Dell EMC Microsoft Storage Spaces Direct Ready Node

Operations Guide for managing and monitoring Ready Node life cycle

DELLEMC

Notes, cautions, and warnings

ⓘ | **NOTE: A NOTE indicates important information that helps you make better use of your product.**

△ | **CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.**

⚠ | **WARNING: A WARNING indicates a potential for property damage, personal injury, or death.**

Contents

Audience and scope

This Operations guide focuses on operational aspects of a hyper-converged infrastructure solution on Dell EMC Microsoft Storage Spaces Direct Ready Nodes for Microsoft Windows Server 2016 with Hyper-V and Storage Spaces Direct. This operations guide includes an overview of the Dell EMC Microsoft Storage Spaces Direct Ready Nodes, guidance to monitor and manage bare-metal, and instructions on performing operations on Storage Spaces Direct cluster and updating the cluster-aware system. This guide is applicable only to infrastructure built by using the validated and certified Dell EMC Ready Nodes for Windows Server 2016 with Hyper-V.

The audience for this document includes, but is not limited to, systems administrators, systems engineers, field consultants, partner engineering team members, and customers with a fair amount of knowledge in deploying hyper-converged infrastructures with Microsoft Windows Server 2016 Hyper-V and Storage Spaces Direct.

ⓘ **NOTE: Instructions in this operations guide are applicable only to the generally available OS build of Windows Server 2016 with the latest applicable updates. These instructions are not validated with Windows Server version 1709. Storage Spaces Direct Ready Nodes do not support the Windows Server Semi-Annual Channel release.**

Topics:

· Assumptions
· Known issues

Assumptions

This operations guide makes certain assumptions about the necessary prerequisite knowledge of the deployment personnel. This includes the prerequisite knowledge of:

· Dell EMC Microsoft Storage Spaces Direct Ready Nodes and deploying and configuring BIOS and iDRAC settings on these ready nodes.
· Deploying and configuring Windows Server 2016 Hyper-V infrastructure.
· Deploying and configuring System Center products such as SCOM, if the deployment involves onboarding Storage Spaces Direct cluster into the existing System Center.

Known issues

For a list of known issues and workaround, see http://en.community.dell.com/techcenter/extras/w/wiki/12392.known-issues.

Dell EMC Microsoft Ready Nodes overview

Dell EMC Microsoft Storage Spaces Direct Nodes encompass various configurations of R740xd and R640 Storage Spaces Direct Ready Nodes to power the primary compute cluster deployed as hyper-converged infrastructure. This hyper-converged infrastructure built by using the Dell EMC Ready Nodes uses a flexible solution architecture rather than a fixed component design. The following figure illustrates one of the flexible solution architectures consisting of compute cluster alongside the redundant top of rack switches, a separate out of band network, and an existing management infrastructure in the data center.

This hyper-converged virtualized solution based on Dell EMC Microsoft Storage Spaces Direct Nodes is available in both hybrid and all-flash configurations. For more information on available configurations, see Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview.
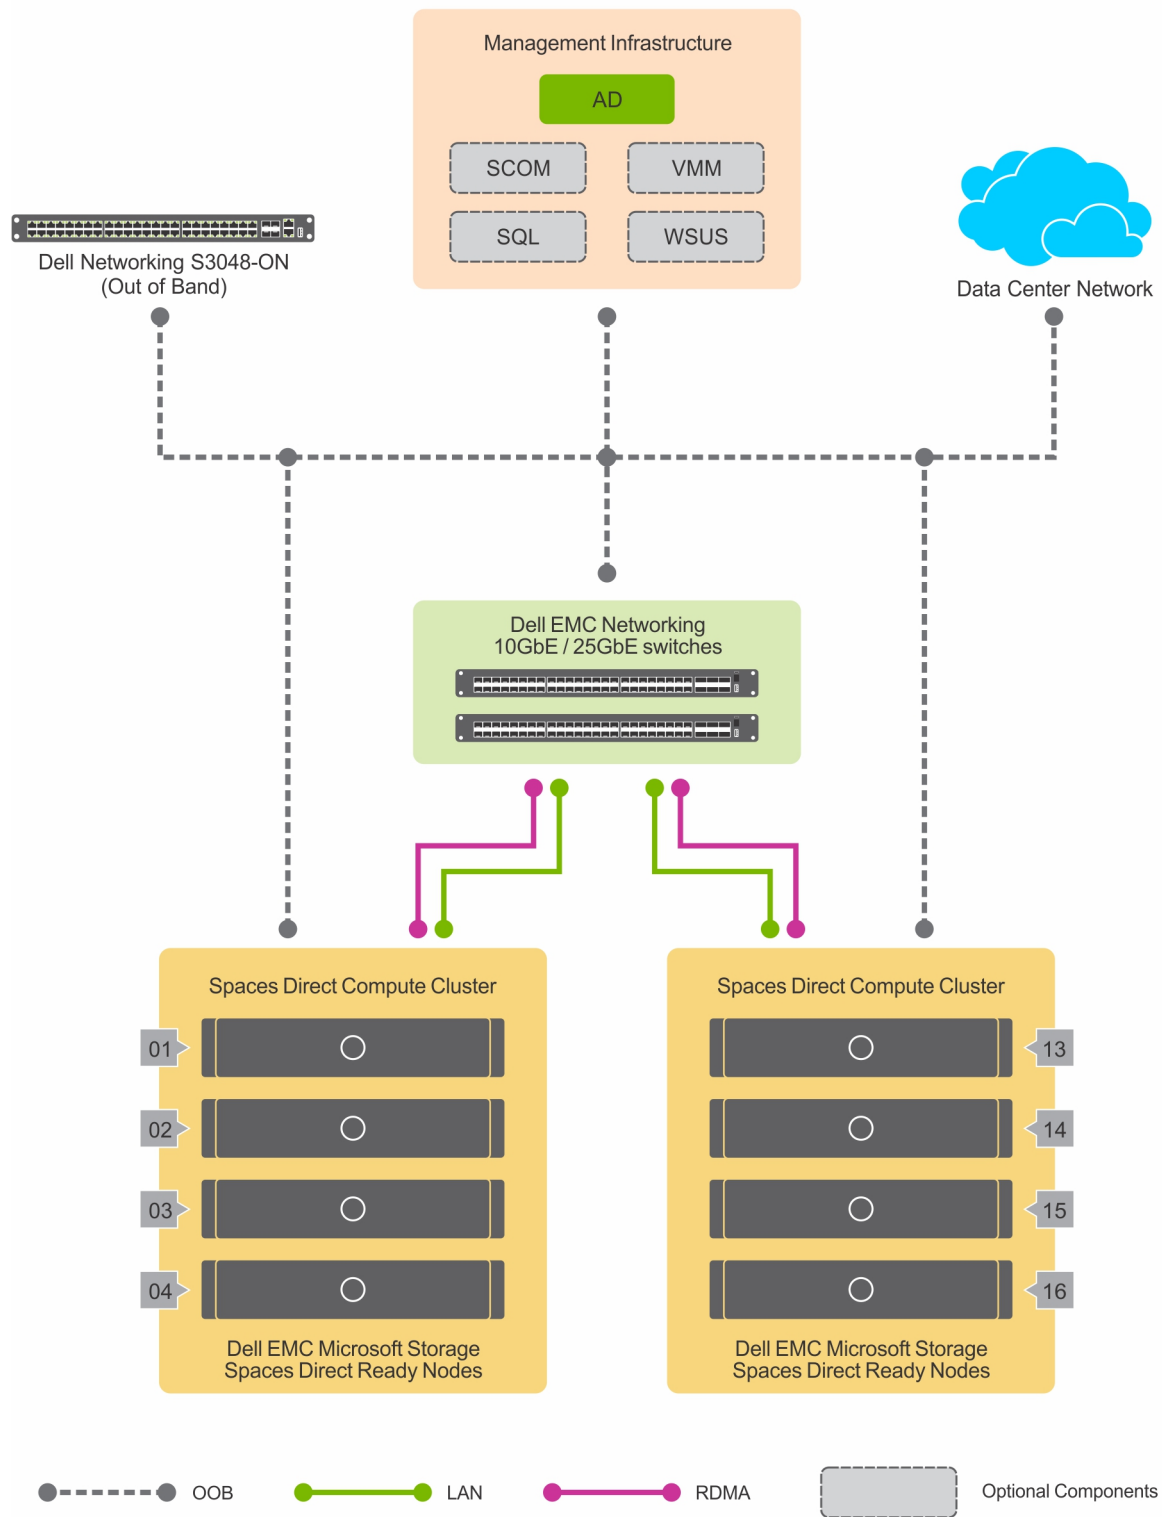
Figure 1. Hyper-converged virtualized solution using Dell EMC Ready Nodes

Deployment guidance

The deployment guidance and instructions for deploying and configuring a cluster using the Dell EMC Microsoft Storage Spaces Direct Ready Nodes is available in http://en.community.dell.com/techcenter/extras/m/white_papers/20444544 and http://en.community.dell.com/techcenter/extras/m/white_papers/20445029. This operations guidance is applicable only to cluster infrastructure that is built using the instructions provided in the Ready Node deployment guidance.

Day 0 checklist

After the Storage Spaces Direct Ready Node cluster is deployed, the day 0 operations must be completed. This section provides the details about the day 0 operations checklist.

Topics:

Day 0 OS updates

Before creating any volumes and deploying any virtualized workloads on the HCI cluster, Dell EMC recommends to perform day 0 updates for the operating system.

In the Server Core operating system, *sconfig* command can be used to enable auto-update for the operating system components. This, however, requires internet connectivity to Windows Update services for downloading and installing applicable updates.

If there is a Windows Software Update Services (WSUS) server in the data center, Group Policy objects can be used to configure the Server Core or Full OS systems in the Storage Spaces Direct HCI cluster to retrieve and download updates from the local WSUS Server.

Since the cluster is operational while performing Day 0 updates, it is recommended to perform cluster-aware updates of the OS. See Cluster-aware updating overview for more information.

Windows activation

When the Server Core operating system is installed using the retail or volume licensing media, the OS license needs to be activated. On the Server Core OS, this can be done by using either the *sconfig* tool or the *slmgr* command.

ⓘ | **NOTE: Windows activation is not required if the OS is factory installed.**

For OS license activation by using *slmgr*, see https://technet.microsoft.com/en-us/library/dn502540(v=ws.11).aspx.

For OS license activation by using *sconfig* command, see https://technet.microsoft.com/en-us/windows-server-docs/get-started/sconfig-on-ws2016.

See the Microsoft documentation on using KMS for volume activation of Windows OS in the data center.

Storage Spaces Direct Management

The following topic describes the concept of Storage Spaces Direct Management.

Creating virtual disks

Cluster creation and enabling Storage Spaces Direct configuration on the cluster creates only a storage pool and does not provision any virtual disks in the storage pool. The New-Volume cmdlet can be used to provision new virtual disks and in turn provision them as the cluster shared volumes for the Storage Spaces Direct HCI cluster.

(i) **NOTE: When creating virtual disks, Dell EMC recommends using 3 copy mirror as the resiliency setting on clusters with 3 or more nodes. There are existing known issues with multi-resilient volumes and therefore it is not recommended to use MRV volumes.**

Before creating volumes in the Storage Spaces Direct cluster infrastructure:

- Ensure that you create multiple volumes. This should be a multiple of number of servers in the cluster. Limit the number of volumes in the cluster to 32.
- Ensure that there is enough reserve capacity left in the storage pool for any in-place volume repairs arising out of failed disk replacement. The reserved capacity should at least be equivalent to the size of one capacity drive per server and up to 4 drives.

For general guidance on planning for volume creation in Storage Spaces Direct, see https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/plan-volumes.

Onboarding Ready Nodes in SCOM

System Center Operations Manager (SCOM) is an infrastructure monitoring tool which, among other things, helps manage alerts, warnings and other items through your environment.

Dell EMC Ready Nodes can be monitored using the operations manager management packs associated with the bare metal hardware and Storage Spaces Direct cluster.

Monitoring a Windows Server 2016 based cluster and Storage Spaces Direct performance requires the following management packs:

Table 1. Management packs

Management pack name	Version	Download location
SC Management Pack for Windows Server 2016.msi	10.0.8.0	https://www.microsoft.com/en-us/download/details.aspx?id=54303
SC Management Pack for Cluster 2016.msi	10.0.6.0	https://www.microsoft.com/en-us/download/details.aspx?id=48256
Microsoft System Center 2016 MP for Storage Spaces Direct.msi	10.0.40.0	https://www.microsoft.com/en-us/download/details.aspx?id=54700

(i) **NOTE: An older version of System Center Management Pack for the Windows Server 2016 Operating System comes bundled with SCOM 2016 RTM install. This can be upgraded to the latest version using Updates and Recommendations workspace under Administration Pane -> Management Packs in the SCOM console. This lists all updates available for any of the installed management packs.**

(i) **NOTE: Instructions for installing System Center Management Pack for Windows Server Failover Cluster and System Center Management Pack for Microsoft Windows Server 2016 Storage Spaces Direct are available along with management pack download.**

Post installation, the management packs need to be imported. For information on importing SCOM management packs, see How to Import an Management Pack in Operations Manager 2016.

Before the Storage Spaces Direct cluster nodes can be monitored, they need to be discovered as agent managed systems. This process involves installing the agents on the nodes that need to monitored. The push method of installing agents using the SCOM discovery wizard,

by default, uses the management server action account specified during SCOM deployment. Therefore, ensure that this account has local administrative privileges on the target nodes.

Run as Profile and run as Accounts

The SCOM monitoring processes on the target nodes use an agent action account associated with a Run As profile for monitoring and collecting management data. Each management pack requires these accounts with certain privileges. The following table lists the Run As Account requirements for the management packs to monitor Windows Server 2016 and Storage Spaces Direct cluster performance.

Table 2. Management packs

Management pack name	Run As Profile and Run As Account Requirements
System Center Management Pack for the Windows Server 2016 Operating System	NA
System Center Management Pack for Windows Server Failover Cluster	When using low privilege account, Windows Cluster Action Account Run As profile must have administrative credentials to access the cluster. Microsoft recommends that this account is configured for more secure distribution.
System Center Management Pack for Microsoft Windows Server 2016 Storage Spaces Direct	Microsoft Storage Library: RunAs Account Run As Profile requires a Run As account with Cluster Administrator privileges. Microsoft recommends that this account is configured for more secure distribution.

Dell EMC recommends using more secure distribution of cluster administrator credentials to the managed nodes. For the **Run As** Profiles that are configured with more secure credential distribution, the credentials must then be manually distributed to the target nodes where needed. Follow the steps listed below to enable secure distribution of Run As Account credentials to target nodes.

1 Open the Operations Manager console and click **Administration**.
2 In the **Administration** window, click **Accounts** under **Run As** configurations.
3 Double-click the **Run As Account** that needs secure distribution.
4 Select the **Distribution** tab.
5 Click **Add** and then click **Search**.
6 Select the **target node (2)** to which the credentials must be distributed. Click **Add** and click **OK**.
7 Click **Apply** and then **OK**.

Discovering Microsoft Windows Server 2016 cluster nodes

The SCOM discovery wizard can be used to perform node discovery and Microsoft monitoring agent install. For the Storage Spaces Direct cluster discovery and monitoring, the agent on the nodes should be configured as proxy. Follow the steps listed below to enable the agent proxy setting on all agents participating in a Windows Failover Cluster.

1 Open the Operations Manager console and click **Administration**.
2 In the Administration Pane, click **Agent Managed**.
3 Double-click an agent in the list.
4 Select the **Security** tab.
5 Select the option **Allow** this agent to act as a proxy and discover managed objects on other computers.
6 Repeat steps 3 through 5 for each agent that is installed on a clustered server.

After the cluster node and cluster discovery is complete, the cluster view under **Monitoring Pane** > **Microsoft Windows Cluster** > **Clusters** are updated.

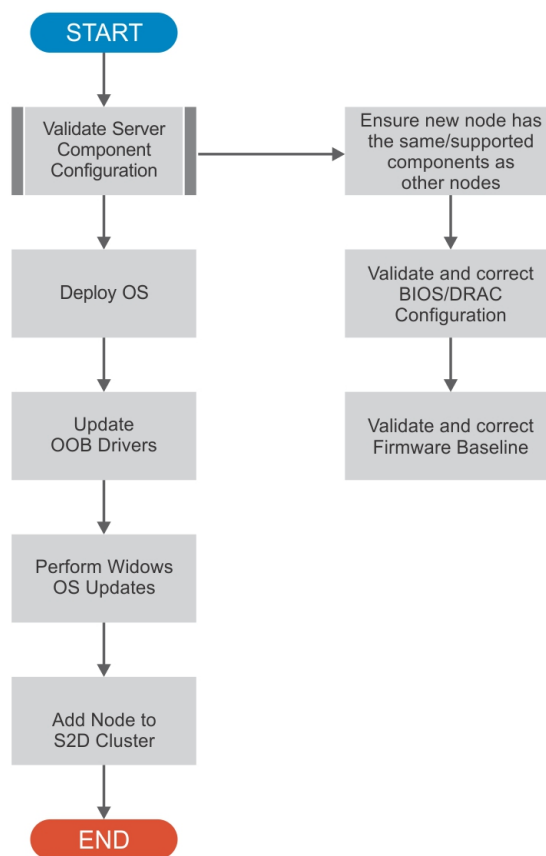ⓘ NOTE: It may take up to 24 hours for this view to get generated completely.

Figure 2. Discovering Microsoft Windows Server 2016 cluster nodes

Dell EMC Microsoft Storage Spaces Direct Ready Node Life Cycle Management

OpenManage Essentials is a hardware management application that provides a comprehensive view of systems, devices, and components in the enterprise's network. With OpenManage Essentials, a web-based and one-to-many systems management application for systems and other devices. To perform the platform updates using OpenManage Essentials:

1   Discover and inventory the system.
2   Monitor the health of the system.
3   View and manage system alerts.
4   Perform system updates and remote tasks.
5   View hardware inventory and compliance reports.
6   Manage the configuration baseline of a server or chassis.

Topics:

Management of Dell EMC Microsoft Storage Spaces Direct Ready Nodes

Administrators can use OME to perform the following operations of managing the Storage Spaces Direct Ready Nodes.

Discovery and inventory of Storage Spaces Direct Ready Nodes: Storage Spaces Direct Ready Nodes can be discovered and inventoried using OME. After discovery and inventory of Ready Nodes in a storage spaces direct cluster, complete hardware configuration information (including disk information), BIOS and firmware versions of all the nodes are available from within a single OME console.

System lockdown: System lockdown is a new capability introduced in 14th generation of Dell EMC PowerEdge servers that prevents unauthorized or unintentional modification of the system configuration. Administrators can put a Storage Spaces Direct Ready Node under System Lockdown to prevent fix and maintenance tasks such as power operations from executing.

Firmware updates and compliance on a Storage Spaces Direct Ready Node: OME can be used to simplify firmware updates of all the ready nodes. Using a catalog, all the Ready Nodes can be updated seamlessly to a qualified set of BIOS firmware.

Enforce device compliance on a Storage Spaces Direct Ready Node: OME can be used to track BIOS and other system configuration. A system configuration baseline can be first established which saves all the system configuration. Whenever the configuration needs to be updated, the baseline is updated and OME can be used to apply this update to all the nodes in the storage spaces direct cluster.

DELLEMC

Discovery and inventory of Storage Spaces Direct Ready Nodes in OME

OME supports agentless discovery and inventory of Storage Spaces Direct Ready Nodes by communicating directly with the server's integrated Dell Remote Access Controller (iDRAC) using WS-MAN protocol.

Accessing OME CLI prompt

To access the OME CLI prompt, click **Start** > **Dell OpenManage Applications** > **Essentials Command Line Interface**.

Install Hyper-V and failover cluster management tools

To manage the S2D cluster/Hyper-V nodes remotely from the OME server, you must install the following management tools on the OME server.

Hyper-V management tools

```
Install-WindowsFeature -Name Hyper-V-PowerShell -Verbose
```

```
Install-WindowsFeature -Name Hyper-V-Tools -Verbose
```

Failover Cluster management tools

```
Install-WindowsFeature -Name Failover-Clustering -IncludeAllSubFeature -IncludeManagementTools -Verbose
```

Creating a discovery range group

Discovery configuration range group PowerShell commands require a discovery profile template and range list template that defines the various parameters and the associated protocol values and are submitted as part of the CLI argument input.

OME installs a sample discovery profile template and range list template files that is located in **C:\Program Files\Dell\SysMgmt\Essentials\Tools\CLI\ and it is named as DiscoveryProfile.xml & RangeList.csv**.

Discovery profile template

Discovery profile template defines the entire variable set that can be entered by using the UI discovery configuration wizard. This template can be individually modified to turn on the desired protocols.

Since the Storage Spaces Direct Ready Nodes are discovered using "out-of- band" network, the following sections were modified in the discovery profile template.

```
"Discovery Configuration "Enabled=True"

Modify subnet mask example "255.255.224.0"

WSmanconfig Enabled=True"

Modify the Idrac username "Example: root"
```

```
  DiscoveryProfile.xml
  1    <?xml version="1.0" encoding="utf-8" ?>
  2    <DiscoveryConfiguration Enabled="True">
  3      <NetMask>
  4        255.255.224.0
  5      </NetMask>
  6      <ICMPConfig>
  7        <Timeout>600</Timeout>
  8        <Retries>1</Retries>
  9      </ICMPConfig>
 10      <SNMPConfig Enabled="False">
 11        <SNMPV1V2CConfig Enabled="True">
 12          <GetCommunity>public</GetCommunity>
 13          <SetCommunity></SetCommunity>
 14        </SNMPV1V2CConfig>
 15        <SNMPV3Config Enabled="True">
 16          <SNMPV3Username>user1</SNMPV3Username>
 17          <SNMPV3AuthenticationProtocol>SHA1</SNMPV3AuthenticationProtocol>
 18          <SNMPV3EncryptionProtocol>AES</SNMPV3EncryptionProtocol>
 19        </SNMPV3Config>
 20        <Timeout>4</Timeout>
 21        <Retries>2</Retries>
 22      </SNMPConfig>
 23      <WMIConfig Enabled="False">
 24        <UserName>Administrator</UserName>
 25      </WMIConfig>
 26      <StoragePowerVaultConfig Enabled="False">
 27
 28      </StoragePowerVaultConfig>
 29      <StorageEMCConfig Enabled="False">
 30        <UserName>Administrator</UserName>
 31        <Port>443</Port>
 32      </StorageEMCConfig>
 33      <WSManConfig Enabled="True">
 34        <UserName>root</UserName>
 35        <Timeout>4</Timeout>
 36        <Retries>6</Retries>
 37        <Port>443</Port>
 38        <SecureMode Enabled="True" SkipNameCheck="True" TrustedSite="True">
 39          <CertificateFile></CertificateFile>
 40        </SecureMode>
 41      </WSManConfig>
 42      <IPMIConfig Enabled="False">
 43        <UserName>Administrator</UserName>
 44        <KGKey></KGKey>
 45        <Timeout>5</Timeout>
 46        <Retries>2</Retries>
 47      </IPMIConfig>
 48      <SSHConfig Enabled="False">
 49        <UserName>Administrator</UserName>
 50        <Timeout>5</Timeout>
 51        <Retries>2</Retries>
 52        <Port>22</Port>
 53      </SSHConfig>
 54    </DiscoveryConfiguration>
```

Figure 3. Discovery profile

Range list template

The commands that create or modify discovery configuration groups can use a file that defines the ranges as an input parameter. This can be a comma-delimited list or ranges.


```
1    Name,SubnetMask,DisplayName,
2    172.18.129.4,255.255.224.0,Node1,
3    172.18.129.5,255.255.224.0,Node2,
4    172.18.129.6,255.255.224.0,Node3,
5    172.18.129.7,255.255.224.0,Node4,
```

Figure 4. Range list template

In this case, the input list is a comma-delimited file that initially defines the name/IP of the Storage Spaces Direct Ready Nodes for the discovery.

Add discovery range group

After the discovery profile template and range list template files are updated with the necessary details, run the following PowerShell commands to add adiscovery group using OME CLI.

```
$path = "C:\Program Files\Dell\SysMgt\Essentials\Tools\CLI\Samples"

Add-DiscoveryRangeGroup -GroupName "Group Name" -Profile $path\DiscoveryProfile.xml -
RangeListCSV $path\RangeList.csv
```
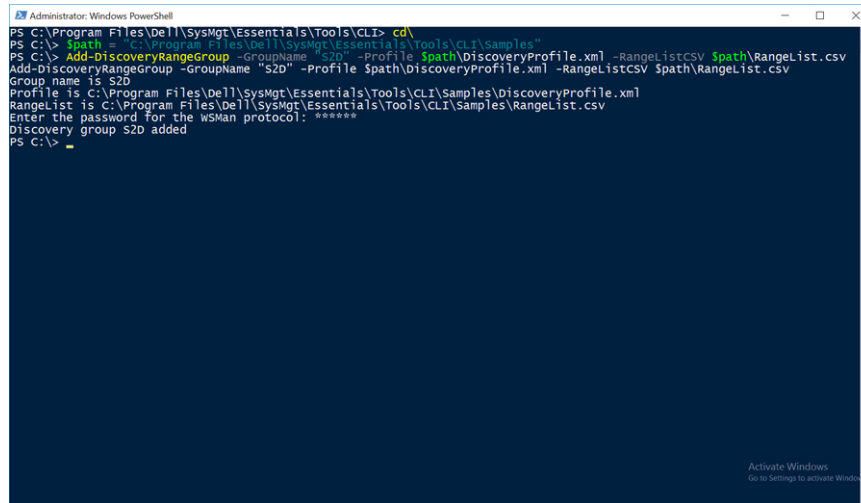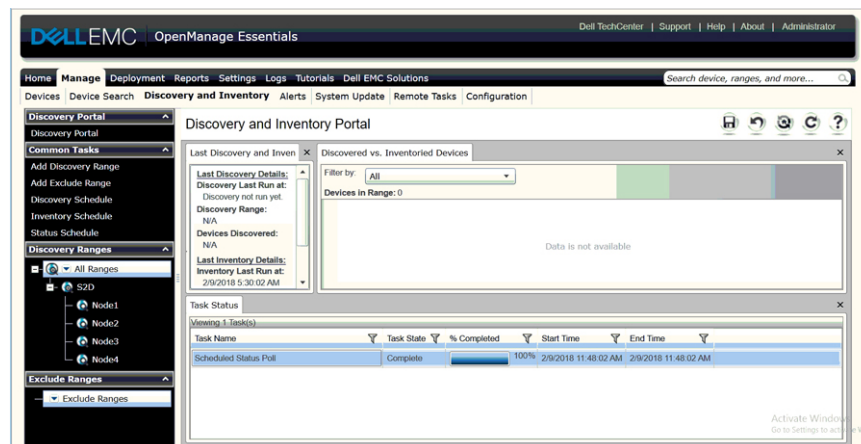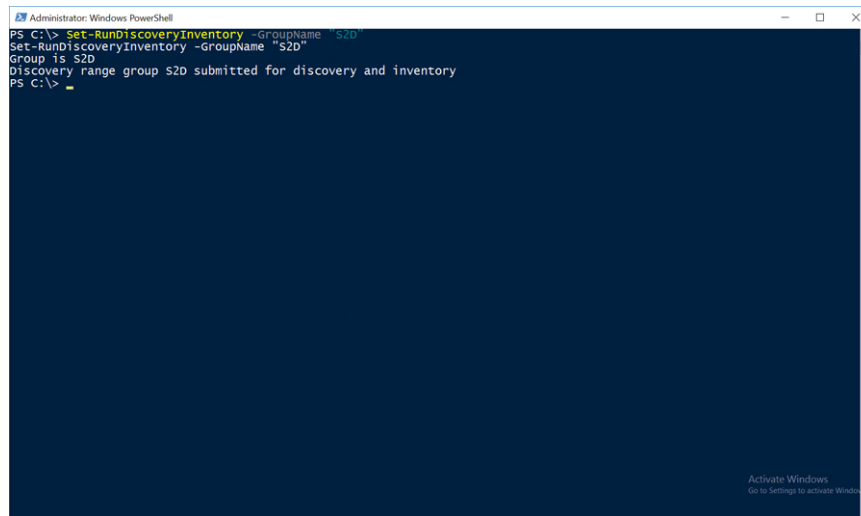
When prompted, enter the WSMAN (iDRAC) credentials.

Figure 5. Discovery range

Figure 6. Discovery and inventory portal
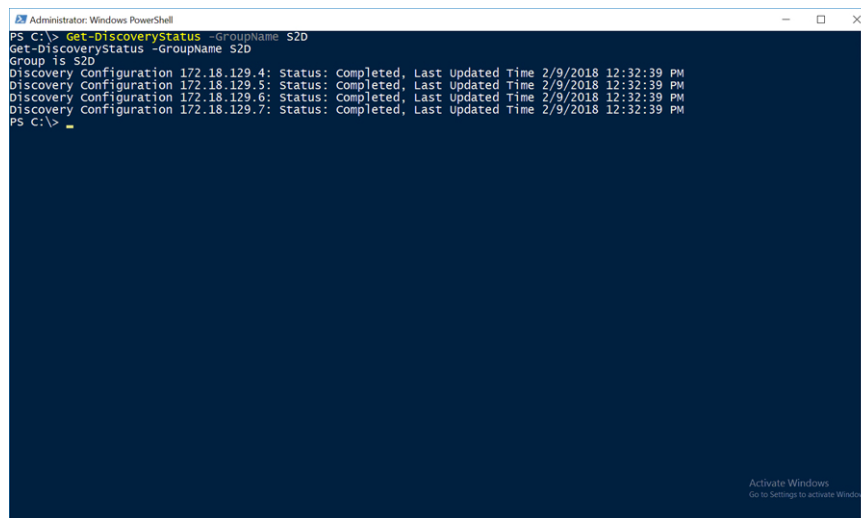
Discovery and inventory of Storage Spaces Direct Ready Node using OME CLI

Unlike adding a range using the OME console, the range added by the CLI will not automatically be submitted for discovery and inventory. To submit the range for discovery and inventory, run the `Set-RunDiscoveryInventory -GroupName "Group Name"`command.

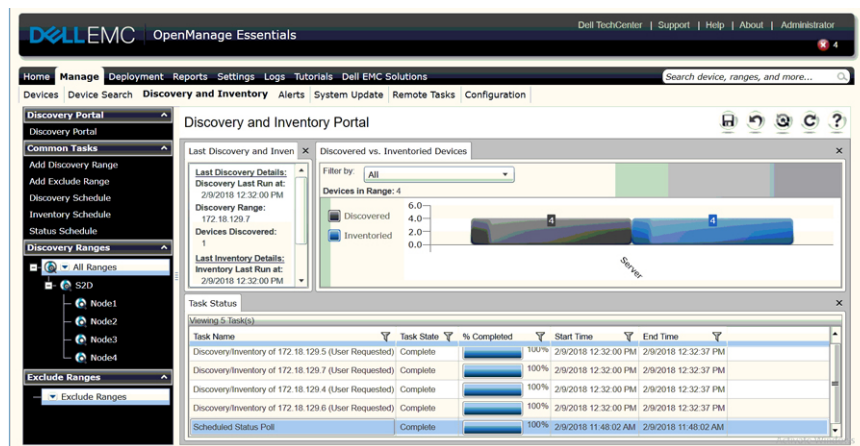**Figure 7. Run discovery inventory**

Figure 8. Discovery status

Figure 9. OME inventory status

Creating custom group and adding discovered devices to the custom group

The following command describes the process to create custom groups and add the discovered devices to a particular group.

```
$Servers = Select-DevicesFromGroup -GroupName RAC

Add-CustomGroup -GroupName "Group Name" -Devices $Servers
```
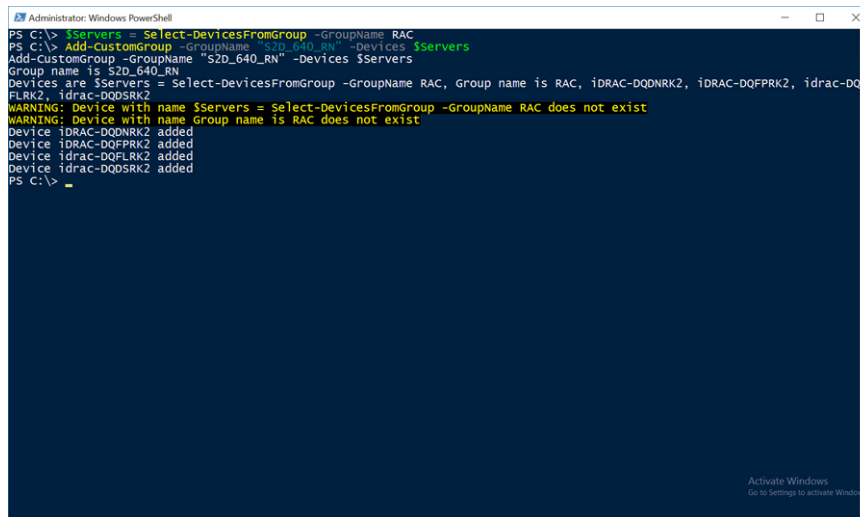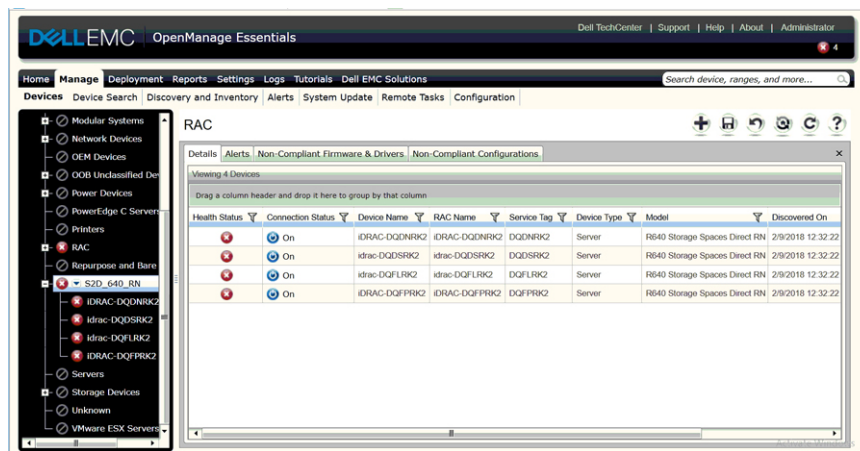

Figure 10. Creating custom group

Figure 11. Custom group created

Obtain firmware catalog for Storage Spaces Direct Ready Nodes

For Storage Spaces Direct Ready Node, the recommended option is to use a storage spaces direct specific catalog for a qualified set of firmware and BIOS. Get the latest catalog from the link below and copy to a location in the OME server.

> ⓘ | NOTE: OME 2.4 supports multiple catalogs. Catalogs can be applied to device groups.

http://en.community.dell.com/techcenter/extras/m/white_papers/20445064.

Perform the following steps to associate a custom catalog to a device group:

1   From OME, click **Manage → Devices**
2   Select the device group for Storage Spaces Direct Ready Nodes. Right click and select **Associate Catalog Baseline.**
     If a catalog baseline is already created, then it is displayed under **List of Catalog Baselines**.
3   If not listed, select the option **Create Catalog Baseline**. Provide a name in the **Baseline Name** field.
4   Select **Use repository manager file**, browse to the location, select Catalog.xml file, and click **Import now**.
5   Click **Finish**.

Storage Spaces Direct Ready Node lockdown

The Dell EMC 14th generation Storage Spaces Direct Ready Node servers offer a new capability called System Lockdown mode. After the initial deployment is complete and the storage spaces direct cluster is functional, Dell EMC recommends locking down the system configuration so that any updates to BIOS/iDRAC configuration settings and firmware updates are blocked. This ensures that the system configuration stays compliant to a validated baseline. System configuration cannot be changed from any of the out of band or in-band interfaces that are supported. This does not prevent from running the routine monitoring and maintenance tasks such as power operations, power budget, profiles, operations such as blinking drive LEDs, and running diagnostics.

Perform the following steps to enforce system lockdown mode by using RACADM command line task:

1   Open the OME console and select **Manage** > **Remote tasks** > **Create Command line task**.
2   In the **Command** field, add `set idrac.lockdown.systemlockdownmode enabled`.
3   Click **Next** and select the servers to enable the system lockdown mode.
4   Click **Next** and enter the iDRAC credentials.

DELLEMC

5   Click **Finish**.

Figure 12. Enabling lockdown

Prerequisites for maintenance

Use the following PowerShell commands to ensure that all the requirements are met before proceeding with the maintenance operation of the Storage Spaces Direct Ready Node in a Storage Spaces Direct cluster. These steps will ensure that all the requirements are met and there are no faults before placing a Storage Spaces Direct Ready Node into Maintenance Mode.

1   Verify if all the nodes in the cluster are available using `Get-clusternode` command.
2   Verify if the all the cluster network are available using `Get-ClusterNetwork`.
3   Verify if the cluster status is healthy using the following commands:
     · `Get-ClusterS2D`
     · `Get-StoragePool`
     · `Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport`
4   Verify if all the physical and virtual disks are healthy using the following commands:
     · `Get-physicaldisk`
     · `Get-virtualdisks`
5   Verify that there are no backend repair jobs running using the `Get-storagejob` command.

Preparing Storage Spaces Direct Ready Node in maintenance mode

After the pre-requisites are met and before performing the platform updates, the Storage Spaces Direct Ready Node must be placed in Maintenance Mode (Pause and Drain). You can move roles or virtual machines and gracefully flush and commit data in storage spaces direct ready node.

1   Run the following command to put the node in maintenance mode (pause and drain):

```
Suspend-ClusterNode -name "Hostname" -Drain.
```

2   Turn off the System Lockdown mode.

System lock down prevents the node for any updates to BIOS or DRAC configuration settings and firmware updates, it is necessary to turn off the system lockdown mode for system updates. For more information on System Lockdown, refer Storage Spaces Direct Ready Node lockdown.

Using Virtual Machine Manager

Perform the following steps to place the Storage Spaces Direct Ready Node in maintenance mode using VMM

1   In the VMM console, click **Fabric** > **Fabric Resources** > **Servers** > **All Hosts**.

2   Select the host to place in maintenance mode. In the **Host** group, click **Start Maintenance Mode**.

3   Select **Move all virtual machines to other hosts in the cluster**.

4   To verify if the host is in maintenance mode, check the status in **Fabric** > **Hosts**.

For system updates, DellEMC recommends to take one node at a time for Maintenance Mode when there is relatively low system activity.

Update Storage Spaces Direct Ready Node

Use the OME console to update the Storage Spaces Direct Ready Node by performing the following steps:

1   Click **Manage** > **Devices**. Verify that the Ready Node is discovered and classified under RAC device group. The discovered iDRAC is present either under the compliant or non-compliant systems section in the compliance pie-chart.

Enter the result of your step here (optional).

2   Click **System Update** > **Summary Tab** > **Advanced Settings**.

3   Set preferred update mode to **Out-of-Band (iDRAC)**.

4   Click **OK** to save the settings and close the Advanced Settings window. The figure below shows a screenshot of the advanced settings window:

Figure 13. Advanced settings

5   If the discovered iDRAC is non-compliant, it is listed under **Non-Compliant** tab. Select the iDRAC that is non-compliant and the package to be updated on the system and click **Apply Selected Updates**.

6   When the User Preferred Delivery Mode is set to iDRAC, the Update Method will display Out-of-Band for all the available components. The System Update Task window is displayed.

7   Provide a name for the task.

8   Select the **Run now** option.

9   Enter the user name and password of the iDRAC.

10 Click **Finish** to create system update task.

The system update task is created with the mentioned name and appended with "- iDRAC" to the task name. This indicates that the preferred mode of delivery is Out-of-Band. After all the selected components (DUPs) are successfully applied on the selected managed system, the task status is set to Complete. About 20 minutes after the task is complete, an auto inventory task is initiated to gather the updated inventory information.

Enforce device compliance on a Storage Spaces Direct Ready Node

OME can be used to track BIOS and other system configuration. A system configuration baseline can be first established which saves all the system configuration. Whenever the configuration needs to be updated, the baseline is updated and OME can be used to apply this update to all the nodes in the storage spaces direct cluster.

The following subsections describe each task in detail:

Configure deployment file share

Before creating or deploying a configuration template from a device, a deployment file share on the server running OME must be configured. The deployment file share temporarily stores the configuration file that is used to retrieve and apply the configuration settings on a target server or chassis. To configure the deployment file share:

1   Click **Deployment**. In the **Common Tasks** pane, click **File Share Settings**.

    The **File Share Settings** window is displayed.

2   Type the domain user name and password of the server running OME in the appropriate field.

3   Click **Apply**. If the file share is configured successfully, the File Share Status is displayed.

Create device configuration baseline

To create a baseline, perform the following steps:

1   Click **Manage** > **Configuration**.

2   In the **Common Tasks** pane, click **Create Baseline**.

    The Create Baseline Wizard is displayed.

3   In the **Name** field, type a name for the baseline.

4   Select **Create from File** to create a baseline by importing the XML template.

Associate devices to baseline

The **Associate Devices to a Baseline** task designates the baseline to be used for verifying the compliance status of target devices. To associate devices to a baseline, perform the following steps:

1   Click **Manage** > **Configuration**.

2   In the **Common Tasks** pane, click **Associate Devices to a Baseline**.

3   The Associate Devices to a Baseline wizard is displayed.

4   On the Select Baseline page, click **Server Baseline**.

5   Select a device configuration baseline from the list, and click **Next**.

6   On the **Select Devices** page, select the target devices from the **All Applicable Devices** tree, and then click **Finish**.

Remediate non-compliant device

Perform the following steps to remediate the Storage Spaces Direct Ready Nodes that are not compliant with the associated device configuration baseline:

1   Click **Manage** > **Configuration** > **Make Device(s) Complaint**.

    The Name page is displayed.

2   Enter the Name for the remediation task, and click **Next**.

3   On the Select Devices page, the list of non-compliant servers with the corresponding non-compliant attributes is displayed. Select the non-compliant device from the list, and click **Next**. Ensure that the selected non-compliant device is already in maintenance mode.

4   On the **Options** page, select **Automatic Server Reboot** and click **Next**.

5   On the **Set Schedule** page:

    a   Select **Run now**.

    b   On the **Execution Credentials** page, type the credentials of the selected device, and click **Next**.

    c   On the **Summary** page, review the information that has been provided, and then click **Finish**.

Verify device compliance

After the Device Configuration Baseline has been applied to a Storage Spaces Direct Ready Node, its compliance should be verified. To view the compliance status of devices on the associated configuration baseline:

1 Click **Manage** > **Configurations** > **Device Compliance Portal**.

The **Device Compliance** graph and grid display the compliance status of the devices.

2 Click the **Device Compliance** graph to verify the compliance status of the device.

Verify compliance

After the server is rebooted and re-inventoried by OME, ensure that the updated Storage Spaces Direct Ready Node is compliant.

1 To view the compliant servers:

    a Click **Manage** > **System Update**.

    b In the**System Update** page, select the **Compliant Systems** tab.

2 To view the non-compliant servers:

    a Click **Manage** > **System Update**.

    b In the **System Update** page, select the **Non-Compliant Systems** tab.

       Systems with drivers and firmware versions that are different from the catalog are displayed.

Exit Storage Spaces Direct Ready Node from maintenance mode

After updating the Storage Spaces Direct Ready Node, exit the node from maintenance mode using `Resume-ClusterNode -Name "Hostname" -Failback Immediate` PowerShell command.
The storage spaces direct operation to rebuild and rebalance the data to ensure load balancing is initiated.

Exit using VMM

Perform the following steps to exit from the maintenance mode using VMM.

1 In the VMM console, click **Fabric** > **Fabric Resources** > **Servers** > **All Hosts**.

2 Select the host to bring a host out of maintenance mode, click **Stop Maintenance Mode**.

Exit using iDRAC out of band

Dell EMC Storage Spaces Direct Ready Node offers device firmware updates remotely through the integrated Dell Remote Access Controller (iDRAC). For Storage Spaces Direct, the recommended option is to use a storage spaces direct specific custom catalog for a qualified set of firmware and BIOS. Download the catalog file from with the necessary firmware as mentioned in Storage Spaces Direct support matrix file to a network share before proceeding with the update procedure.

1 Login to the iDRAC web interface.

2 Click **Maintenance** > **System Upadate**.

The **Firmware Update** page is displayed.

3 On the **Update** tab, select **Network Share** as the File Location.
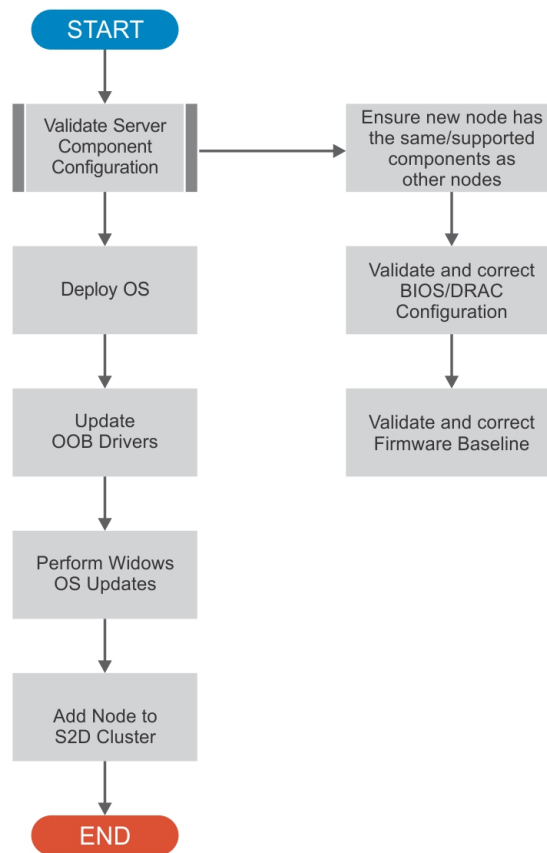
4 Provide the details of the network share as shown here:

Figure 14. Check for updates

5   Click **Check for updates**.

6   A list of available updates are displayed as shown here.

Figure 15. Select updates

7   Select the updates and click on **Install and Reboot** to install and reboot the system.

Expanding Ready Node cluster

Expanding cluster compute or storage capacity is one of the tasks performed during cluster operations. This section provides instructions on these tasks.

Figure 16. Expanding Ready Node cluster

Adding server nodes

Before adding new server nodes into a HCl cluster, complete the following requirements:

1   Verify that the drives are of the same configuration as the current nodes on the cluster. It should have the same number of drives, and also the different tiers (SSD, NVM, HDD) should be the same size drive as the node in use.

2   Apply the BIOS configuration to the node and configure iDRAC.

    To configure the node, the OS deployment and node network configuration can be done by using the steps mentioned in the section Hyper-converged Infrastructure Deployment in the Dell EMC Microsoft Storage Spaces Direct Ready Node Deployment Guide.

    The PowerShell commands in the following sections should not be executed again as the cluster is already created, Storage Spaces direct already enabled, and the management network already excluded:

    · Configuring Storage Spaces Direct
    · Create Host Cluster
    · Remove Host Management Network from Live Migration

3   Verify that the nodes are compliant with the firmware baseline.

4   Update the hardware timeout configuration for the Spaces port.

5   After the node is configured, update Windows to bring the node to the same level as the cluster.

> ⓘ | **NOTE: Before adding the node to the existing cluster, execute the *Get-PhysicalDisk* command on the new node and verify the output to ensure that all disks are in healthy state and there are equal number of disks as the other cluster node.**

Adding server nodes manually

To manually add server nodes to the cluster, see https://technet.microsoft.com/EN-US/WINDOWS-SERVER-DOCS/STORAGE/STORAGE-SPACES/ADD-NODES.

ⓘ | NOTE: **The procedure is applicable only if the cluster and Storage Spaces Direct configuration is done manually.**

Adding server nodes using VMM

If the initial cluster creation and Storage Spaces Direct configuration was done using System Center VMM, the following commands can be run on the VMM server to add the new node into the existing cluster:

```
$logicalSwitchName = 'S2DSwitch'
$hostGroupName = 'S2DHosts'
$runAsAccountName = 'Administrator'
$uplinkPortProfileName = 'S2D_UPP'
$clusterName = 'S2D4NCluster.hci.lab'

#Get RunAs Account
$runAsAccount = Get-SCRunAsAccount -Name $runAsAccountName

# Get Host Group
$hostGroup = Get-SCVMHostGroup -Name $hostGroupName
Add-SCVMHost -ComputerName $nodeName -VMHostGroup $hostGroup -Credential $runAsAccount

#Job Group GUID
$jobGroup = (New-Guid).Guid

# Get Host 'S2D4Node004.hci.lab'
$vmHost = Get-SCVMHost | Where-Object { $_.Name -eq $NodeName }

# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter'
$networkAdapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName
-eq 'SLOT 1 Port 1' }
$uplinkPortProfileSet = Get-SCUplinkPortProfileSet -Name $uplinkPortProfileName
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $networkAdapter -UplinkPortProfileSet
$uplinkPortProfileSet -JobGroup $jobGroup

# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter #2'
$networkAdapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName
-eq 'SLOT 1 Port 2' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $networkAdapter -UplinkPortProfileSet
$uplinkPortProfileSet -JobGroup $jobGroup
$networkAdapter = @()
$networkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 1' }
$networkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 2' }
$logicalSwitch = Get-SCLogicalSwitch -Name S2dSwitch

#Management
$vmNetwork = Get-SCVMNetwork -Name 'Management'
$vmSubnet = Get-SCVMSubnet -Name 'Management_0'
New-SCVirtualNetwork -VMHost $vmHost -VMHostNetworkAdapters $networkAdapter -LogicalSwitch
$logicalSwitch -JobGroup $jobGroup -CreateManagementAdapter -ManagementAdapterName "Management"
-ManagementAdapterVMNetwork $vmNetwork -ManagementAdapterVMSubnet $vmSubnet

#Storage1
$vmNetwork = Get-SCVMNetwork -Name 'Storage1'
$vmSubnet = Get-SCVMSubnet -Name 'Storage1_0'
$ipV4Pool = Get-SCStaticIPAddressPool -Name 'Storage1-IPpool'
New-SCVirtualNetworkAdapter -VMHost $vmHost -Name "Storage1" -VMNetwork $vmNetwork -
```

```
LogicalSwitch $logicalSwitch -JobGroup $jobGroup -VMSubnet $vmSubnet -IPv4AddressType "Static" -
IPv4AddressPool $ipV4Pool -MACAddressType "Static" -MACAddress "00:00:00:00:00:00"

#Storage2
$vmNetwork = Get-SCVMNetwork -Name 'Storage2'
$vmSubnet = Get-SCVMSubnet -Name 'Storage2_0'
$ipV4Pool = Get-SCStaticIPAddressPool -Name 'Storage2-ippool'
New-SCVirtualNetworkAdapter -VMHost $vmHost -Name "Storage2" -VMNetwork $vmNetwork -
LogicalSwitch $logicalSwitch -JobGroup $jobGroup -VMSubnet $vmSubnet -IPv4AddressType "Static" -
IPv4AddressPool $ipV4Pool -MACAddressType "Static" -MACAddress "00:00:00:00:00:00"

#Set the host properties
Set-SCVMHost -VMHost $vmHost -JobGroup $jobGroup -RunAsynchronously

#Add to cluster
# Get Host Cluster 'S2D4NCluster.hci.lab'
$VMHostCluster = Get-SCVMHostCluster -Name $clusterName

# Get Host 'S2D4Node004.hci.lab'
$VMHosts = @()
$VMHosts += (Get-SCVMHost).Where({ $_.FQDN -eq $nodeName })

Install-SCVMHostCluster -VMHostCluster $VMHostCluster -VMHost $VMHosts -Credential
$runAsAccount -RunAsynchronously
```

Adding disks

When new disks are added to extend the overall storage capacity per node, the storage spaces direct cluster starts claiming the physical disks into an existing storage pool.

After the drives are added, they will be shown as available for pooling (CanPool set to True) in the output of the `Get-PhysicalDisk` command.

Within a few minutes, the newly added disks will be claimed into the existing pool and storage spaces direct will start a rebalance job. Run the following command to verify that the new disks are a part of the existing pool:

```
PS C:\> Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport
CPUUsageAverage            :    2.66 %
CapacityPhysicalPooledAvailable :   8.01 TB
CapacityPhysicalPooledTotal    :  69.86 TB
CapacityPhysicalTotal       :  69.86 TB
CapacityPhysicalUnpooled     :      0 B
CapacityVolumesAvailable     :  15.09 TB
CapacityVolumesTotal        :  16.88 TB
IOLatencyAverage          : 908.13 us
IOLatencyRead            :      0 ns
IOLatencyWrite           : 908.13 us
IOPSRead              :      0 /S
IOPSTotal             :      1 /S
IOPSWrite             :      1 /S
IOThroughputRead          :      0 B/S
IOThroughputTotal         :  11.98 KB/S
IOThroughputWrite         :  11.98 KB/S
MemoryAvailable          : 472.87 GB
MemoryTotal            :    768 GB
```

After all available disks are claimed into the storage pool, the **CapacityPhysicalUnpooled** should be 0 B.

The storage rebalance job may take a few minutes, and the process can be monitored by using the `Get-StorageJob` cmdlet.

Extending volumes

Volumes created in Spaces Direct storage pools can be resized using the Resize-VirtualDisk cmdlet. For information and for the commands used to perform this task, see https://technet.microsoft.com/en-us/windows-server-docs/storage/storage-spaces/resize-volumes.

Performing Ready Node recovery

In case of a cluster node failure, node OS recovery should be done in a systematic manner to ensure that the node is brought up with configuration consistent with other cluster nodes. The following sections provide details about OS recovery and post OS recovery configuration that is needed to bring the node into an existing Storage Spaces Direct cluster.

ⓘ **NOTE: To perform node recovery, ensure that the OS is re-installed. See the instructions in the OS RAID configuration section in the Dell EMC Microsoft Storage Spaces Direct Ready Node Deployment Guide to create an OS RAID disk.**

OS RAID configuration

The Dell EMC PowerEdge servers offer Boot Optimized Storage Solution (BOSS) S-1 controller as an efficient and economical way to segregate operating system and data on the internal storage of the server. The BOSS S-1 solution in the 14th generation of PowerEdge servers uses one or two BOSS M.2 SATA devices to provide RAID 1 capability for the operating system drive.

ⓘ **NOTE: All Dell EMC Microsoft Storage Spaces Direct Ready Nodes are configured with hardware RAID 1 for the OS drives on BOSS M.2 SATA SSD devices. The steps in this section are required only when recovering a failed cluster node. Before creating a new RAID, the existing or failed RAID must be deleted.**

The following steps describe the process to create OS volumes.

1. Login to iDRAC web console.
2. Navigate to **Storage** > **Controllers** and ensure that the BOSS S.1 controller is visible.
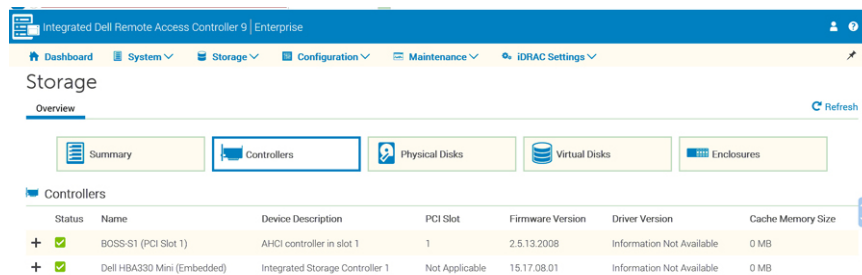
Figure 17. View controllers

3. Navigate to **Configuration** > **Storage Configuration**. Expand **Virtual Disk Configuration** and click **Create Virtual Disk**.
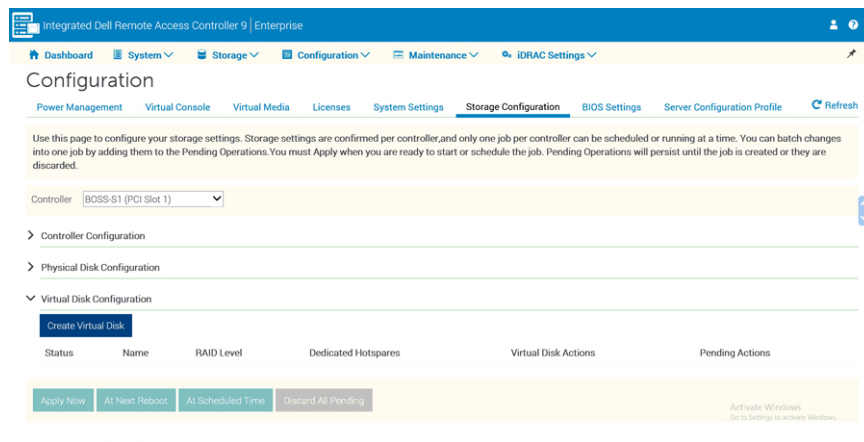
Figure 18. Creating virtual disk

4   Provide a VD name and select BOSS M.2 devices in the physical disks.

Figure 19. Provide virtual disk name

Figure 20. Select physical disks

5   Click **Add Pending Operations**.

6   Navigate to **Configuration** > **Storage Configuration** > **Virtual Disk Configuration**. Select the VD and select **Initialize: Fast in Virtual Disk Actions**.

Figure 21. Initialize configuration

7   Reboot the server. The VD creation process may take several minutes to complete.

8   After the initialization is completed successfully, the VD health status is displayed.
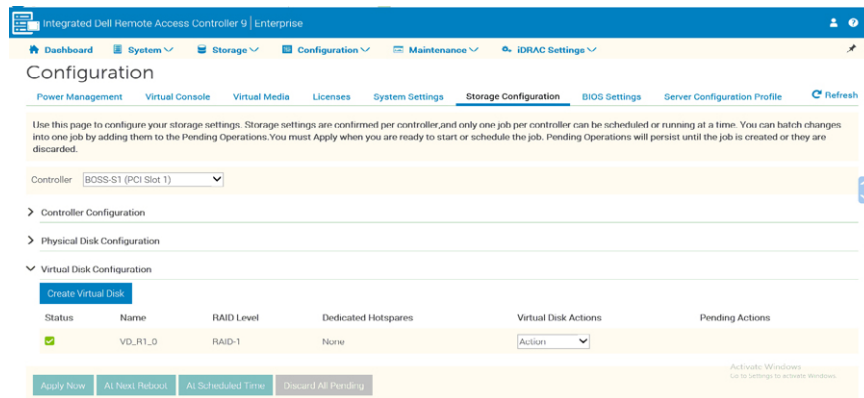
Figure 22. View status

Operating system recovery

This section provides an overview of steps involved in OS recovery on the Dell EMC Microsoft Storage Spaces Direct Ready Nodes.

ⓘ NOTE: **Ensure that the RAID 1 VD created on the BOSS M.2 drives is re-initialized.**

ⓘ NOTE: **Do not reinitialize or clear the data on the disks that were a part of storage Spaces Direct storage pool. This will help in reducing repair times when the node is added back to the same cluster after recovery.**

ⓘ NOTE: **Before performing OS recovery on the failed node, ensure that the node is evicted from the existing cluster and VMM host groups, if the node was managed using VMM. Ensure that the computer account and DNS entries related to the node are deleted from Active Directory and DNS.**

Manual Operating system recovery

For manually deployed nodes, the OS recovery on the node can be done by using any of the methods that were used for OS deployment.

After the node is added to the existing Storage Spaces Direct cluster, the optional configuration described in section Onboarding Ready Nodes in SCOM can be performed to enable monitoring and management of the recovered Storage Spaces Direct node by using System Center OM.

Factory Operating system recovery

For the factory installed OEM license of the operating system, it is recommended to use the OS recovery media shipped with the PowerEdge server. Using this media for OS recovery ensures that the OS stays activated after the recovery. Using any other OS media will trigger the need for activation post OS deployment. OS deployment using the recovery media is same as either retail or other OS media based installation.

After the OS deployment using the recovery media is complete, perform the following steps to bring the node into an existing Storage Spaces Direct cluster:

1 Update OOB Drivers
2 Configure host networking
3 Change host name
4 AD Domain Join
5 QoS Policy (for RoCE for RDMA only)

6   RDMA configuration

7   Configure firewall

8   Day 0 OS updates

9   Add Server Nodes to the cluster

For instructions on updating OOB drivers, configuring host networking, changing host name, AD domain join, QoS policy, RDMA configuration, configuring firewall, see the Dell EMC Storage Spaces Direct Ready Node deployment guide.

Some instructions for updating OOB drivers, configuring host networking, changing host name, AD domain join, QoS policy, RDMA configuration, configuring firewall may be different for the 2-Node ROBO configuration. For more information see the Dell EMC Storage Spaces Direct Ready Node deployment guide for ROBO for deployment instructions specific to ROBO cluster.

DELLEMC