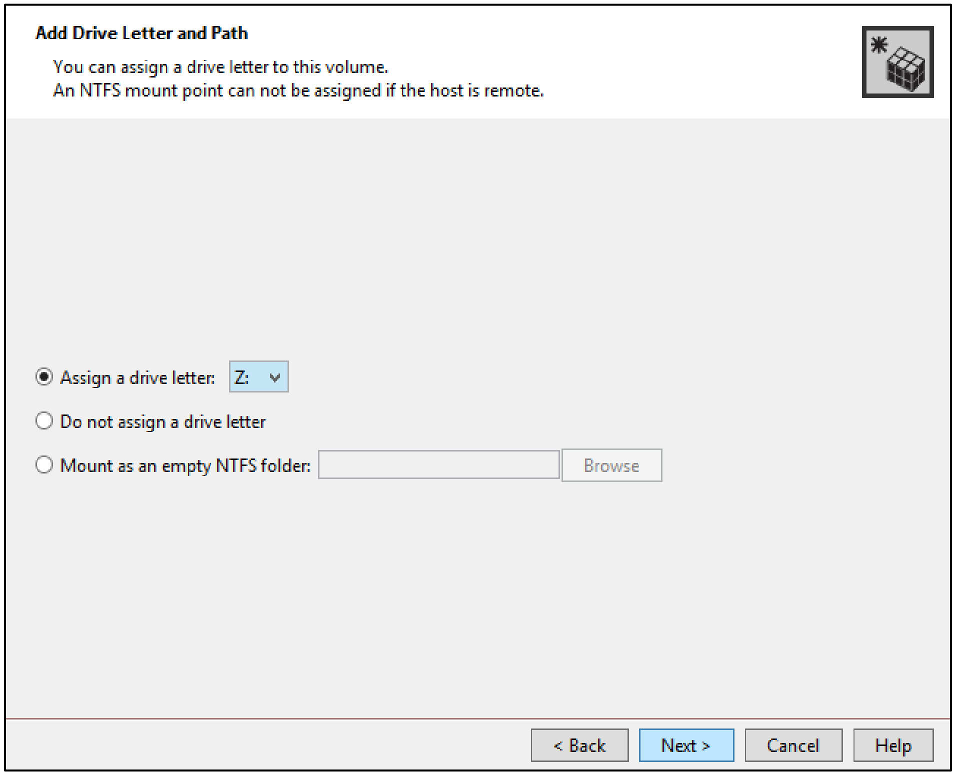

Add Drive Letter and Path

You can assign a drive letter to this volume.
An NTFS mount point can not be assigned if the host is remote.

● Assign a drive letter: Z:

○ Do not assign a drive letter

○ Mount as an empty NTFS folder:

Follow the wizard till the Finish page to complete the workflow.

Working with CVM using the ffconfigcli utility

Command	Description
`ffconfigcli -autoconfigure`	Configures a CVM cluster with a name same as that of the Microsoft Failover Cluster
`ffconfigcli -addnode <NodeName>`	Adds a node to a configured CVM cluster
`vxclustadm startnode`	Starts CVM on the node where you run this command
`vxclustadm [-f] stopnode`	Stops CVM on the node where you run this command
`vxclustadm [-v] nodestate`	Gives information about the state of a particular node in the CVM cluster
`vxclustadm nidmap`	Gives information about the node mapping and role (Master or Slave) of the nodes in the CVM cluster

Note

- Use the optional `-verbose` option at the end of a command to display information about the tasks that the command performs.
- Wherever the name of the node is required, you cannot optionally provide the node's IP address as it is not supported.
- Using the `-addnode` command, you can add only one node at a time. To add multiple nodes follow any of the following method:
    - Add the nodes one-by-one to both Microsoft failover cluster and CVM.
    - Add all the nodes to Microsoft failover cluster first and then add them one-by-one to CVM.

      If you use this method, CVM will auto-start only on the node added last. On the other previously-added nodes, you must manually start CVM.

Formatting and converting a cluster-shared volume to NTFS

Perform this task only if you have created a cluster-shared volumes using CLI.

Creating a volume using CLI creates a RAW and unformatted volume. To add a cluster-shared volume to a Microsoft Failover Cluster, you must first format it and then convert it to NTFS.

The following workflow represents the main tasks for formatting and converting a cluster-shared volume (here, VolDec10 with drive letter X) to NTFS:

1 On a cluster node, launch VEA. In the tree-view, right-click the volume and from the context menu select **Online Volume**.

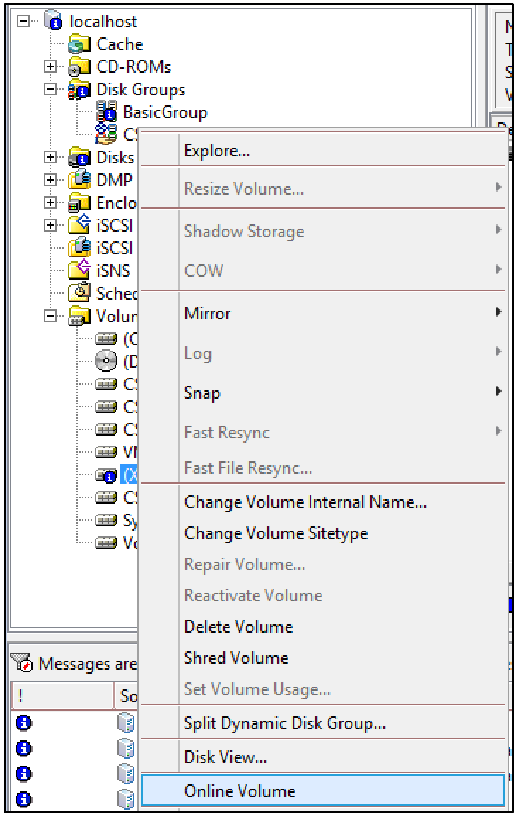

2 Right-click the volume and from the context menu select **File System > New File System**.

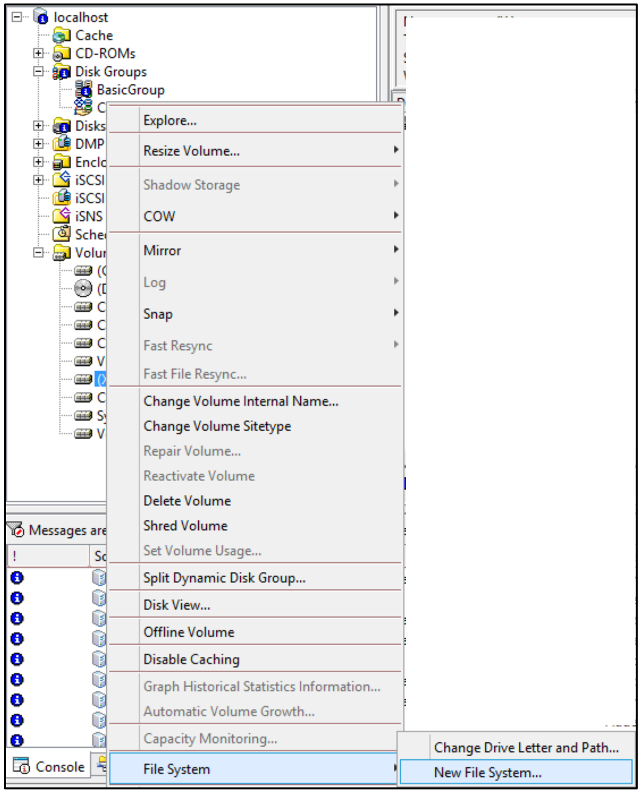

3 On the Create File System panel, select **NTFS** and then click **OK**.

Note: The **Perform a quick format** check box is selected by default. Do not clear the selection.

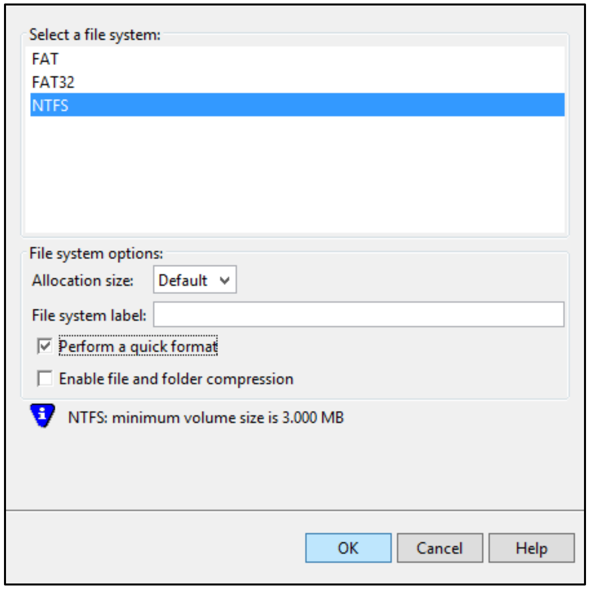

Select a file system:

FAT
FAT32
NTFS

File system options:

Allocation size: Default

File system label:

☑ Perform a quick format

☐ Enable file and folder compression

NTFS: minimum volume size is 3.000 MB

Creating a shared volume resource in Microsoft Failover Cluster

Perform this task only in case of the following scenarios:

- You upgraded a disk group to the latest version and import it as a cluster-shared disk group (CSDG)
- You created a cluster-shared volume using the VEA GUI and choose not to format the volume (by clearing the Format this volume check box)
- You created a cluster-shared volume using the SFW CLI

Note: Before you add the cluster-shared volume to a Microsoft Failover Cluster, you must format it and convert it to NTFS.

1 On a cluster node, launch Failover Cluster Manager.
In the tree-view, right-click **Roles** and then select **Create Empty Role**.

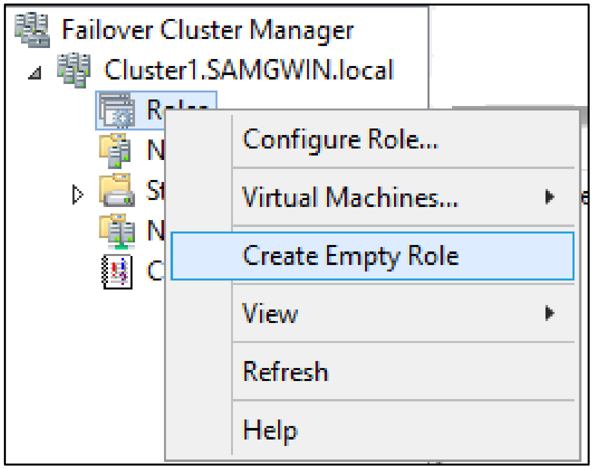

The wizard creates an empty role.

In the upper right pane, right-click the empty role and click **Properties** to specify a name and select a desired cluster node to own the role.

Click **OK**.

2 In the upper right pane, right-click the role and click **Add Resource** > **More Resources** > **Volume Manager Shared Volume**.

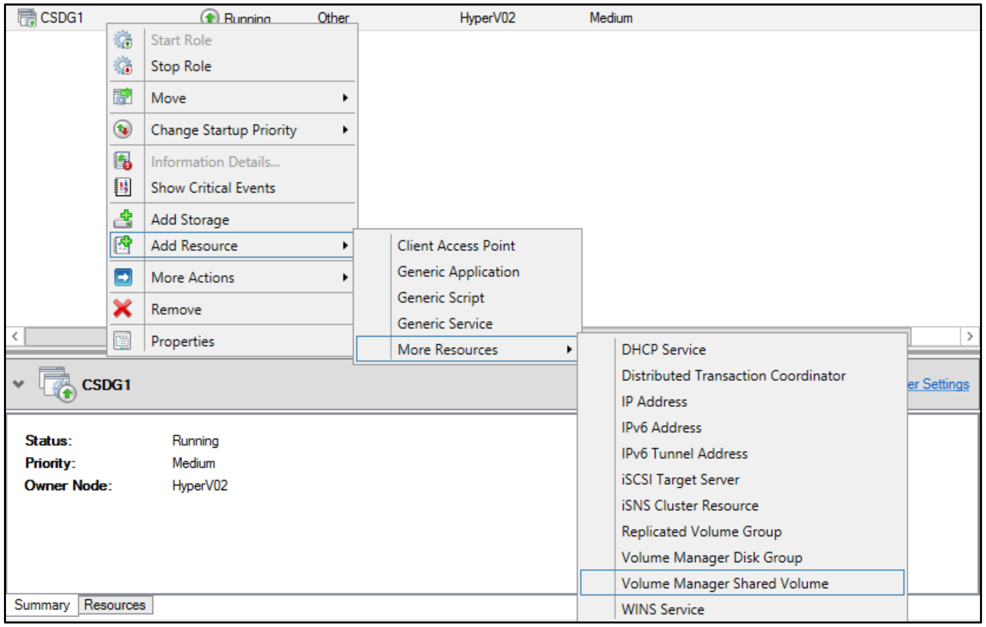

Status: Running
Priority: Medium
Owner Node: HyperV02

The wizard adds a Volume Manager Shared Volume resource to the role.

From the lower right pane, select the **Resources** tab and then right-click the resource added. From the context menu click **Properties**.

(Contd…)

Creating a shared volume resource in Microsoft Failover Cluster

-contd-

3

On the New Volume Manager Shared Volume Properties panel, perform the following tasks:

1. On the **General** tab, specify the name for the resource.
2. On the **Properties** tab, for MountPoint attribute, specify the drive letter that you assigned to the volume created.
3. Click **OK**.

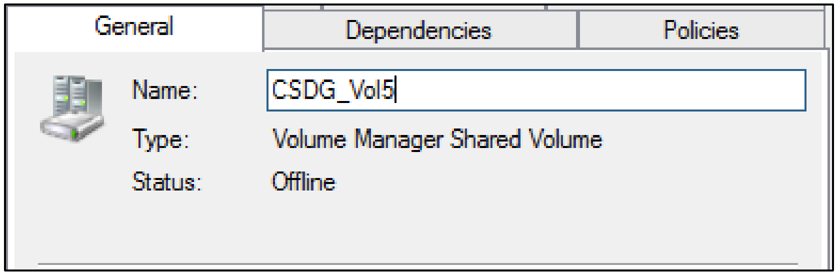

Name: CSDG_Vol5

Type: Volume Manager Shared Volume

Status: Offline

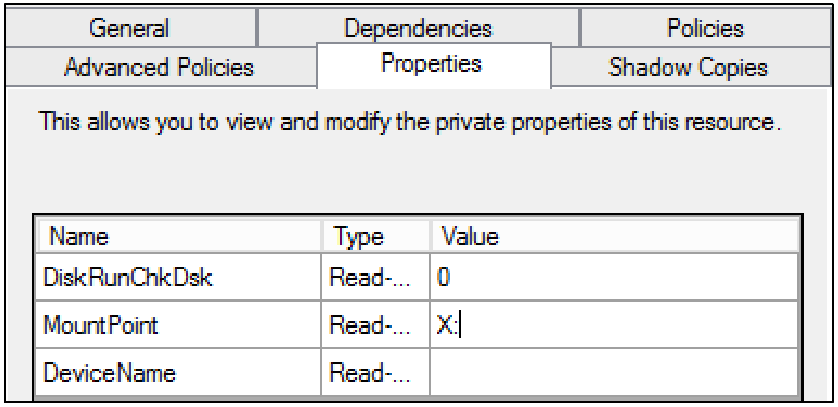

This allows you to view and modify the private properties of this resource.

Name	Type	Value
DiskRunChkDsk	Read-...	0
MountPoint	Read-...	X:
DeviceName	Read-...	

4

In the lower right pane, select the Resources tab and right-click the resource added. From the context menu click **Bring Online**.

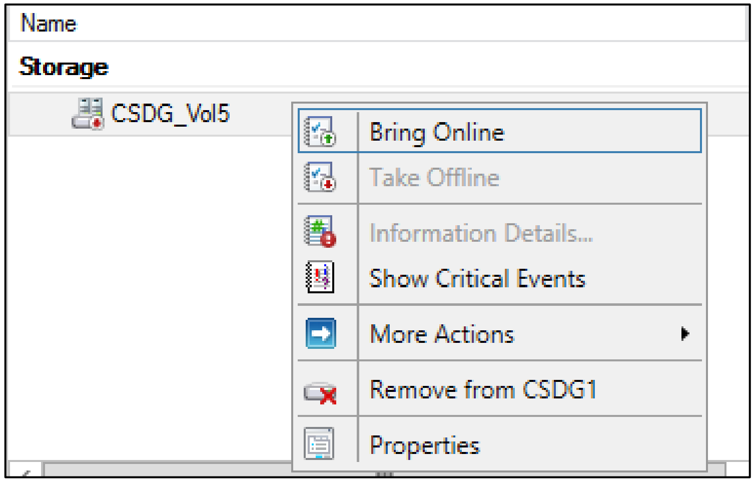

Configuring a role for volume-level failover

Perform this task to configure a role for high availability and to set dependencies between the resources. This task prepares the role for a volume-level failover.

This task involves the following steps:

1. Configure a role for each Hyper-V virtual machine or an application that you want to configure for high availability. A resource is created for the configured role.
2. Add a cluster-shared volume resource to the configured role
3. Set dependency between the two resources.

The following workflow represents the main tasks for configuring a role in Microsoft Failover Cluster, adding a cluster-shared volume as its resource, and setting a dependency between the cluster-shared volume resource and the resource created for the configured role:

1

On a cluster node, launch Failover Cluster Manager.
In the tree-view, right-click **Roles** and then select **Configure Role**.

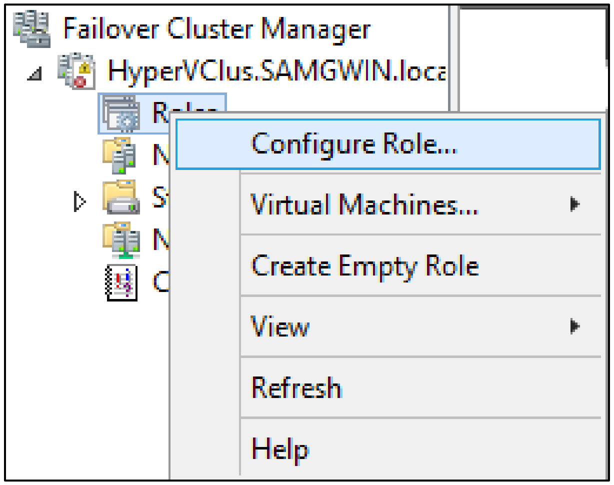

The High Availability Wizard appears.

On the Before You Begin panel, click **Next**.

2

On the Select Role panel, select the role you want to create and click **Next**.

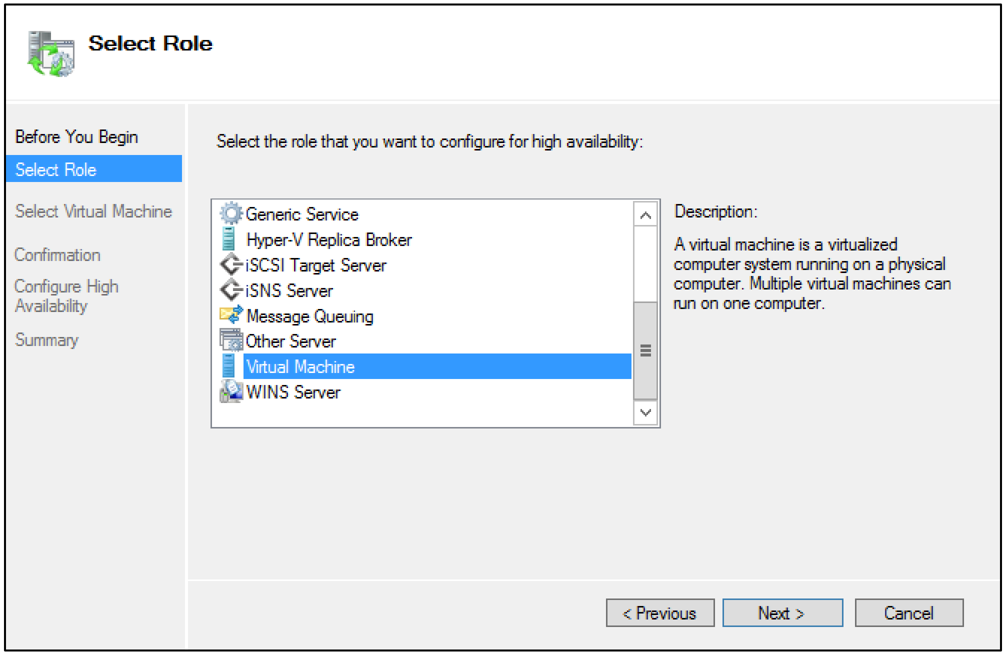

Select the role that you want to configure for high availability:

Description:

A virtual machine is a virtualized computer system running on a physical computer. Multiple virtual machines can run on one computer.

Follow the wizard till the Finish page to complete the workflow. The wizard configures the selected role and adds a corresponding resource.

(Contd…)

Configuring a role for volume-level failover

-contd-

3 In the upper right pane, right-click the configured role and select **Add Storage**.

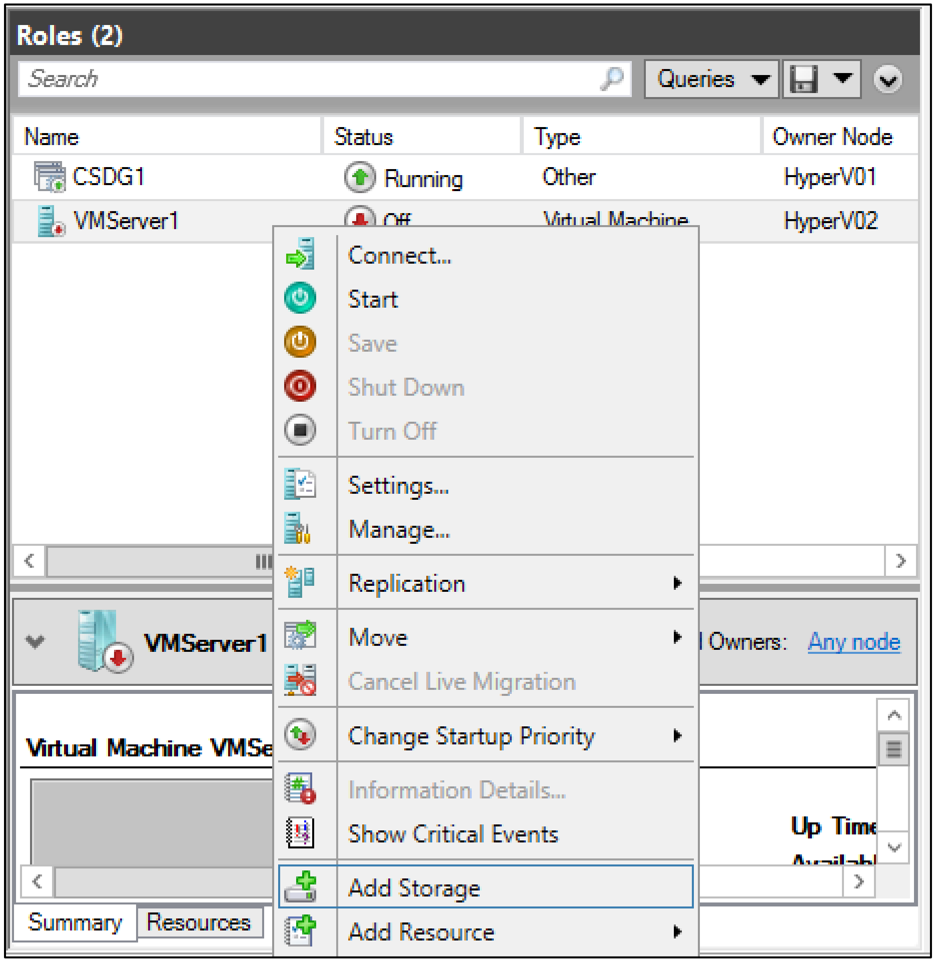

Roles (2)

Name	Status	Type	Owner Node
CSDG1	Running	Other	HyperV01
VMServer1	Off	Virtual Machine	HyperV02

4 On the Add Storage panel, select the required cluster-shared volume and click **OK**.

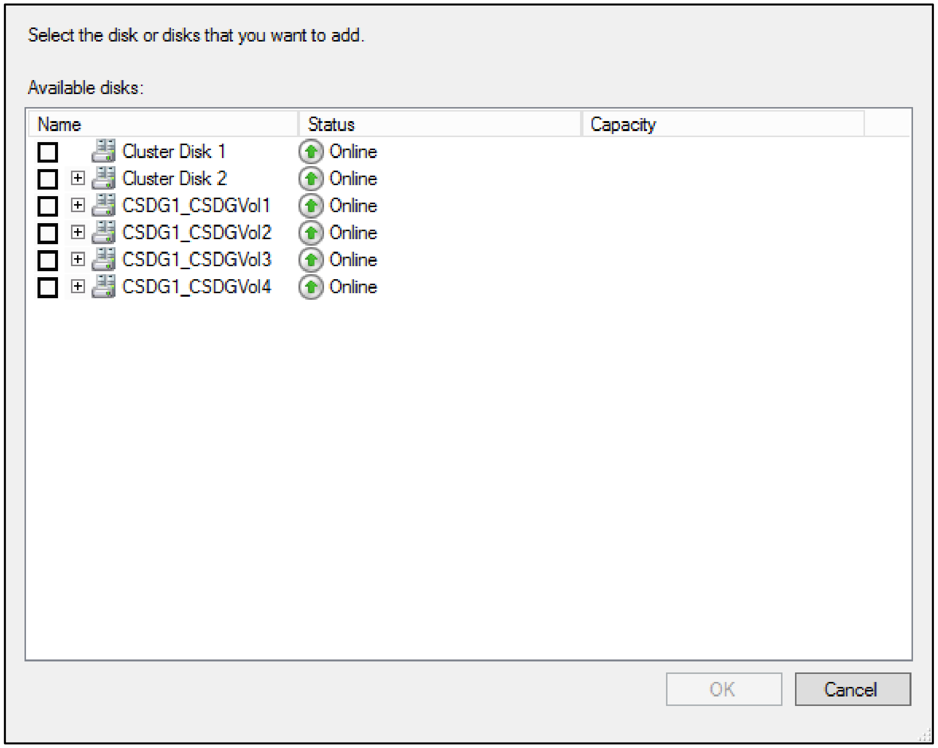

Select the disk or disks that you want to add.

Available disks:

Name	Status	Capacity
□ Cluster Disk 1	Online	
□ ⊞ Cluster Disk 2	Online	
□ ⊞ CSDG1_CSDGVol1	Online	
□ ⊞ CSDG1_CSDGVol2	Online	
□ ⊞ CSDG1_CSDGVol3	Online	
□ ⊞ CSDG1_CSDGVol4	Online	

The wizard adds the selected cluster-shared volume as a storage resource for the configured role.

(Contd…)

Configuring a role for volume-level failover

-contd-

5 In the lower right pane, right-click the resource for the configured role and select **Properties**.

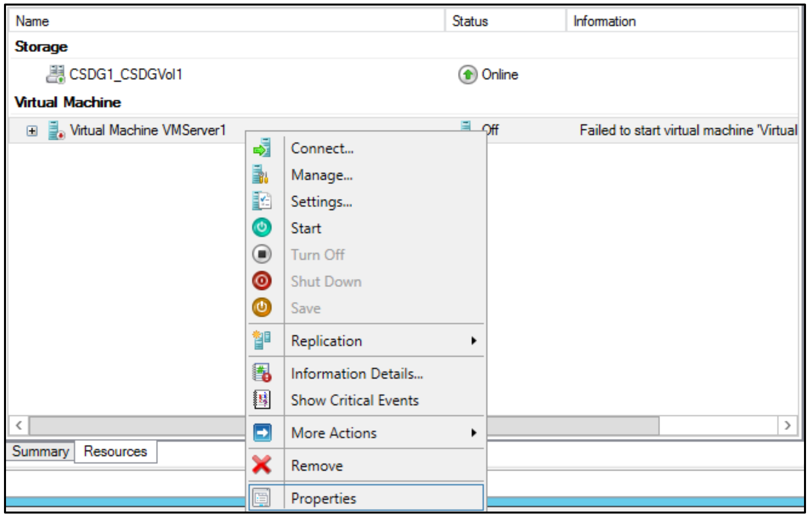

Name	Status	Information
Storage		
CSDG1_CSDGVol1	Online	
Virtual Machine		
Virtual Machine VMServer1	Off	Failed to start virtual machine 'Virtual

6 On the resource properties panel, select the **Dependencies** tab and from the Resource drop-down list select the cluster-shared volume resource. Click **OK**.

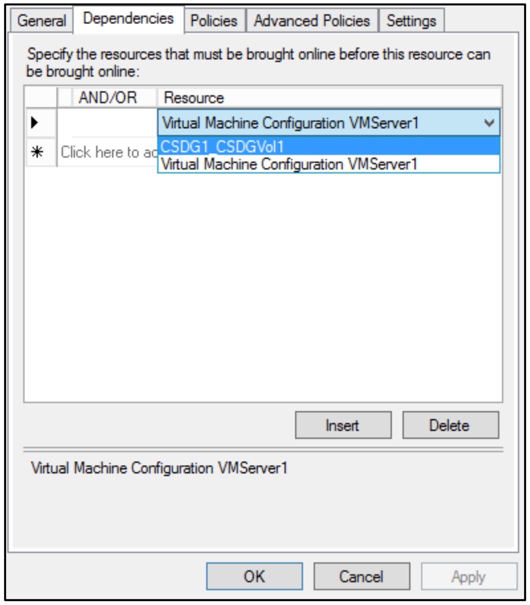

Specify the resources that must be brought online before this resource can be brought online:

AND/OR	Resource
	Virtual Machine Configuration VMServer1
Click here to ad	CSDG1_CSDGVol1 Virtual Machine Configuration VMServer1

Virtual Machine Configuration VMServer1

7 In the upper right pane, right-click the role and select Start.

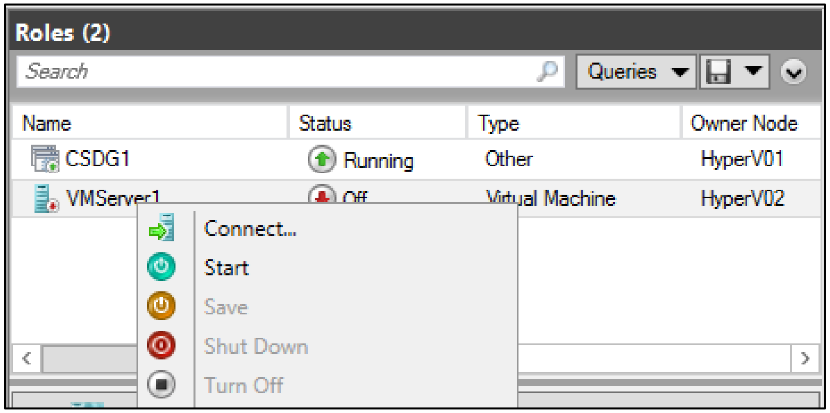

Roles (2)

Name	Status	Type	Owner Node
CSDG1	Running	Other	HyperV01
VMServer1	Off	Virtual Machine	HyperV02

This completes the tasks required for configuring a role for volume-level failover.

References

Refer to the product guides, marketing recordings and compatibility lists for a deeper understanding of Cluster Volume Manager.

- Storage Foundation and High Availability Solution Installation and Upgrade Guide
- Symantec Storage Foundation Administrator's Guide
- Hardware Compatibility List (HCL)
- Software Compatibility List (SCL)

The Storage Foundation for Windows documentation for other releases and platforms can be found on the SORT website.