

Dell EMC SC Series Storage with SAS Front-end Support for Microsoft Hyper-V

Abstract

This document describes how to configure Microsoft® Hyper-V® hosts equipped with supported SAS HBAs to access SAN storage on select Dell EMC™ SC Series arrays with SAS front-end ports.

October 2017

Revisions

Date	Revision
October 2015	Initial release with support for Dell SCv2000
February 2016	Updated technical support contact information
July 2016	Updated to include support for Dell SC4020
October 2017	Updated to include support for Dell EMC SCv3000 and Dell EMC SC5020

Acknowledgments

Author: Marty Glaser

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2015-2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgments.....	2
Executive summary.....	4
Audience	4
Feedback	4
1 Introduction.....	5
1.1 Supported hardware	5
1.2 Microsoft Windows Server clustering and Hyper-V	6
1.3 Transport options.....	6
2 SAS FE host path configuration options	8
2.1 Multipath configuration	8
3 Configure Hyper-V hosts to access SC Series arrays with SAS FE ports	16
3.1 Prepare the environment.....	16
3.2 Install and configure a SAS HBA in each Windows host server	17
3.3 Configure SAS FE fault domains to support MPIO hosts.....	20
3.4 Connect Windows host servers to the SC Series array with SAS cables	23
3.5 Create server objects	24
3.6 Create cluster server object on the storage array	27
3.7 Create and map storage volumes to the server cluster.....	28
3.8 Configure MPIO on the host servers	31
4 Create a Hyper-V cluster.....	34
4.1 Validate the servers to be clustered	34
4.2 Create a new Hyper-V cluster	36
4.3 Convert the cluster disk to a cluster shared volume.....	37
5 Support for guest VMs with SAS pass-through disks	38
5.1 Windows Server 2008 R2 and 2012.....	38
5.2 Windows Server 2012 R2 and newer	39
A Additional resources.....	40

Executive summary

Select Dell EMC™ SC Series arrays support serial-attached SCSI (SAS) front-end (FE) ports for connecting host servers equipped with a supported SAS host bus adapter (HBA) directly to SC Series array SAN storage. SAS FE is a simple, cost-effective transport option that is ideal for locations such as a branch office with a limited number of host servers.

The focus of this paper is SAS FE support for Microsoft® environments with Microsoft Windows Server® Hyper-V®, and includes cabling diagrams and step-by-step configuration guidance. SAS FE support in VMware® environments is addressed in the document, [Dell EMC SC Series Storage with SAS Front-end Support for VMware vSphere](#).

Audience

This document is for administrators to who want to learn more about how to configure Microsoft Hyper-V hosts equipped with supported SAS HBAs to access SAN storage on select Dell SC Series and Dell EMC SC Series arrays with SAS FE ports. Readers should have working knowledge of SC Series storage and Microsoft Hyper-V environments.

Feedback

We value customer feedback as we strive to provide high quality documentation in support of Dell EMC products. Please send feedback or recommendations on how we can improve this document to StorageSolutionsFeedback@dell.com.

1 Introduction

All SC Series arrays can be configured to support iSCSI or Fibre Channel (FC) for connecting host servers to SAN storage. These are flexible, robust, and highly scalable transport options that are the best choice for most SAN environments. For a location such as a small branch office, SAS FE is an attractive, simple, cost-effective transport option because it does not require additional switch hardware or support expertise. However, with SAS FE, the number of physical host servers supported with each SC array is limited to four hosts. The figures in this document present cabling options.

1.1 Supported hardware

The following SC Series arrays support SAS FE:

- SCv2000
- SCv3000
- SC4020
- SC5020

The SCv2000 and SCv3000 arrays are affordable, entry-level systems that offer many of the same enterprise-class features as other SC Series arrays, and the SC4020 and SC5020 are fully featured arrays. Each of these arrays supports three different front-end transport options: FC, iSCSI, or SAS FE. Customers can decide at the time of purchase the type of front-end connectivity that is right for each SC Series array in their environment.




Figure 1 Front and rear views of the SCv2000/SC4020 with SAS FE ports




Figure 2 Front and rear views of the SCv3000/SC5020 with SAS FE ports

For more information about the SC Series arrays discussed in this paper — including release notes, getting started guides, system deployment guides, and owner's manuals — see the resources available at [Dell Support](#).

1.2 Microsoft Windows Server clustering and Hyper-V

Microsoft Windows Server clustering and Hyper-V provide the foundation for creating highly available (HA) host server and VM configurations in Microsoft environments. HA provides load balancing and redundancy to enhance performance and protect against single points of failure.

Hyper-V is the Microsoft virtualization platform for Windows Server. Introduced with Windows Server 2008, each new release has incorporated many new features and enhancements. Support for SAS FE with SC Series arrays requires Windows Server 2008 R2 Hyper-V or newer.

This document assumes the reader has working knowledge of Microsoft Windows Server, Windows clustering, and Microsoft Hyper-V. For more information, see the [Windows IT Center](#).

1.3 Transport options

For most environments, FC or iSCSI are the preferred transport methods for configuring front-end connectivity between host servers and SC Series storage. Fibre Channel and iSCSI are mature, robust, proven technologies that can scale to include a very large number of hosts and SC Series arrays across multiple locations. When configured with redundant fabrics, these transports offer highly resilient and reliable data transfer between SAN storage and hosts. However, they require additional hardware components including switches, HBAs, and cabling, along with the technical expertise to support the technology.

For small environments with a limited number of physical host servers, SAS FE connectivity provides comparable performance and resiliency to FC or iSCSI but without the extra cost and complexity of additional hardware components. There are important scale and design factors to consider if choosing SAS FE instead of FC or iSCSI:

Scale: With SAS FE, the number of physical hosts per SC Series array is limited to a maximum of four MPIO hosts.

Design: With SAS FE, physical hosts must be in close proximity to the SC Series array, typically in the same or an adjacent rack that is within reach of SAS cabling (typically one to six meters in length).

2 SAS FE host path configuration options

When an SC Series array is equipped with SAS FE ports, a total of eight ports (four per controller) are available to connect host servers. To support MPIO, each SC Series array supports a maximum of four host servers (two SAS ports per host).




Figure 3 SCv2000/SC4020 with SAS FE ports (4 ports per controller)



Figure 4 SCv3000/SC5020 with SAS FE ports (4 ports per controller)

Note: Dell EMC recommends configuring all hosts to use MPIO as a best practice regardless of the type of transport in order to provide path redundancy. Carefully consider the risks before configuring any host to use single-path. The customer assumes all risks associated with the design of their environment.

2.1 Multipath configuration

To provide path redundancy and ensure uptime, each host should be configured with at least two data paths to external storage. With MPIO, if a data path fails, another data path provides the host uninterrupted data access to SAN storage. This is particularly important for standalone Hyper-V or other Windows hosts that do not have node-level redundancy.

Microsoft offers native MPIO support with SC Series storage that is quick and easy to configure. For more information on MPIO for Microsoft Windows including Hyper-V, see the [Dell EMC SC Series Storage: Microsoft Multipath I/O Best Practices](#).

With SAS FE, up to four host servers can be connected to each SC Series array in any combination of standalone hosts or cluster nodes. Figures 5 through 10 show several different MPIO cabling options for Hyper-V hosts and clusters.

Each color in Figures 5 through 10 represents a separate SAS FE fault domain. Fault domains protect host servers against a single path or single controller failure.

Each SAS FE fault domain consists of two SAS FE ports. To group the SAS FE ports, group SAS FE port 1 on each controller to create SAS fault domain 1, group SAS FE port 2 on each controller to create SAS fault domain 2, and so on. A corresponding SAS FE port from each controller must be used when creating a SAS FE fault domain as shown in the figures.

Note: SAS FE supports multiple dual-port SAS HBAs per physical host server (up to four per host if the host server has available PCIe slots). If more than one SAS HBA is installed in a host server, the HBAs must be identical. Each dual-port HBA must be connected to only one fault domain as shown in Figure 11 and Figure 12.

Note: The SCv2000 and SC4020 arrays utilize a 2U chassis. The SCv3000 and SC5020 arrays utilize a 3U chassis. SAS FE functionality is similar on supported SC Series models.




Figure 5 SCv2000/SC4020 with a 4-node MPIO Hyper-V cluster




Figure 6 SCv3000/SC5020 with a 4-node MPIO Hyper-V cluster

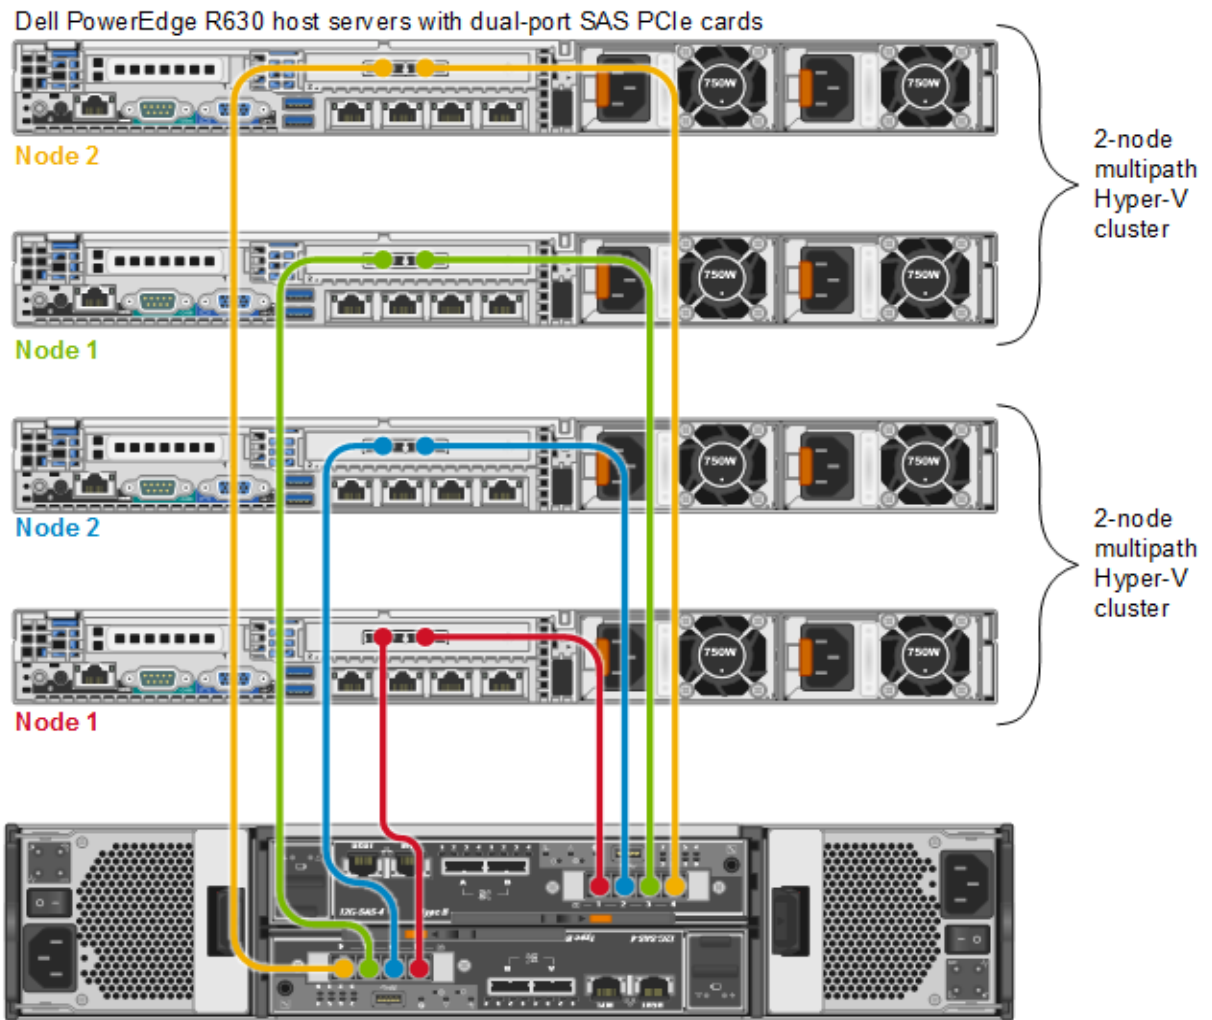


Figure 7 SCv2000/SC4020 with two 2-node MPIO Hyper-V clusters




Figure 8 SCv3000/SC5020 with two 2-node MPIO Hyper-V clusters

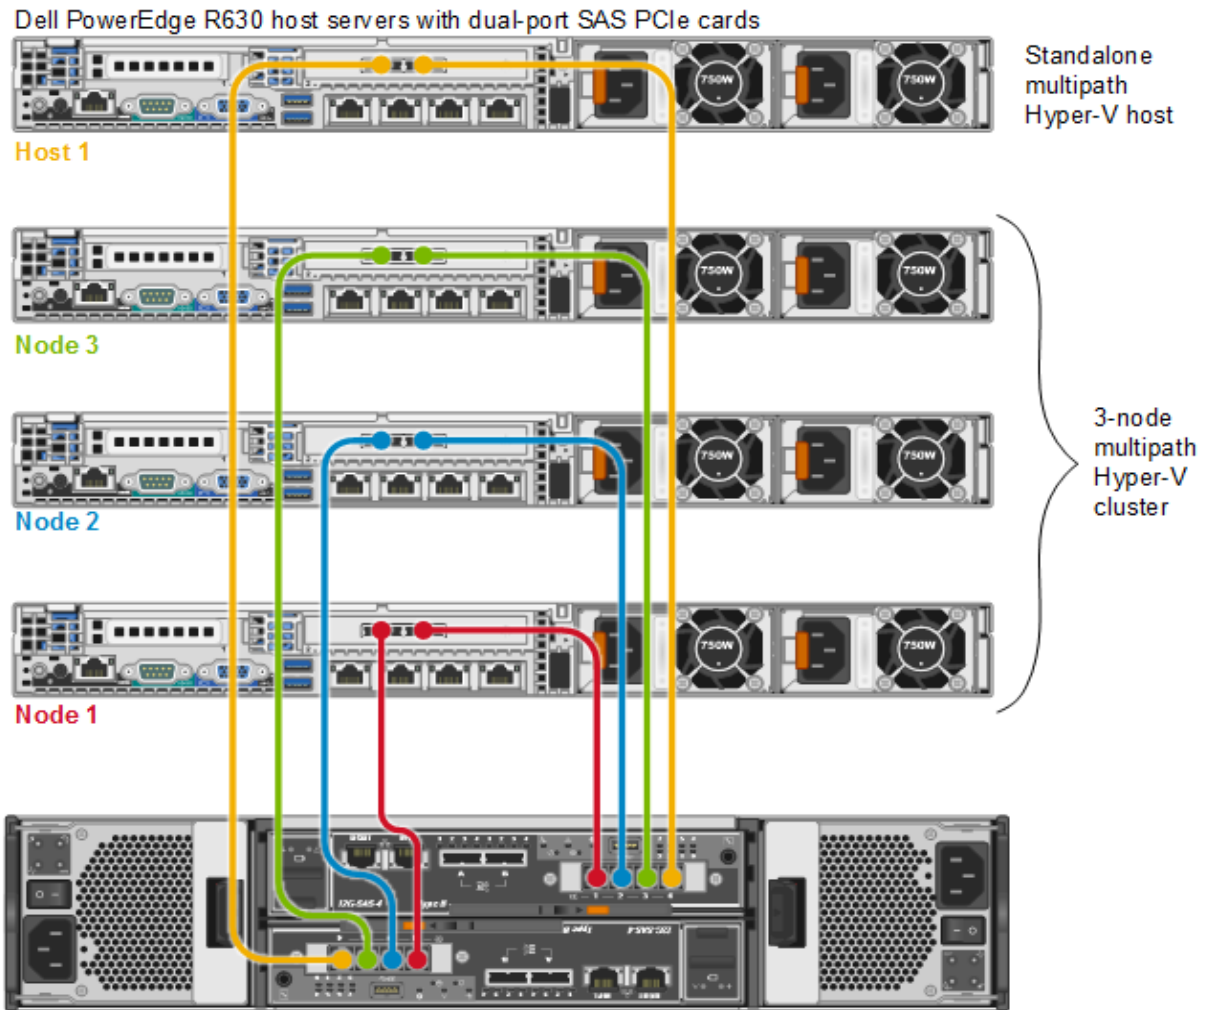


Figure 9 SCv2000/SC4020 with a 3-node MPIO Hyper-V cluster and MPIO standalone host




Figure 10 SCv3000/SC5020 with a 3-node MPIO Hyper-V cluster and MPIO standalone host

Dell PowerEdge R630 host servers with two dual-port SAS PCIe cards each




Figure 11 SCv2000/SC4020 with a 2-node MPIO Hyper-V cluster with two SAS PCIe cards per host

Dell PowerEdge R630 host servers with two dual-port SAS PCIe cards each




Figure 12 SCv3000/SC5020 with a 2-node MPIO Hyper-V cluster with two SAS PCIe cards per host

3 Configure Hyper-V hosts to access SC Series arrays with SAS FE ports

This section provides instructions for configuring Microsoft hosts with SAS HBAs to access SAN storage on an SC Series array equipped with SAS FE ports. Step-by-step guidance is also provided for configuring a Hyper-V cluster to use cluster shared volumes (CSVs) presented to Hyper-V nodes with SAS FE.

3.1 Prepare the environment

Before configuring Windows Server hosts with SAS HBAs to access storage on an SC Series array with SAS FE ports, verify compliance with the prerequisites listed in Table 1.

Table 1 Prerequisites checklist

✓	Tasks
	Verify that at least one SC Series array equipped with SAS FE ports is configured and available.
	<p>Verify that software versions are compatible with SAS-FE:</p> <ul style="list-style-type: none"> • SCv2000 arrays require Dell Enterprise Manager (EM) Client software version 2015 R1 or newer to support SAS FE. Use this software to discover, configure, and manage the SCv2000 array. Installing a Data Collector is supported but not required. Refer to the release notes and administrator's guide as needed. • The SC4020 array requires Dell Storage Center Operating System (SCOS) 7.1 or newer, and Dell Storage Manager (DSM) version 2016 R2 or newer. • The SCv3000 and SC5020 arrays require SCOS 7.2.10 or newer, and DSM 2016 R3 or newer. <p>Running the latest version of SCOS and DSM for the SC Series array is recommended to take advantage of the latest enhancements and bug fixes.</p>
	<p>Ensure that at least one supported SAS PCIe HBA interface card is present for each Windows host, along with SAS cables of the appropriate length. Each host server must have one or more compatible full or half-height PCIe slots available for one or more HBAs.</p> <p>Only SAS HBAs listed in the Dell EMC Storage Compatibility Matrix are supported. The example configurations shown in this guide use the Dell 12Gb SAS HBA. If more than one SAS HBA is installed in the same host, the HBAs must be identical.</p> <p>Note: Only Dell EMC PowerEdge™ servers (13G or newer) are supported if using the Dell 12Gb SAS HBA.</p> <p>For more information, refer to the Dell PowerEdge Controller 9 HBA User's Guide.</p>
	Ensure that servers have an on-board disk controller and disk configuration to support a local boot disk. Boot-from-SAN is not supported with SAS FE.
	Ensure that servers are located in close proximity to the SC Series array. SAS cables are typically one to six meters in length.

	<p>Stage the host servers with a supported Windows Server OS version and patch them to the desired level. Windows Server 2008 R2 or newer is required to support SAS FE HBA drivers. Windows Server 2016 (with Desktop) is used in the examples shown in this document.</p> <p>As a best practice, use the Dell PowerEdge Server Lifecycle Controller (press F10 at boot) to update the internal hardware components of your host servers and to broker the OS install using the latest PowerEdge Server OS driver pack.</p>
	<p>Join the Windows server Hyper-V hosts to a Microsoft Active Directory® domain (recommended), and install the following roles and features as a minimum:</p> <ul style="list-style-type: none">• Hyper-V• Failover Clustering• Multipath I/O

3.2 Install and configure a SAS HBA in each Windows host server

After the prerequisite steps are complete, install a supported SAS HBA in an available PCIe slot in each host server, and update the firmware and Windows driver by following the steps in this section. If installing more than one SAS card per host, the SAS HBAs must be identical. The example configuration in this document shows one HBA per server.




Figure 13 Install a SAS HBA in an available full- or half-height PCIe slot

1. While observing safe electrostatic discharge (ESD) precautions, power off the Windows host server and install a supported SAS HBA PCIe card into an available full or half-height PCIe slot. In this example, a Dell 12Gb SAS HBA is installed in a full-height PCIe slot in a Dell PowerEdge R630 (13G) server.
2. Power on the Windows host and press **F10** at boot to access the Dell Server Lifecycle Controller (LC). Use the LC to verify and update the firmware version (preferred method). Optionally, boot into Windows and follow the steps in the [Dell PowerEdge Controller 9 HBA User's Guide](#) to update the SAS HBA Windows driver and firmware.

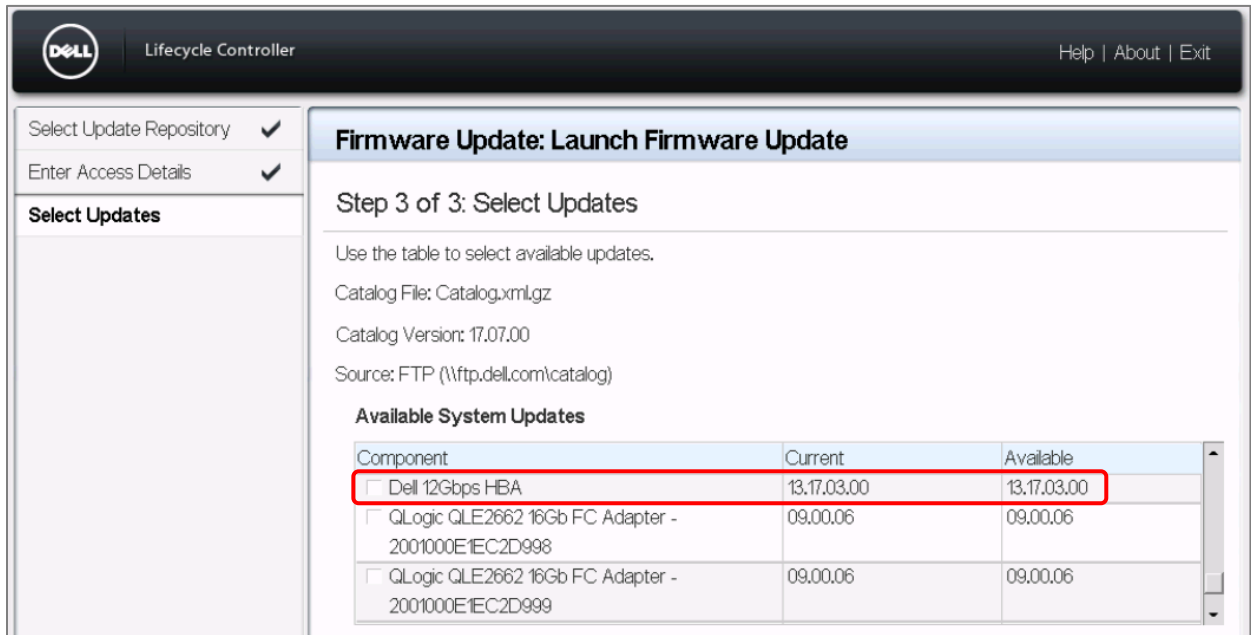


Figure 14 Dell server Lifecycle Controller

3. After updating the firmware, reboot the server and log in to Windows. Launch the Device Manager, and under **Storage controllers**, open the SAS HBA properties and note the driver version.

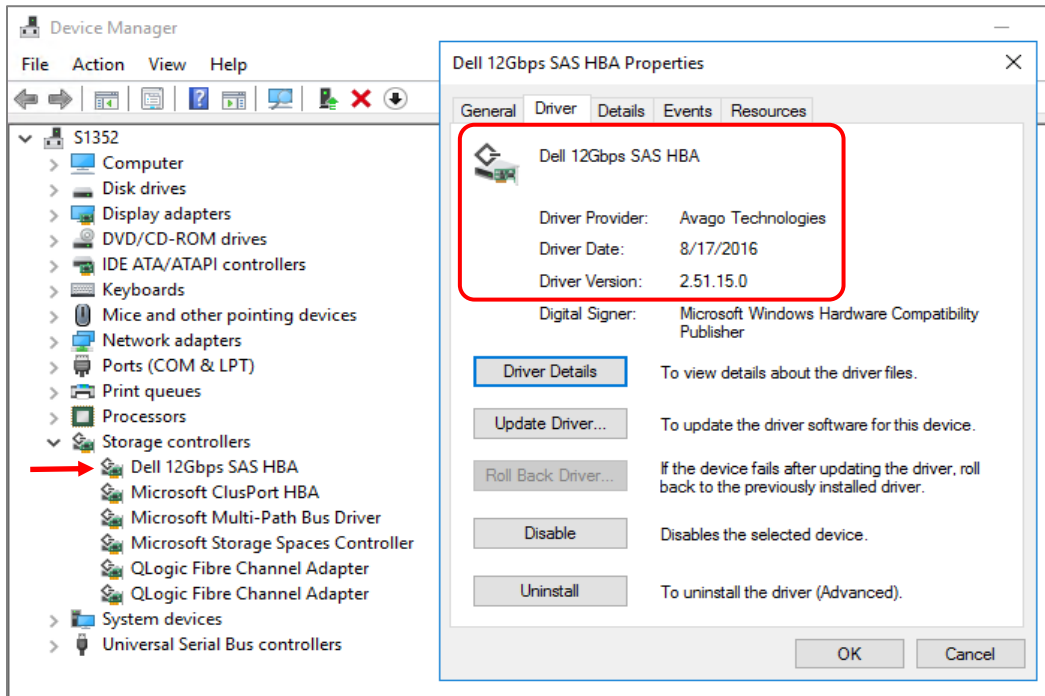


Figure 15 SAS HBA listed in Device Manager

4. Compare the driver version with the version that is available online for your server hardware and OS. If there is a newer driver available, download and install it. In this example, the latest driver and firmware for the **Dell 12Gbps SAS HBA** is located under the **SAS Non-RAID** section for the Dell PowerEdge R630 server.

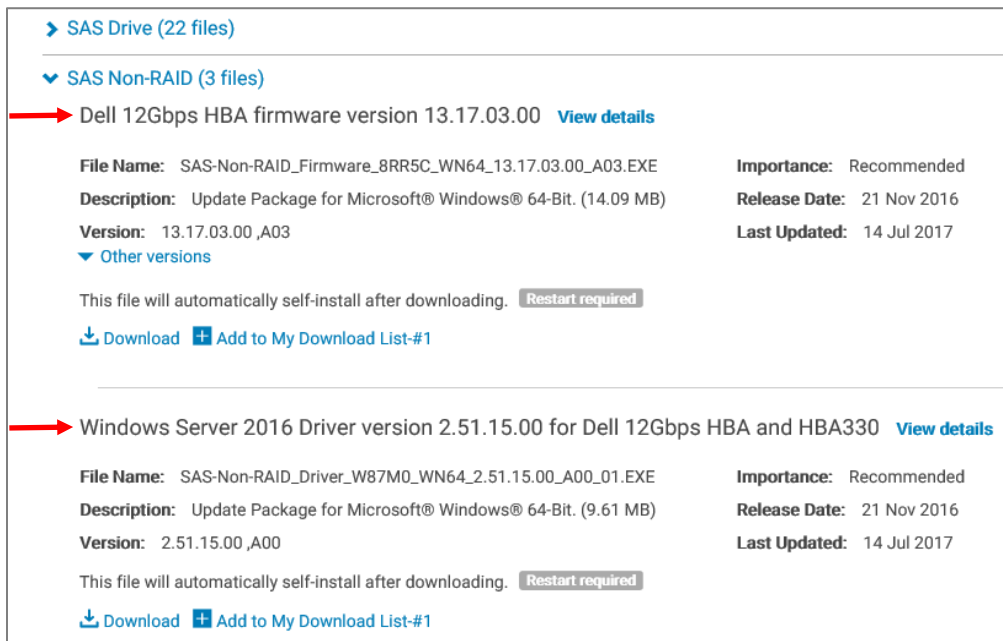


Figure 16 Driver and firmware downloads


Note: The default driver installed by Windows may not be the correct driver or may be outdated. It is important to install the latest driver available from Dell EMC. Verify that the HBA firmware is current when the Windows driver is updated.

- Repeat steps 1–4 to install SAS HBAs in additional Windows host servers. Verify that the SAS HBA firmware and drivers are current on all hosts before continuing.

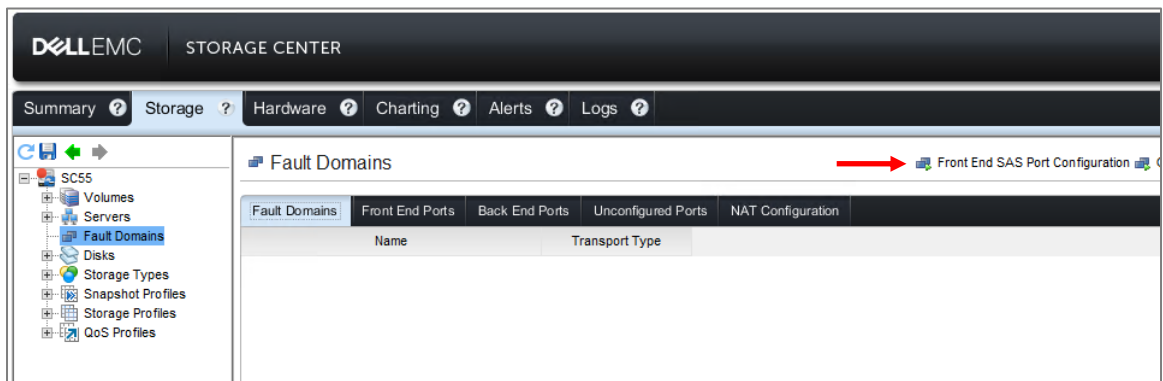
3.3 Configure SAS FE fault domains to support MPIO hosts

An SC5020 array is used in this example to show how to complete the configuration. The configuration steps for other SC arrays with SAS FE are similar.


- Launch the **Dell EMC Storage Manager Client** and connect to the SC Series array directly or to a Data Collector. In this example, the connection is made directly to the SC5020 by entering the management IP address.



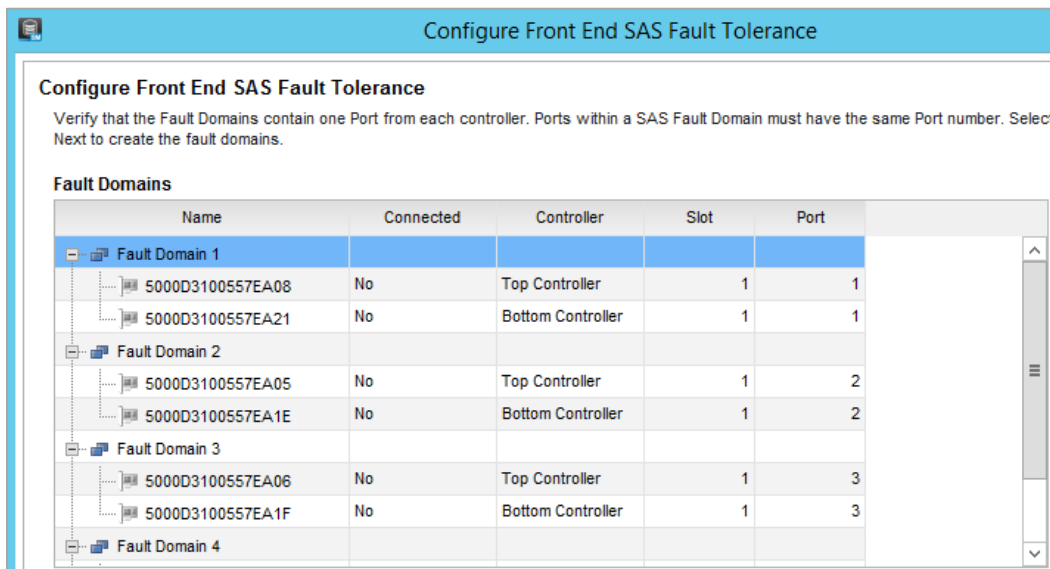
- Click **Storage > Fault domains**. In this example, no fault domains are configured on the SC5020. Click **Front End SAS Port Configuration** to start the configuration wizard.



- Set the number of fault domains to create to 4 (the maximum possible) and click **Next**.




- Review the fault domain information. A matching port on each SAS controller is automatically paired to create each fault domain. Click **Next**.

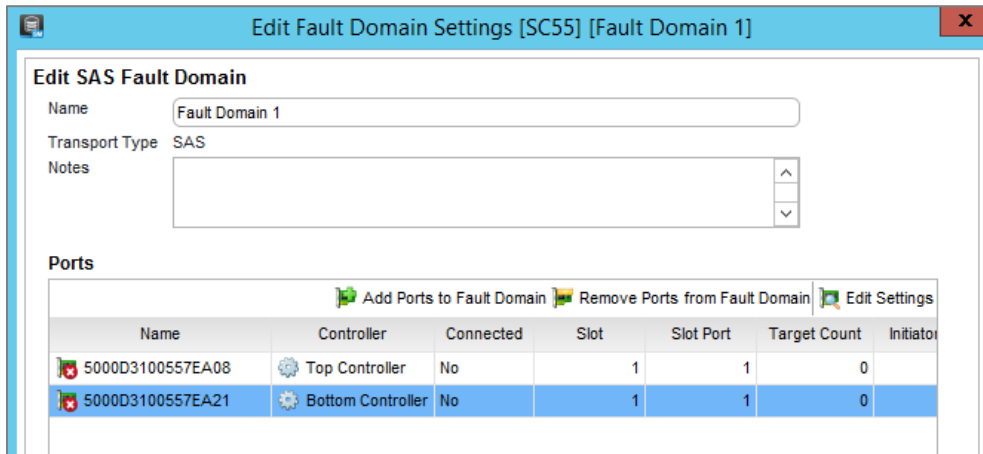


Note: SAS FE ports must be assigned to a fault domain before the ports become available to host servers.

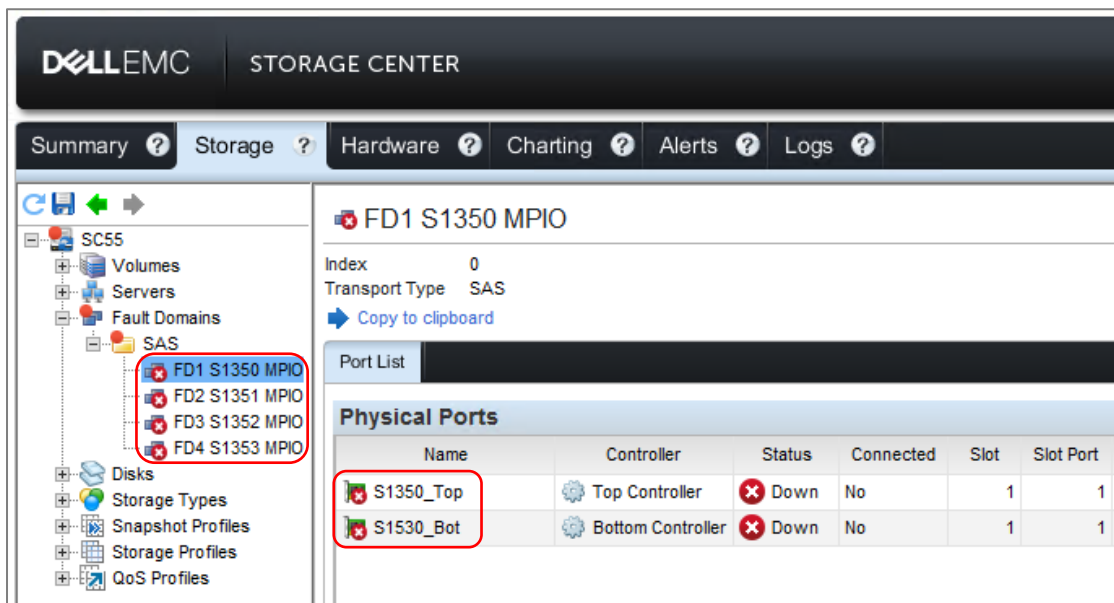
- Allow the wizard to complete and click **Finish**.
- Expand **Fault Domains > SAS** and click each fault domain to view the port details. The fault domains and ports will display a status of **Down** until the SAS cables are connected to host server SAS ports.



- Right-click a fault domain and select **Edit** to view additional information or to perform actions such as renaming the fault domain or assigning friendly names to each physical controller port.



- As a best practice, modify the fault domain names and port names. This enables intuitive administration in later steps and makes troubleshooting easier. In this example, a host server name is added to each component (the servers are named **S13xx**).



3.4 Connect Windows host servers to the SC Series array with SAS cables

The following example provides step-by-step guidance for configuring a two-node MPIO Hyper-V cluster using two Dell PowerEdge R630 servers, two Dell 12Gbps SAS HBAs, four SAS cables, and an SC5020 array equipped with SAS FE ports.


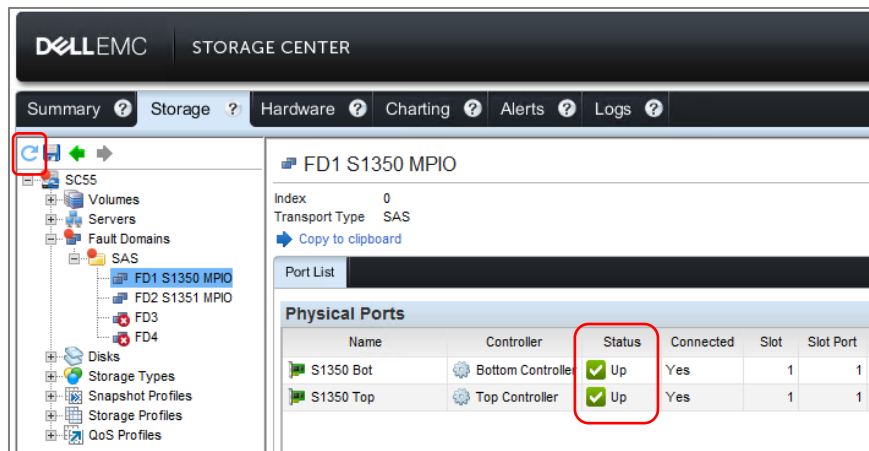


Figure 17 Configuration example with R630 hosts and SC5020

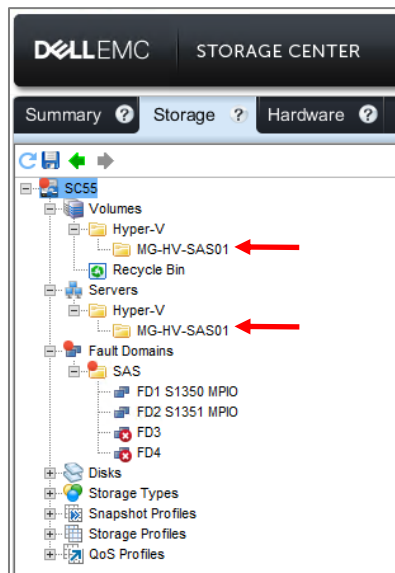
Modify these steps to fit the design of your environment.

1. Starting with the first Windows Hyper-V host server, connect a SAS cable from SAS port 1 on the server to SAS FE port 1 in the top controller of the SC Series array. The server does not need to be turned off before connecting or removing SAS cables.
2. Connect a second SAS cable from the other server SAS port to SAS FE port 1 on the bottom controller of the SC array. Make sure both host server SAS ports are cabled to the same SAS FE fault domain.
3. Repeat steps 1 and 2 for each additional host server until all are connected.
4. Refresh the view in DSM and verify that the connected SAS ports show a status of **Up**. Resolve any connectivity issues before continuing.

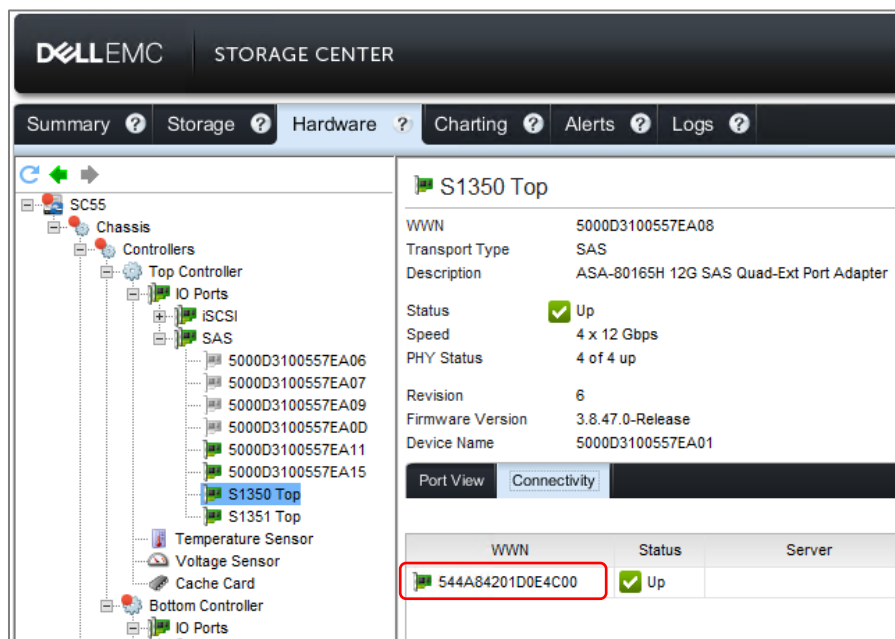


3.5 Create server objects

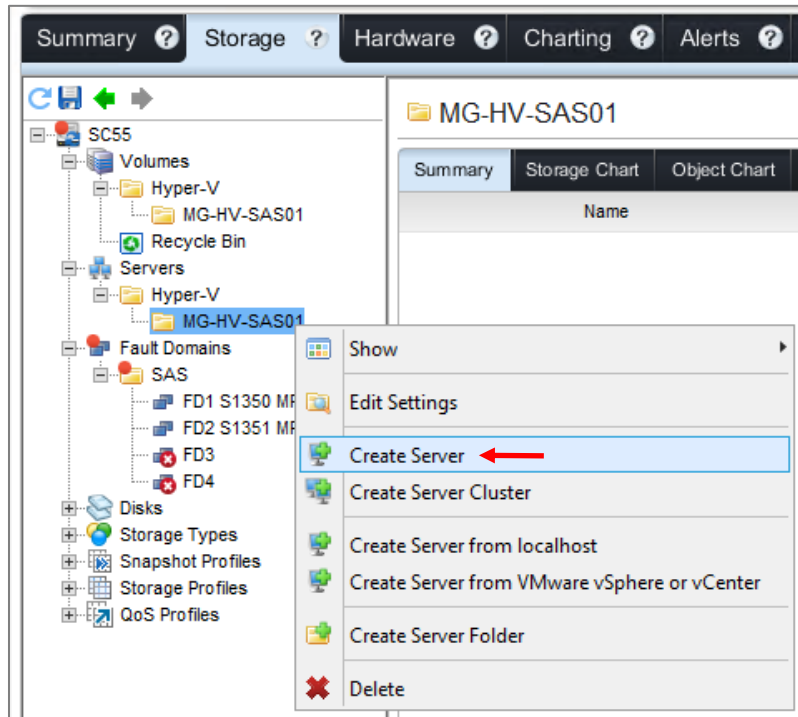
1. In the DSM client, click **Storage > Volumes** and create folders and subfolders to logically group your volumes. Do the same under **Storage > Servers** to logically group your server objects. In this example, a simple tree is created for the objects associated with a Hyper-V cluster named **MG-HV-SAS01**.



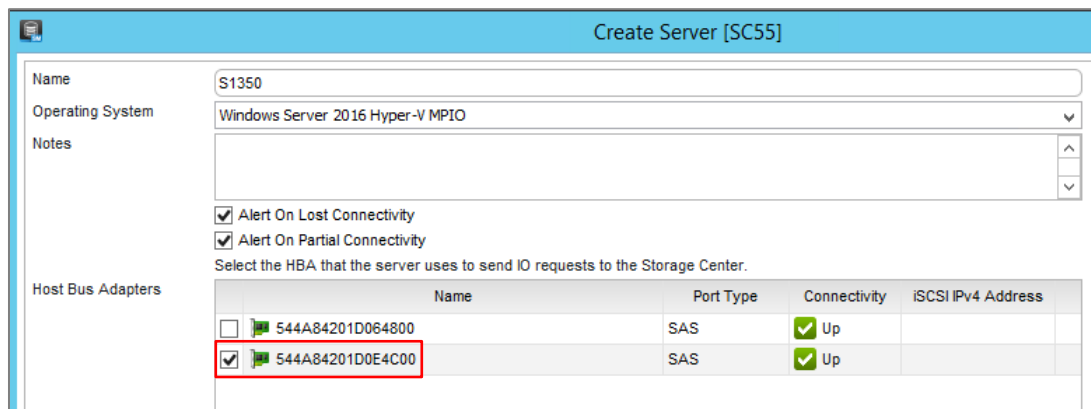
2. Click **Hardware**, expand **Chassis > Controllers > Top Controller > IO Ports > SAS**, and click the SAS port that is connected to the first host server (**S1350 Top** in this example). Ports that are assigned intuitive names before this step will make this process easier. Click the **Connectivity** tab and note the initiator WWN for the host SAS port (ending in **4C00** in this example).



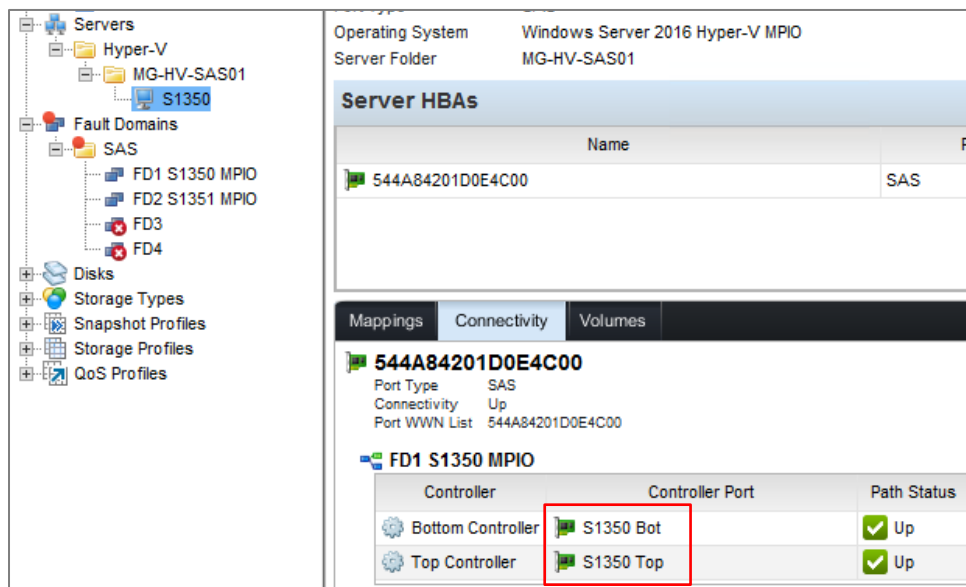
3. Repeat step 2 and note the initiator WWNs for the other host server SAS ports. In this example, the initiator WWN for port **S1351 Top** ends in **4801**.
4. Repeat the process for the second (bottom) controller. In this example, the initiator WWNs end in **4C01** and **4800**.
 - Host Server S1350: Top = **4C00** Bot = **4C01**
 - Host Server S1351: Top = **4801** Bot = **4800**
5. Under **Storage**, right-click the desired Servers subfolder and click **Create Server**.



6. In the wizard, configure the following:
 - a. Provide a name for the host server. In this example, the server is named **S1350**.
 - b. Select the correct MPIO operating system from the drop-down list. **Windows Server 2016 Hyper-V MPIO** is used in this example.
 - c. Use the information from step 4 to determine the correct initiator SAS port (HBA). In this example, the WWN ending in **4C00** is correct for the MPIO host **S1350**.
 - d. Click **OK**.

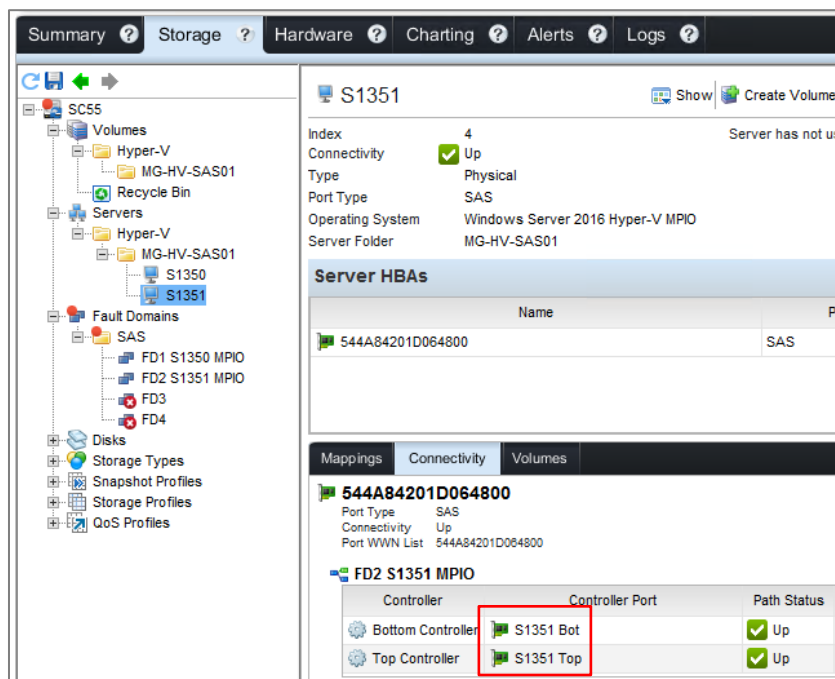


7. Click the new server object. Under the **Connectivity** tab, verify that the top and bottom controller ports are correct. Assigning intuitive names to these ports ahead of time facilitates easy verification.



8. Repeat the steps 5–7 to create additional host servers on the SC Series array. In this example, a second MPIO server is added named **S1351**.


- Verify that the correct controller ports are listed for each host server and that they are up before continuing.



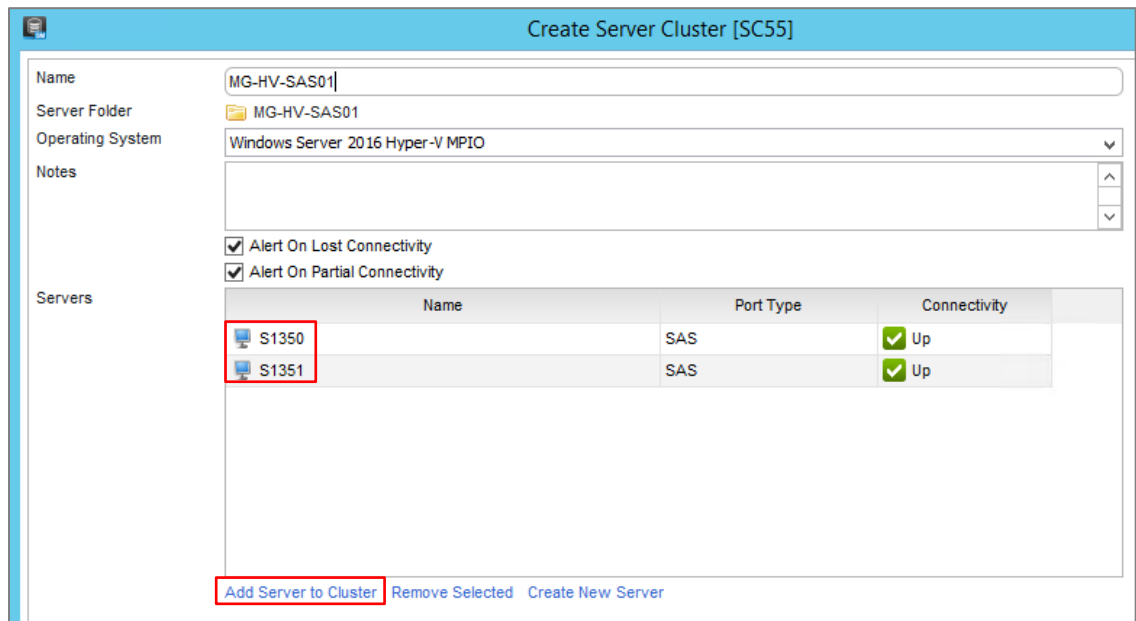
3.6 Create cluster server object on the storage array

To simplify managing cluster volumes on the SC Series array, create a server cluster object with the desired host servers as members of the cluster. In this example, the member servers are the two MPIO hosts, **S1350** and **S1351**.


- In DSM, right-click the desired server folder and select **Create Server Cluster**.



2. Click **Add Server to Cluster**, provide a name for the cluster, and add the desired hosts (in this example, **S1350** and **S1351**).



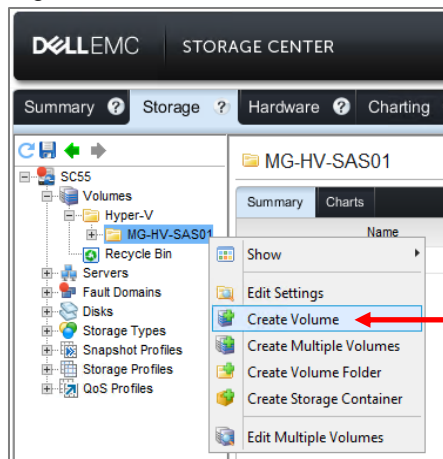
3. The selected hosts are now listed below the cluster server object.



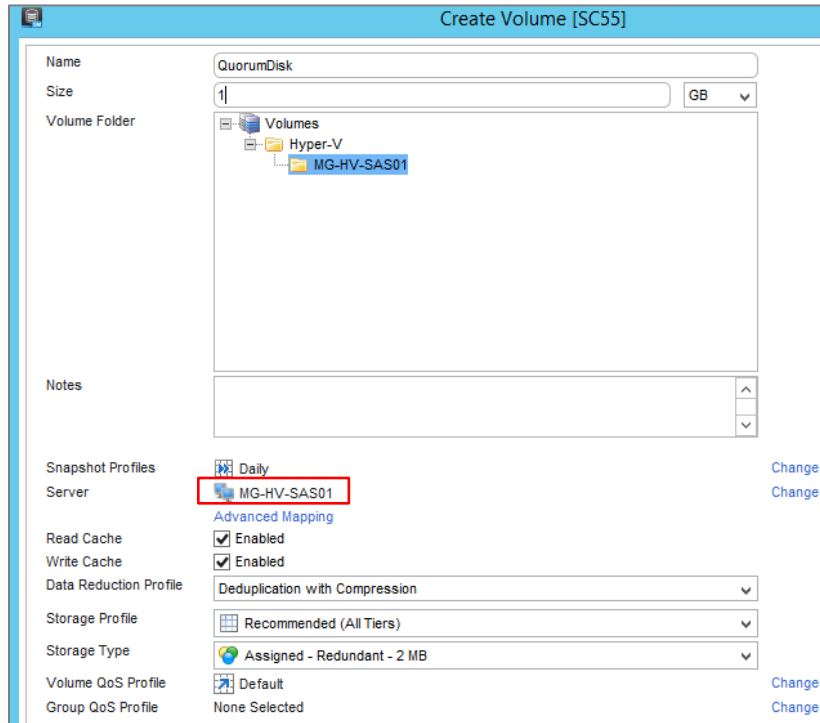
3.7 Create and map storage volumes to the server cluster

Now that the cluster object is created on the storage array, the next step is to create and map storage to the cluster object. In this example, a quorum disk and a cluster volume are created and mapped to the cluster server object.

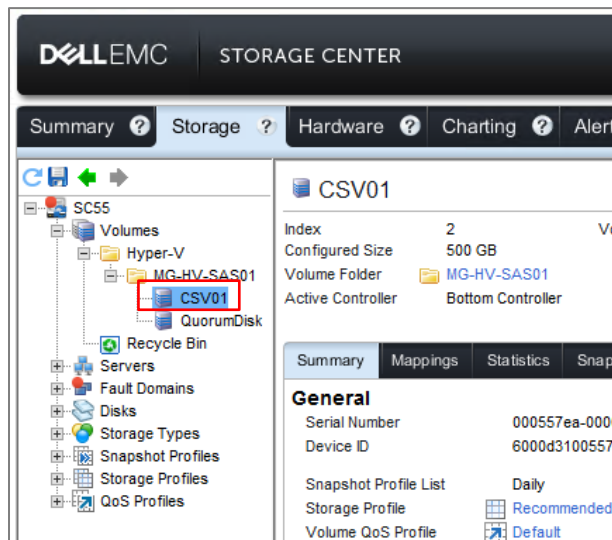
1. Right-click the desired **Volumes** subfolder and select **Create Volume**.



- In this example, the first volume created is a quorum disk. Provide an intuitive name for the volume, and set a volume size and a snapshot profile. For **Server**, select the cluster server object. Configure the other volume settings as desired, and click **OK**.



- Repeat steps 1 and 2 to create and map at least one additional data volume to the server cluster object. In this example, a 500 GB volume named CSV01 is created and mapped to the cluster object. This volume will be configured as a cluster shared volume for the new Hyper-V cluster.



- Click the server cluster object and look under the **Mappings** tab to view the two new volumes along with the mapping details. Each volume will have two paths listed for each host in the cluster for a total of four paths.

The screenshot displays the Dell EMC Storage Management Console interface. On the left is a navigation tree with categories like Servers, Disks, Storage Types, and Snapshots. The main area is divided into several sections:

- Servers:** A table listing two servers, S1350 and S1351, both connected via SAS ports with a status of 'Up'.
- Mappings > Volumes:** A table showing two volumes: CSV01 and QuorumDisk. CSV01 is mapped via the Server Cluster with LUN 11, while QuorumDisk is mapped via the Server Cluster with LUN 10.
- Mapping Details:** A detailed table for the CSV01 volume showing four paths, each associated with a specific server and controller port.

Name	Port Type	Connectivity
S1350	SAS	Up
S1351	SAS	Up

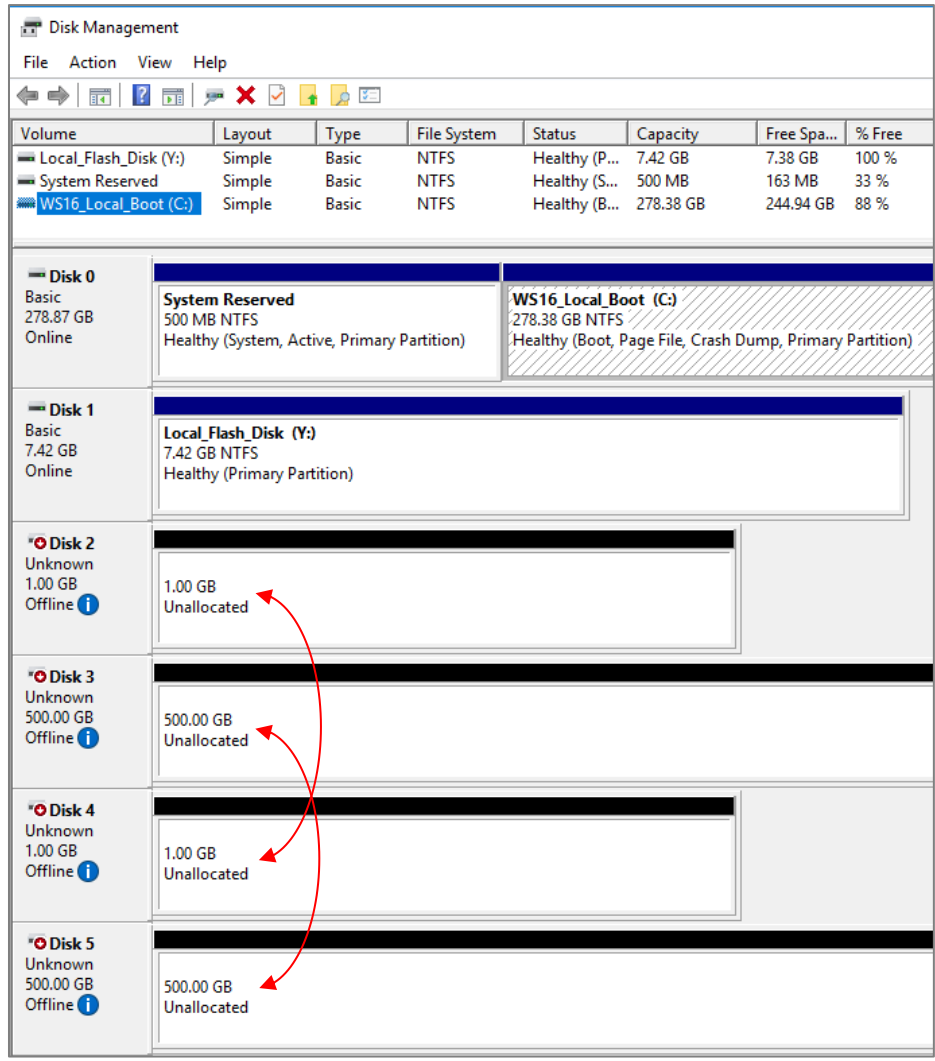
Volume	Connectivity	Volume Folder Path	Mapped Via	LUN Requested	LUN Used	Re
CSV01	Up	Hyper-V/MG-HV-SAS01/	Server Cluster	11	11	No
QuorumDisk	Up	Hyper-V/MG-HV-SAS01/	Server Cluster	10	10	No

Volume	Status	Transport	Server HBA	Controller Port	LUN	Read Only	Operational State
CSV01	Up	SAS	544A84201D0E4C00	S1350 Bot	11	No	Active/Optimized
CSV01	Up	SAS	544A84201D0E4C00	S1350 Top	11	No	Standby
CSV01	Up	SAS	544A84201D064800	S1351 Bot	11	No	Active/Optimized
CSV01	Up	SAS	544A84201D064800	S1351 Top	11	No	Standby

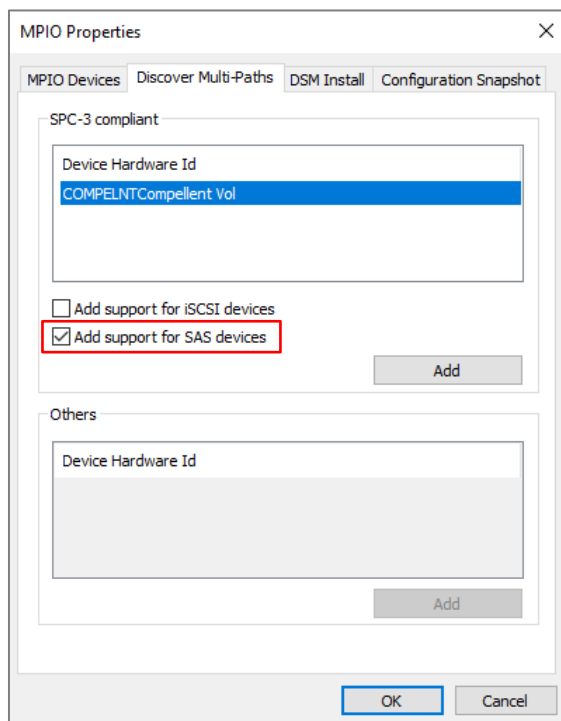
3.8 Configure MPIO on the host servers

Now that two volumes are mapped to both servers through the cluster server object, the next step is to enable MPIO on the Windows host servers.

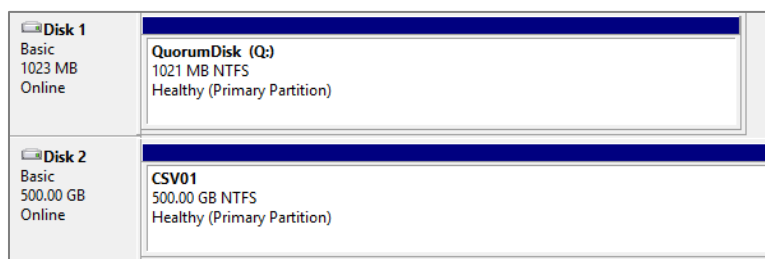
1. Log in to the first Windows host and launch **Disk Management**. From the **Action** drop-down menu, select **Rescan Disk**. Because MPIO has not yet been configured for SAS on the Windows host, Disk Management lists two instances for each disk instead of one instance.



- To correct this so that only one instance of each disk is shown, launch **MPIO Properties** on the Windows host. Under the **Discover Multi-Paths** tab, select **Add support for SAS devices** and click **Add**.



- When prompted, reboot the Windows host.
- After rebooting, launch **Disk Management** again and verify that only one instance of each drive is displayed.
- Initialize each disk and bring it online.
- Format each volume.
 - For the **Quorum Disk**, assign a drive letter of your choice (Q in this example) and assign an intuitive volume label such as QuorumDisk.
 - For the cluster volume, do not assign a drive letter or path when formatting the volume. Assign an intuitive volume label such as CSV01.



- Now that MPIO is enabled on the Windows host servers, follow the recommendations in the [Dell EMC SC Series Microsoft Multi-path I/O Best Practices Guide](#) to configure optimal MPIO settings.

Note: This a critical step to ensure that host server MPIO timeout thresholds are adjusted to tolerate paused I/O (up to 60 seconds) which is typical during planned maintenance such as staggered controller upgrades.

8. Repeat steps 1–4 and step 7 for each additional host in the cluster. In this example, the second host is named **S1351**. It is not necessary to initialize or format the volumes (steps 5 and 6) on other cluster nodes because these steps have already been completed on the first node.

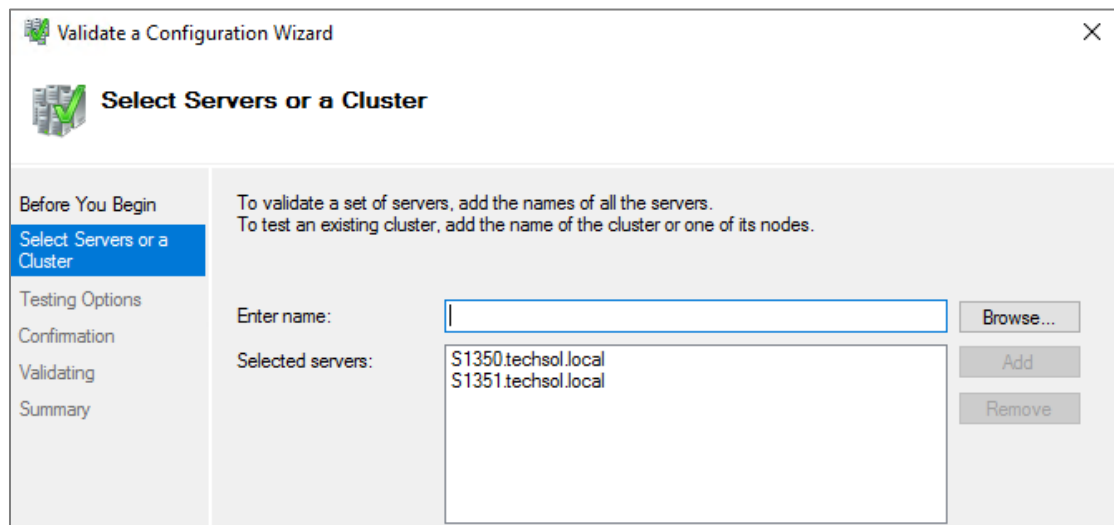
4 Create a Hyper-V cluster

The two configured Windows hosts in this example are now ready to form a new Windows Hyper-V cluster. This document assumes that the reader is already familiar with creating Hyper-V clusters. While it is beyond the scope of this document to cover all of the design considerations and best practices associated with designing and creating a Hyper-V cluster, the basic steps to create a new Hyper-V cluster from the two Windows hosts used in this paper are shown in this section. For more information about Hyper-V design and best practices, there are many resources online such as [Microsoft TechNet](#).

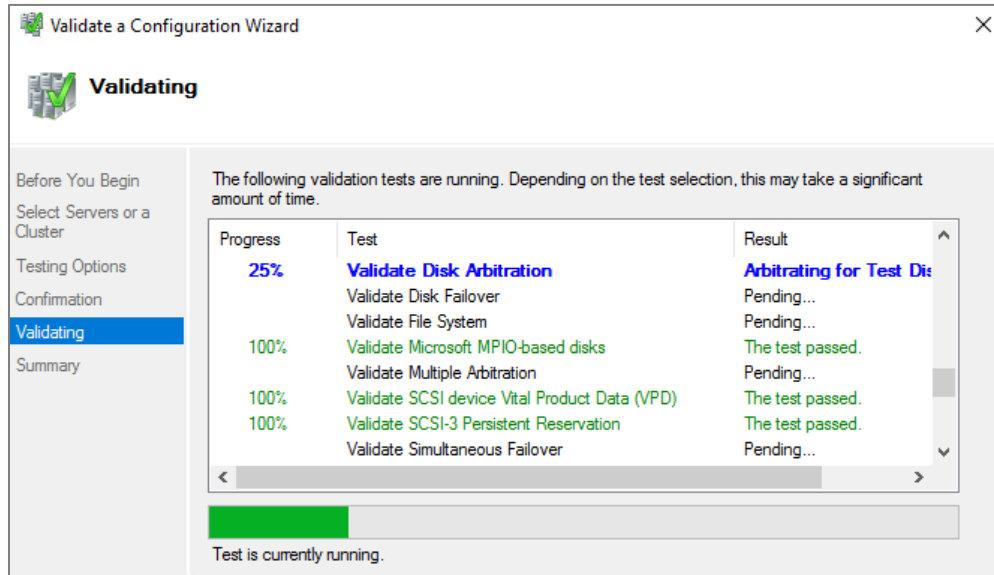
4.1 Validate the servers to be clustered

Run cluster validation to verify that the prospective nodes for a new cluster meet all the requirements for clustering. The validation process runs a number of tests and produces a detailed report showing the results. Correct any deficiencies that prevent clustering before proceeding. After correcting any deficiencies, re-run cluster validation to verify that the tests pass.

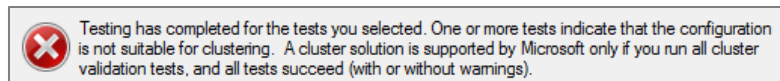
1. Log in to the first Windows node of the cluster and launch **Failover Cluster Manager**. In this example, **S1350** is the first node.
2. Under the **Actions** pane on the right, click **Validate Configuration** to launch the **Validate a Configuration Wizard**. Review the information on the **Before You Begin** page and click **Next**.
3. Add the servers to be clustered. In this example, **S1350** and **S1351** are added. Click **Next**.



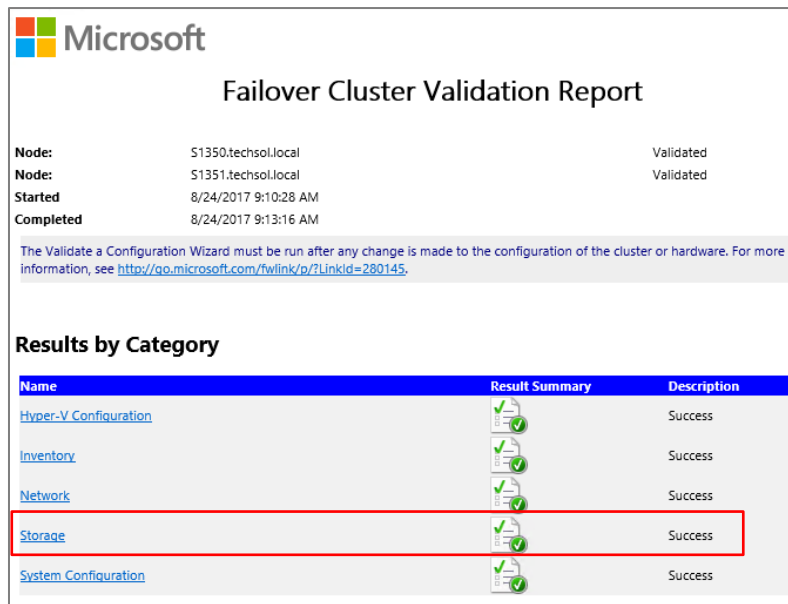
4. Select **Run All Tests** and click **Next** and allow the validation to complete.



5. Examine the results by viewing the report. If any failures are displayed that prevent clustering, correct these deficiencies and then re-run cluster validation. It may be difficult to pass all tests. Minor deficiencies will not prevent the nodes from being clustered.



6. In this example, all the tests pass. Of particular interest are the results shown under the Storage area of the report, since the storage resides on the SC Series array. Ensure that all the storage tests pass before continuing.



4.2 Create a new Hyper-V cluster

1. If the nodes are suitable for clustering, enable **Create the cluster now using the validated nodes** on the **Summary** screen and click **Finish** to start the Create Cluster wizard. The Create Cluster wizard can also be run from the **Actions** pane in Failover Cluster Manager if not continuing the steps from the previous section.
2. Provide a name for the new cluster along with an IP address. In this example, the cluster is named **MG-HV-SAS01**. Click **Next**.

Create Cluster Wizard

Access Point for Administering the Cluster

Before You Begin | **Access Point for Administering the Cluster** | Confirmation | Creating New Cluster | Summary

Type the name you want to use when administering the cluster.

Cluster Name:

i The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.

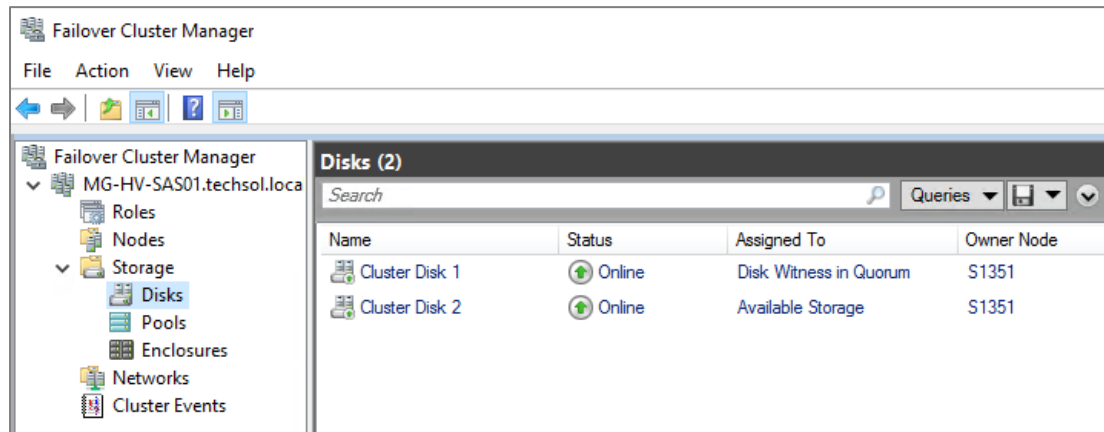
	Networks	Address
<input checked="" type="checkbox"/>	100.88.144.0/20	100 . 88 . 146 . 127

3. On the **Confirmation** screen, verify that **Add all eligible storage to the cluster** is selected. Click **Next**. Monitor the progress bar until it completes.
4. Review the **Summary** page and click **Finish**.

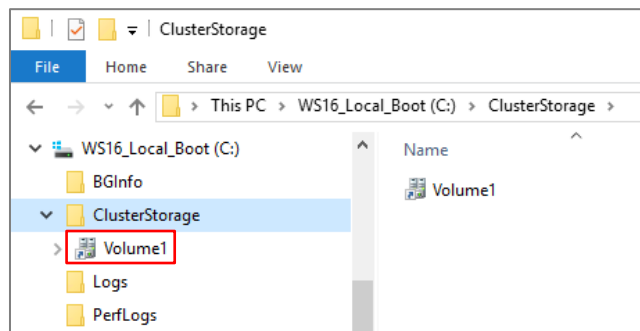
4.3 Convert the cluster disk to a cluster shared volume

As a final step, after creating a Hyper-V cluster, convert the 500 GB cluster disk to a cluster shared volume before creating VM roles.

1. Open **Failover Cluster Manager** and the new cluster will appear in the left pane. Click each of the objects in the tree (Roles, Nodes, Storage, Networking, and Cluster Events) and view information about each one in the center pane.
2. Expand **Storage** and click **Disks**. In this example, the cluster was created with two shared volumes from the SC Series array. The Create Cluster Wizard automatically assigned the smallest disk as a quorum disk, since this is a two-node cluster, and therefore a quorum witness is needed as a tiebreaker.



3. To convert the 500 GB volume to a cluster shared volume, right-click the volume and select **Add to Cluster Shared Volumes**. The volume will now be presented as a share named **Volume1** on both nodes under C:\ClusterStorage.



4. At this point, the cluster is ready to support a workload such as one or more guest VMs.
5. Create and assign additional cluster volumes from the SC Series array to the new cluster as needed.
6. Expand the cluster by adding Windows Server Hyper-V hosts as new nodes using additional SAS ports on the SC Series array.

5 Support for guest VMs with SAS pass-through disks

Use of pass-through (PT) disks is a legacy Hyper-V configuration that (while still supported) is discouraged by Microsoft and Dell EMC. If a guest VM has the MPIO feature installed, it is necessary to change the `AllowFullSCSICommandSet` attribute from “false” to “true” in order for a guest VM to support SAS volumes that are presented as PT disks. This step is not required if the guest VM does not have the MPIO feature installed.

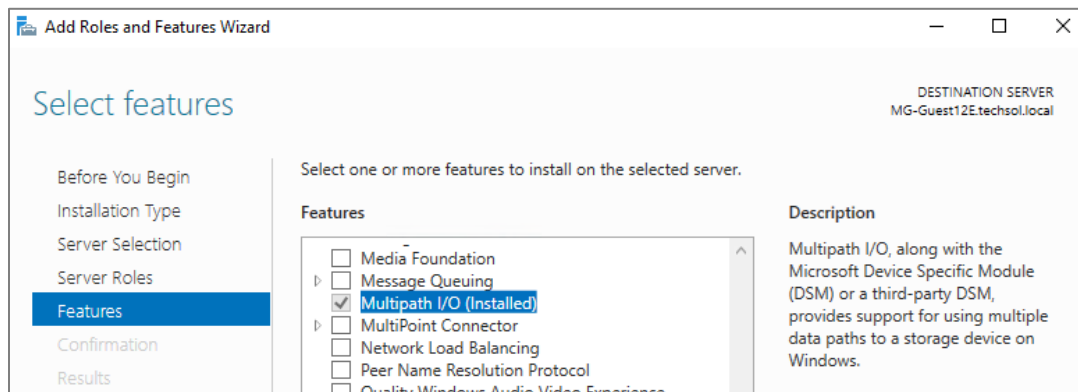


Figure 18 Guest VM with MPIO installed

Modifying the `AllowFullSCSICommandSet` attribute can be done with PowerShell.

Note: The creation of guest VM clusters that utilize SAS PT disks as cluster disks is not supported.

5.1 Windows Server 2008 R2 and 2012

For MPIO-enabled guest VMs running on Windows Server 2008 R2 or 2012, complete the following steps for each guest VM, substituting the name of the VM for `Vm_Name`. For guest VMs running on Windows Server 2012 R2 and newer, see section 5.2.

1. From the server that owns the VM, launch PowerShell with elevated (administrator) privileges, and run the following commands:

```
$HyperVGuestName = "Vm_Name";

$VMManagementService = gwmi -Class "Msvm_VirtualSystemManagementService" -
Namespace "root\virtualization"

$Vm = gwmi -Namespace "root\virtualization" -Query "Select * from
Msvm_ComputerSystem Where ElementName = '$HyperVGuestName'"

$SettingData = gwmi -Namespace "root\virtualization" -Query "Associators
of {$Vm} Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"

$SettingData.AllowFullSCSICommandSet = $true

$VMManagementService.ModifyVirtualSystem($Vm, $SettingData.GetText(1))
```

2. Reboot the VM.
3. Present one or more SAS pass-through disks to the guest VM.
4. Repeat steps 2–3 for additional VMs.

5.2 Windows Server 2012 R2 and newer

For MPIO-enabled guest VMs running on Windows Server 2012 R2 and newer, complete the following steps for each guest VM, substituting the name of the VM for `Vm_Name`. For guest VMs running on Windows Server 2008 R2 or 2012, see section 5.1.

1. From the server that owns the VM, launch PowerShell with elevated (administrator) privileges, and run the following commands:

```
$HyperVGuestName = "Vm_Name";

$VMMManagementService = gwmi -Class "Msvm_VirtualSystemManagementService" -
Namespace "root\virtualization\v2"

$Vm = gwmi -Namespace "root\virtualization\v2" -Query "Select * from
Msvm_ComputerSystem Where ElementName = '$HyperVGuestName'"

$SettingData = gwmi -Namespace "root\virtualization\v2" -Query
"Associators of {$Vm} Where ResultClass=Msvm_VirtualSystemSettingData"

$SettingData.AllowFullSCSICommandSet = $true

$VMMManagementService.ModifySystemSettings($SettingData.GetText(1))
```

2. Reboot the VM.
3. Present one or more SAS pass-through disks to the guest VM.
4. Repeat steps 2–3 for additional VMs.

A Additional resources

Dell.com/support is focused on meeting customer needs with proven services and support.

Dell EMC TechCenter is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware, and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC storage platforms.

Referenced or recommended Dell EMC publications:

- [Dell EMC SC Series Storage product page](#)
- [Dell EMC SC Series Storage with SAS Front-end Support for VMware vSphere](#)
- [Dell EMC SC Series Microsoft Multipath I/O Best Practices](#)
- [Dell EMC SC Series Virtual Fibre Channel for Hyper-V Demo Video](#)
- [Dell EMC SC Series Replay Manager 7 and Hyper-V Demo Video](#)
- [Windows Server 2016 Best Practices Guide for Dell EMC SC Series Storage](#)
- [Windows Server 2012 R2 Best Practices Guide for Dell EMC SC Series Storage](#)
- [Dell EMC SC Series Data Reduction with Deduplication and Compression](#)
- [Dell EMC SC Series Storage: Synchronous Replication and Live Volume](#)
- [Dell EMC SC Series Live Volume with Auto Failover Support for Microsoft Demo Video](#)
- [Dell EMC SC Series Predefined DR Plans for Hyper-V Demo Video](#)
- [Dell EMC SC Series PowerShell SDK Cookbook](#)
- [Dell EMC SC Series and SMI-S Integration with Microsoft SCVMM](#)
- [Dell EMC SC Series Microsoft System Center VMM Rapid Provisioning Demo Video](#)

Referenced or recommended Microsoft publications:

- [Microsoft TechNet Library](#)
- [Microsoft PowerShell Developer Network \(MSDN\)](#)