Dell EMC Solutions for Microsoft Azure Stack HCI

Operations Guide for managing and monitoring solution infrastructure life cycle

DELL EMC

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.

2019 - 04

Rev. A02

Contents

1

Audience and scope

This Operations guide focuses on operational aspects of a hyper-converged infrastructure solution on Azure Stack HCI with Hyper-V and Storage Spaces Direct. This operations guide includes an overview of the Solutions for Microsoft Azure Stack HCI, 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 Solutions for Windows Server 2016 and Windows Server 2019 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 and Windows Server 2019 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 and Windows Server 2019 with the latest applicable updates. These instructions are not validated with Windows Server version 1709. Dell EMC Ready Solutions do not support the Windows Server Semi-Annual Channel release. Dell EMC recommends updating the host OS with latest cumulative updates from Microsoft before starting the cluster creation and configuration tasks.

Topics:

Assumptions

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

- Dell EMC Solutions for Azure Stack HCI and deploying and configuring BIOS and iDRAC settings on these ready nodes.
- Deploying and configuring Windows Server 2016 and Windows Server 2019 Hyper-V infrastructure.

Known issues

Before starting the cluster deployment, ensure that you review the known issues and workarounds. For a list of known issues, see https://www.dell.com/support/article/sln313305.

2

Dell EMC Solutions for Azure Stack HCI overview

Dell EMC Solutions for Azure Stack HCI encompasses various configurations of R740xd, R740xd2, and R640 Storage Spaces Direct Ready Node and PowerEdge R440 servers to power the primary compute cluster deployed as a hyper-converged infrastructure. This hyper-converged infrastructure built by using these Ready Nodes uses a flexible solution architecture rather than a fixed component design. The following illustration 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.

The Dell EMC Solutions for Azure Stack HCI based on Dell EMC Storage Spaces Direct Ready Nodes is available in both hybrid and all-flash configurations. For more information about available configurations, see Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview.

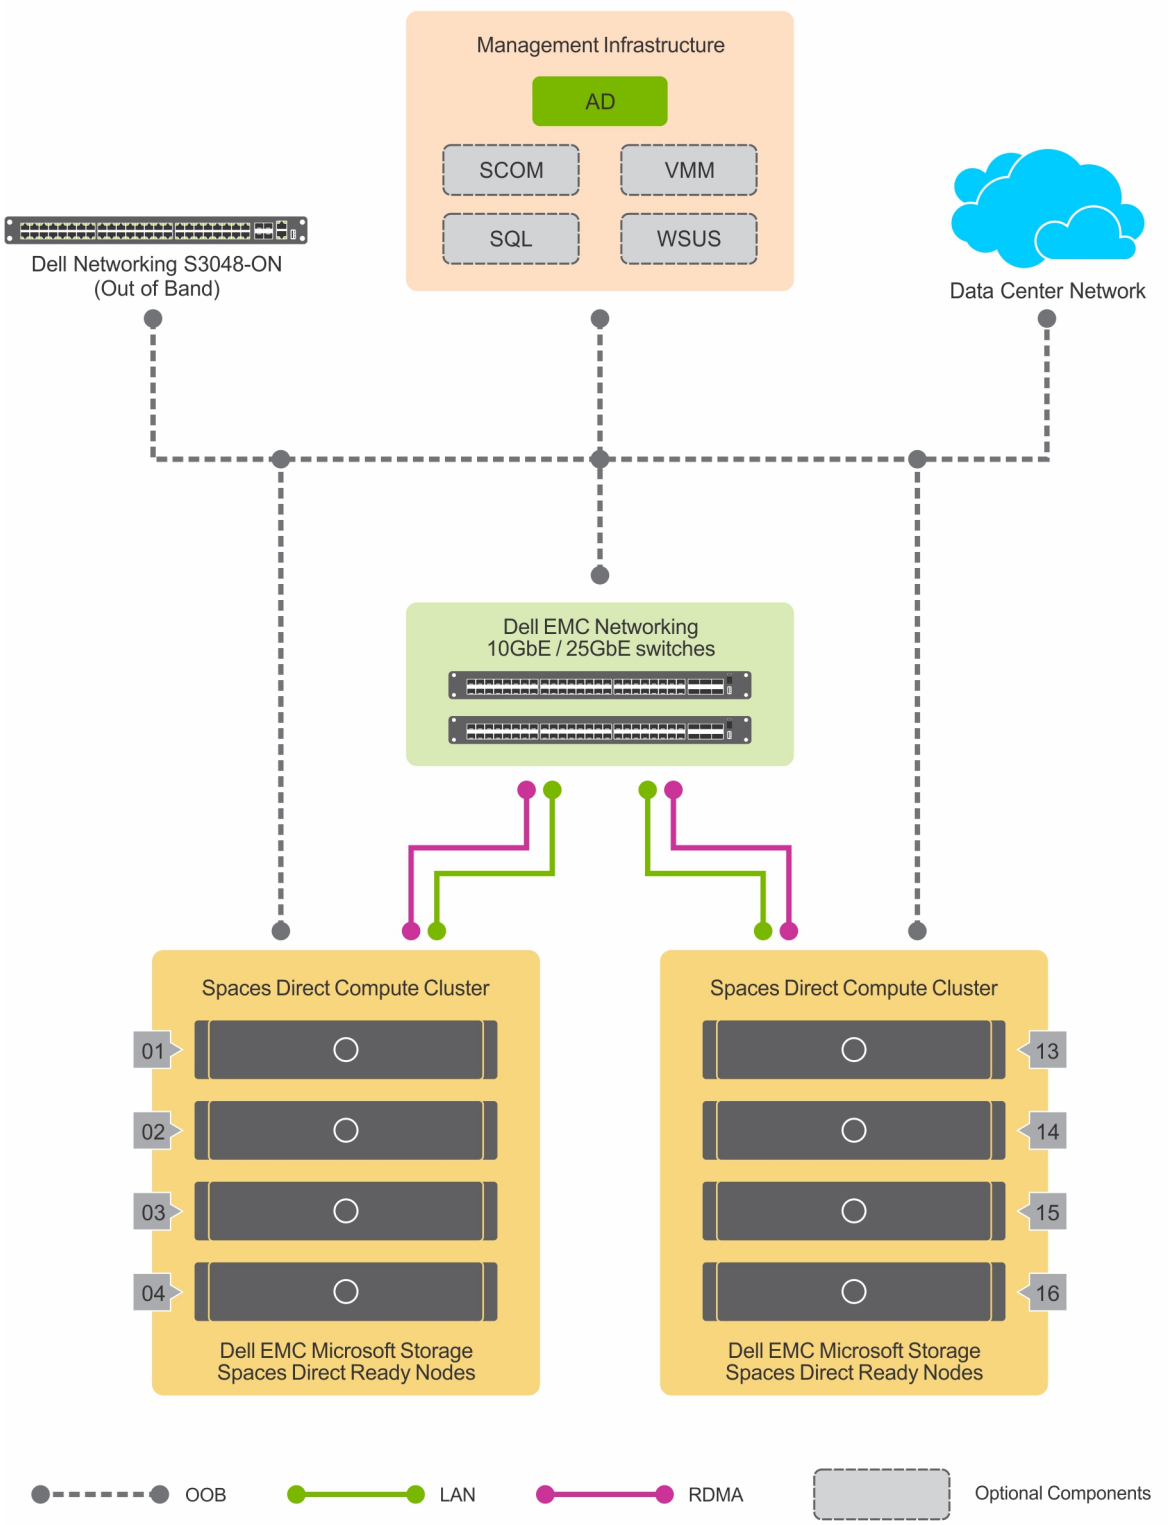

Figure 1. Hyper-converged virtualized solution using Dell EMC Ready Nodes

Deployment guidance

The deployment guidance and instructions for configuring a cluster using the Dell EMC Solutions for Azure Stack HCI is available at dell.com/wssdmanuals. This operations guidance is applicable only to cluster infrastructure that is built using the instructions provided in the Ready Node deployment guidance.

3

Day 0 checklist

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

Topics:

Windows activation

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

NOTE: Windows activation is not required if the operating system is factory installed.

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

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

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

Storage Spaces Direct management

The following sections provide instructions on day 0 Storage Spaces Direct management.

Create 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.

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 the 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/windows-server/storage/storage-spaces/plan-volumes.

Managing and monitoring Dell EMC Solutions for Azure Stack HCI cluster using Windows Admin Center

Windows Admin Center (WAC) is the new browser-based management tool developed by Microsoft to monitor and manage Windows servers, failover clusters, and hyper-converged clusters.

The Dell EMC Microsoft Storage Spaces Direct Ready Nodes offer Software-Defined Storage (WSSD) building blocks to create highly available and highly scalable hyper-converged Infrastructure (HCI). These ready nodes are preconfigured with certified components and validated as a S2D solution that includes Ready Nodes with Dell EMC networking switches as building blocks with simplified ordering and reducing deployment risks. Dell EMC offers different configuration options within these building blocks to meet different capacity and performance points. The HCI clusters built on these building blocks can be seamlessly monitored and managed using the WAC.

Install Windows Admin Center

You can download Windows Admin Center (WAC) from Microsoft download center and install on Windows 10, Windows Server 2016, Windows Server 2019, or Windows Server version 1709. You can install WAC directly on a managed node to manage itself. You can also install WAC on other nodes in the infrastructure or on a separate management station to manage the HCI nodes remotely. It is possible to implement high-availability (HA) for WAC by using failover clustering. When WAC is deployed on nodes in a failover cluster, it acts as a active-passive cluster providing highly available WAC instance.

WAC installer wizard performs the configuration required for WAC functionality. This includes creating a self-signed certificate and configuring trusted hosts configuration for remote node access. Optionally, you can supply the certificate thumb print that is already present in the target node local certificate store. By default, WAC listens on port 443 which can be changed during the installation process.

NOTE: The automatically generated self-signed certificate expires in 60 days. Ensure that you use a certificate authority (CA) provided SSL certificate if you intend to use WAC in production environment.

For complete guidance on installing WAC on Windows Server 2016 and Windows Server 2019 with desktop experience or Server Core, see https://docs.microsoft.com/windows-server/manage/windows-admin-center/deploy/install.

NOTE: This section assumes that you have deployed the Dell EMC Microsoft Storage Spaces Direct Ready Node Cluster using the deployment guidance available at: https://dell.com/wssdmanuals.

Once the installation is complete, you can access WAC at https://managaementstationname:<PortNumber>

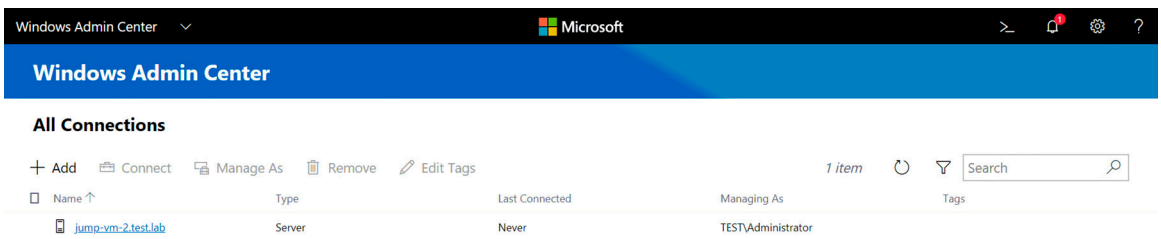

Figure 2. WAC Initial start screen

Add HCI cluster connection

The hyper-converged cluster based on Dell EMC Solutions for Azure Stack HCI can be added as a connection for monitoring and management purpose.

To add HCI cluster connection in WAC:

1 Go to **Windows Admin Center** > **Hyper-Converged Cluster Manager**.

2 Click **Add**.

The **Add Hyper-Converged Custer Connection** window is displayed.

3 Enter the cluster FQDN and select the **Also add servers in the cluster** check box.

The WAC discovers the cluster and the nodes that are part of the cluster.

4 Click **Add**.

The cluster is added to the connection list and WAC is configured to monitor and manage the HCI cluster.

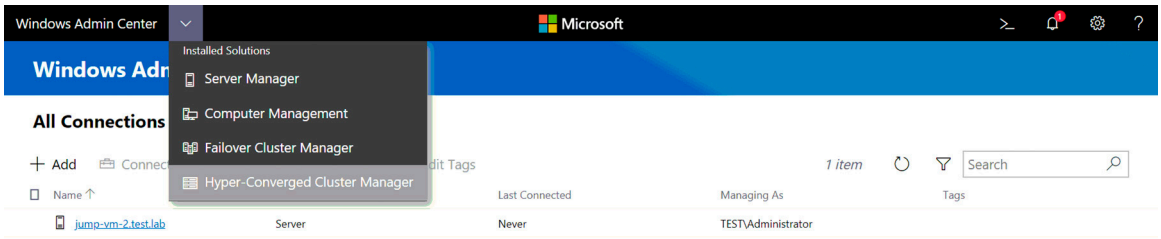

Figure 3. HCI cluster navigation

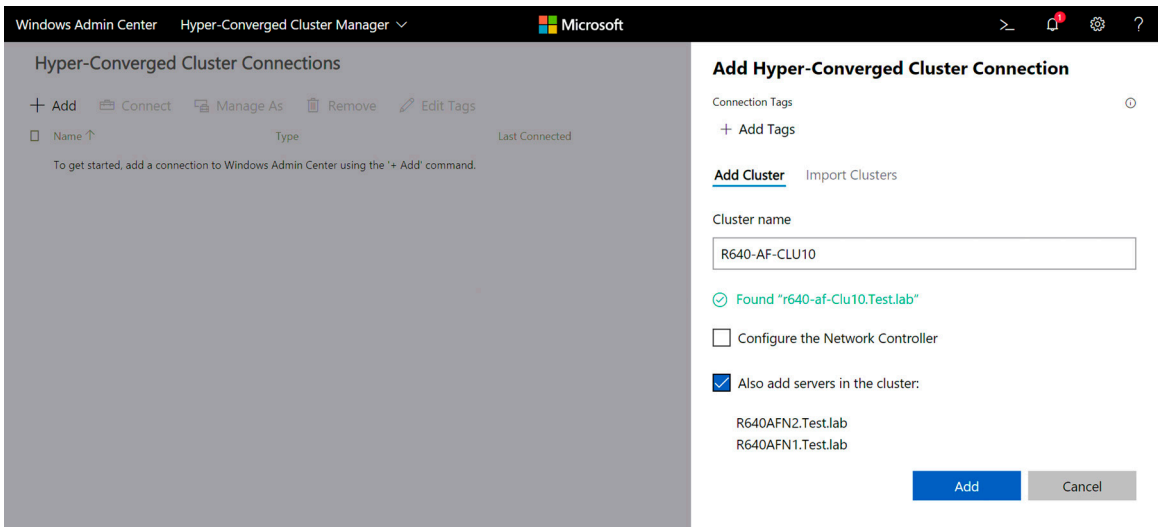

Figure 4. Adding HCI Cluster

Accessing HCI cluster

To view the dashboard for the HCI cluster that you have added to WAC, click the name of the cluster shown in the hyper-converged Cluster Connections window.

This dashboard provides the real-time performance view from the HCI cluster. This view includes total IOPS, average latency values, throughput achieved, average CPU usage, memory usage, and storage usage from all cluster nodes. It also provides a summarized view of the Ready Nodes with drives, volumes, and virtual machines health.

You can drill down into any alerts by clicking on the alerts tile in the dashboard.

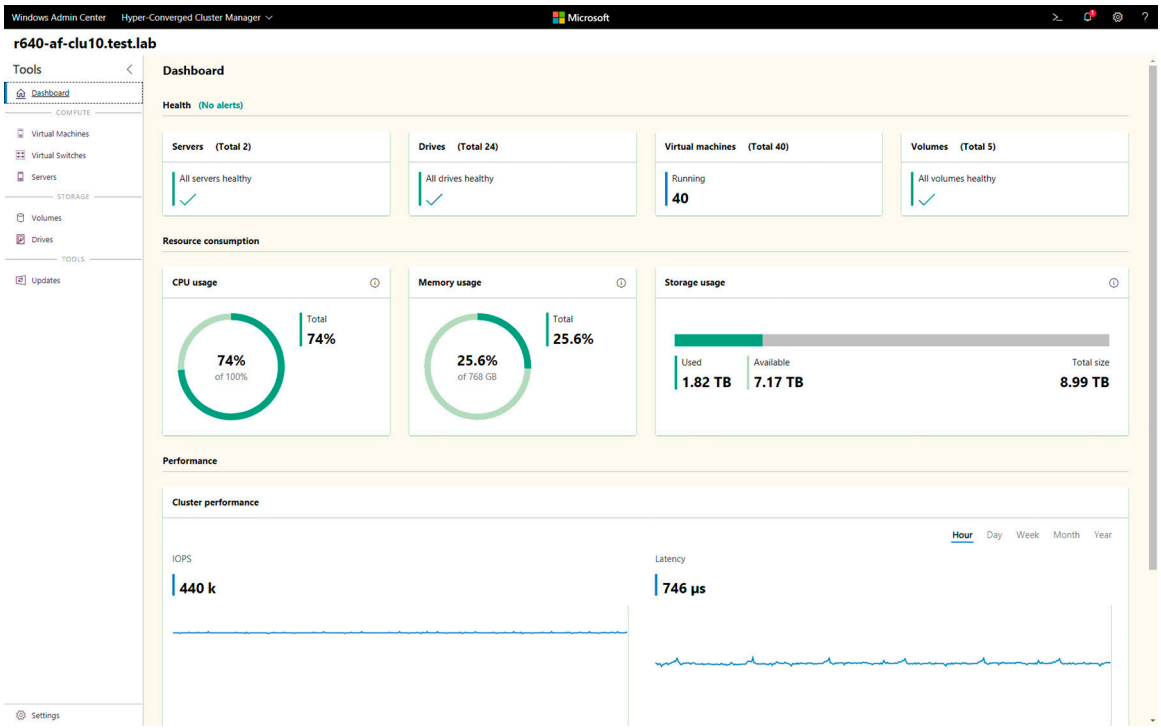

Figure 5. HCI dashboard in WAC

View servers details

To view the server details, click the tools pane and go to **Servers** > **Inventory**.

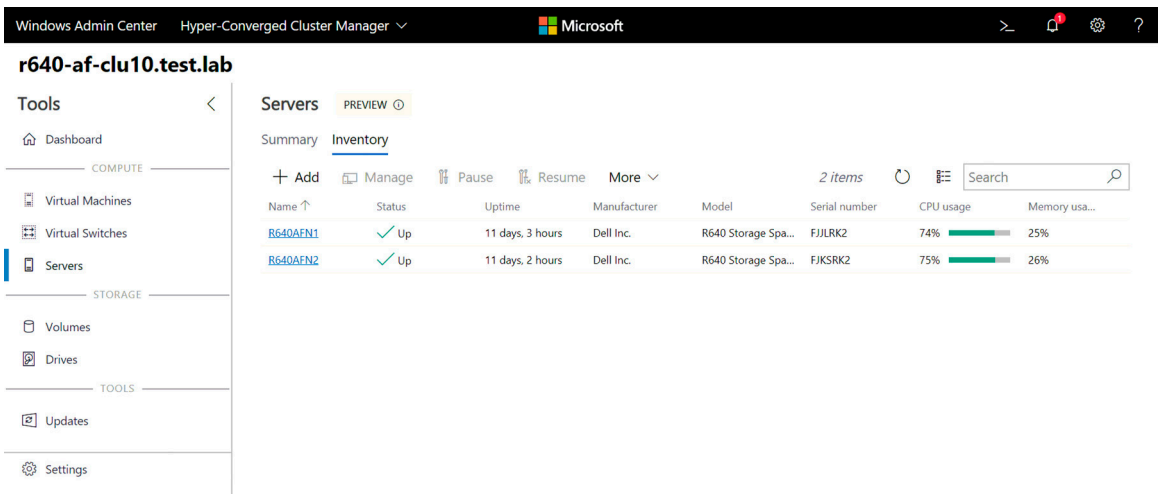

Figure 6. Servers Inventory

NOTE: The performance numbers shown in the illustration are for a 2-node Dell EMC Microsoft Storage Spaces Direct Ready Node cluster with all-flash drive configuration.

View drives details

To view the total number of drives of the cluster, health status of the drives, used, available, and reserve storage of the cluster, go to **Drives** > **Summary** from the left pane.

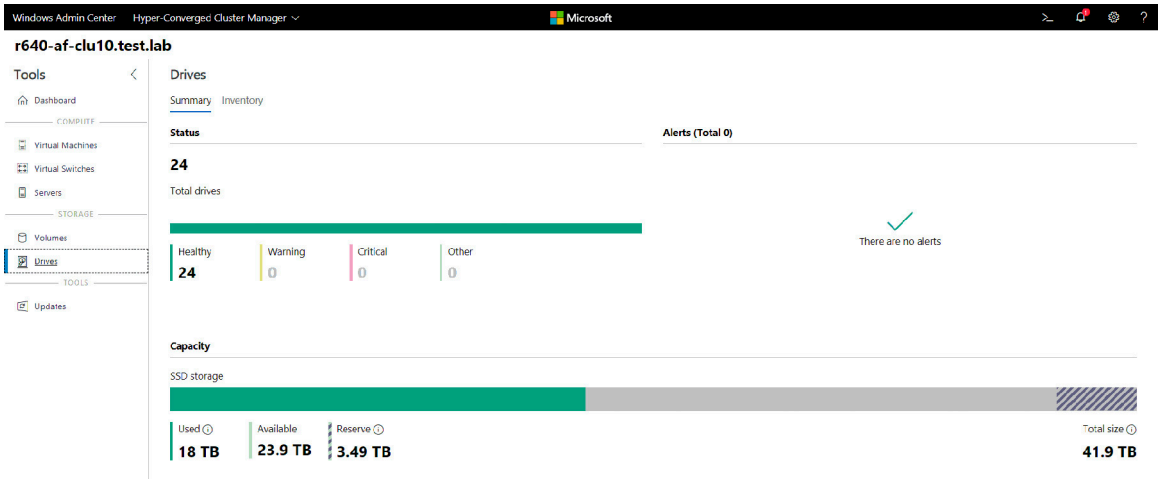

Figure 7. Drives Summary in WAC

To view the drives inventory from the cluster nodes, go to **Drives > Inventory** from the left pane.

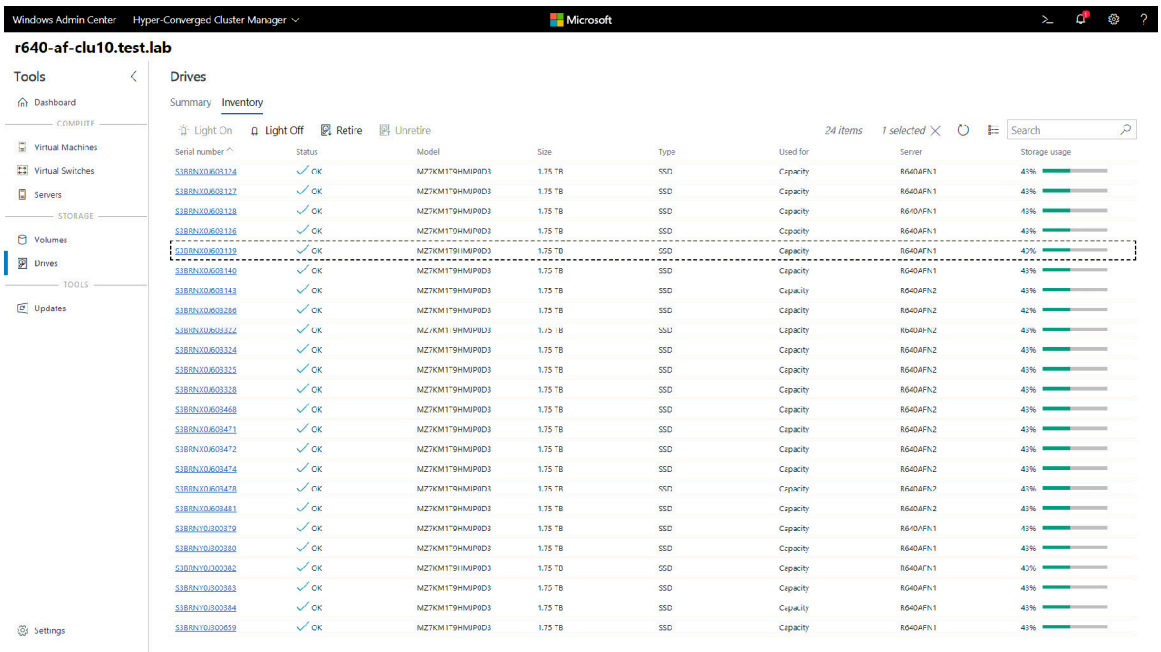

Figure 8. Drives inventory in HCI cluster

The HCI cluster is built on 2 x Dell EMC Solutions each with 12 x 1.92-TB SSD drives.

By clicking the serial number of the drive, you can view the drive information which includes health status, slot location, size, type, firmware version, IOPS, used or available capacity, and storage pool of the drive.

Also, you can set the drive options as **Light On** or **Light Off**, or **Retire** or **Unretire** from the storage pool by selecting the appropriate option from the dashboard.

Manage and monitor volumes

You can manage and monitor the Storage Spaces Direct volumes using WAC.

The following features are supported in WAC:

- Create volume
- Browse volume
- Expand volume
- Delete volume
- Make volume offline or online

To access the volumes created on HCI cluster, click the HCI cluster and go to **Volumes**. On the right pane, the **Summary** and **Inventory** tabs are displayed.

The summary tab provides the information of number of volumes that are created in the cluster, health status of the volumes, alerts, total IOPS, latency, and throughput information of the available volumes.

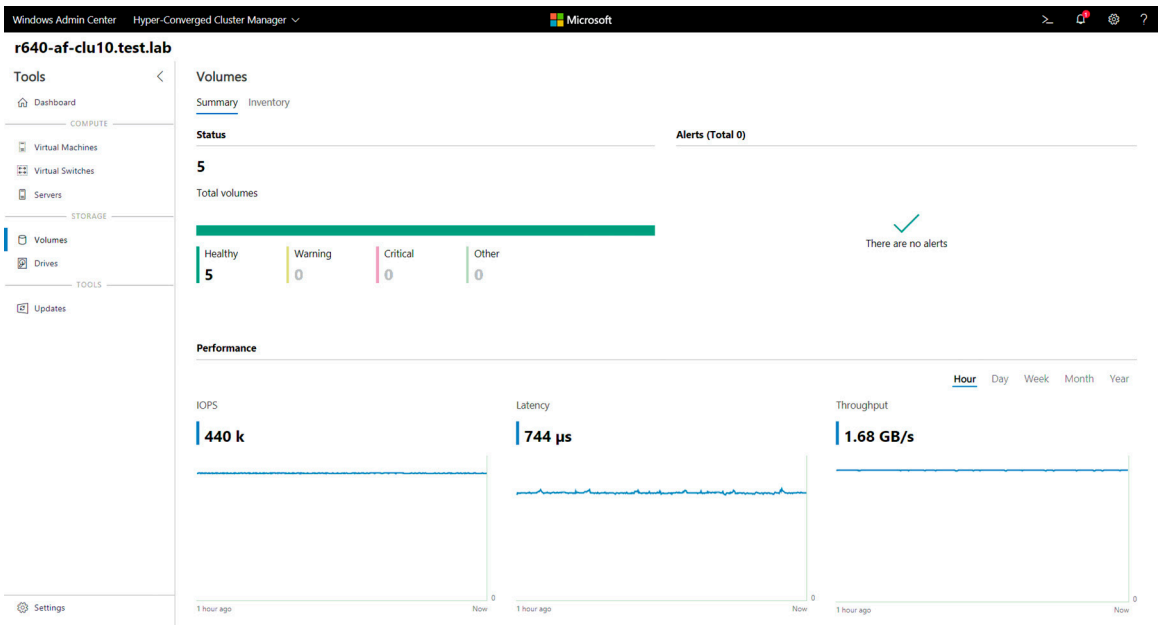

Figure 9. Volumes Summary

The Inventory tab provides the volumes inventory from the HCI cluster nodes. The volumes can be managed and monitored.

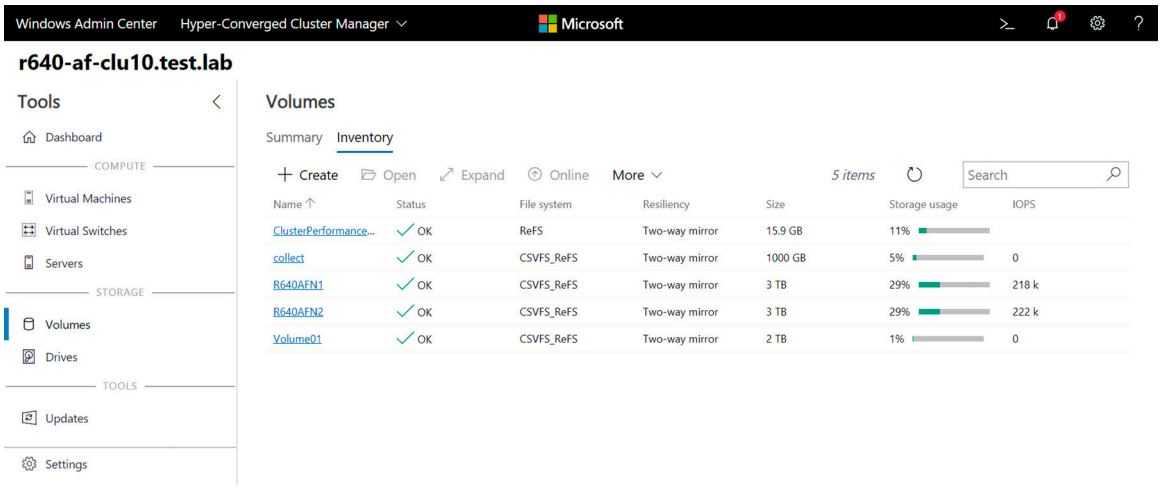

Figure 10. Volumes Inventory

Create volume in Storage Spaces Direct

To create volume in Storage Spaces Direct in WAC:

1 Go to **Volumes** > **Inventory**.
2 Click **Create**.
The **Create volume** window is displayed.
3 Enter the volume name, resiliency, size of the volume, and click **Create**.
The new volume is created.

Manage volumes

To open, expand, delete, or to make the volume offline:

1 Go to **Volumes** > **Inventory**.
2 Click the volume name.
3 Click **Open** to open the volume folder.
4 Click **Offline** or **Delete** to make the volume offline, or to delete the volume.
5 Click **Expand** to expand the volume.
The **Expand volume** window is displayed.
6 Enter the additional size of the volume.
7 Select the volume size from the drop-down list and click **Expand**.

Enable data deduplication on Storage Spaces Direct

Data deduplication helps to maximize free space on the volume by optimizing duplicated portions on the volume without compromising data fidelity or integrity.

NOTE: To enable data deduplication on an HCI cluster, ensure that the data deduplication feature is enabled on all the cluster nodes.

To enable data deduplication feature, run the following PowerShell command *Install-WindowsFeature FS-Data-Deduplication*.

To enable data deduplication and compression on Storage Spaces Direct volume, perform the following:

1 Go to **Volumes** > **Inventory**.
2 Click the volume for which data deduplication needs to be enabled.
3 In the optional features, switch the ON button to enable deduplication and compression on that volume.
The **Enable Deduplication** window is displayed.
4 Click **Start** and select **Hyper-V** from the drop-down list.
5 Click **Enable Deduplication**.
Deduplication is enabled and the Storage Spaces Direct volume is compressed.

Monitor and manage virtual machines

You can monitor and manage the virtual machines hosted on the HCI cluster using WAC.

To access the virtual machines hosted on HCI cluster, click the HCI cluster and go to **Virtual machines**. On the right pane the summary tab and the inventory tab are displayed.

The Summary tab provides the following information of the virtual machine environment of the HCI cluster:

- Total number of virtual machines, their state, and alerts.
- Host and guest CPU utilization.
- Host and guest memory utilization.
- Virtual machine total IOPS, Latency, and throughput information.

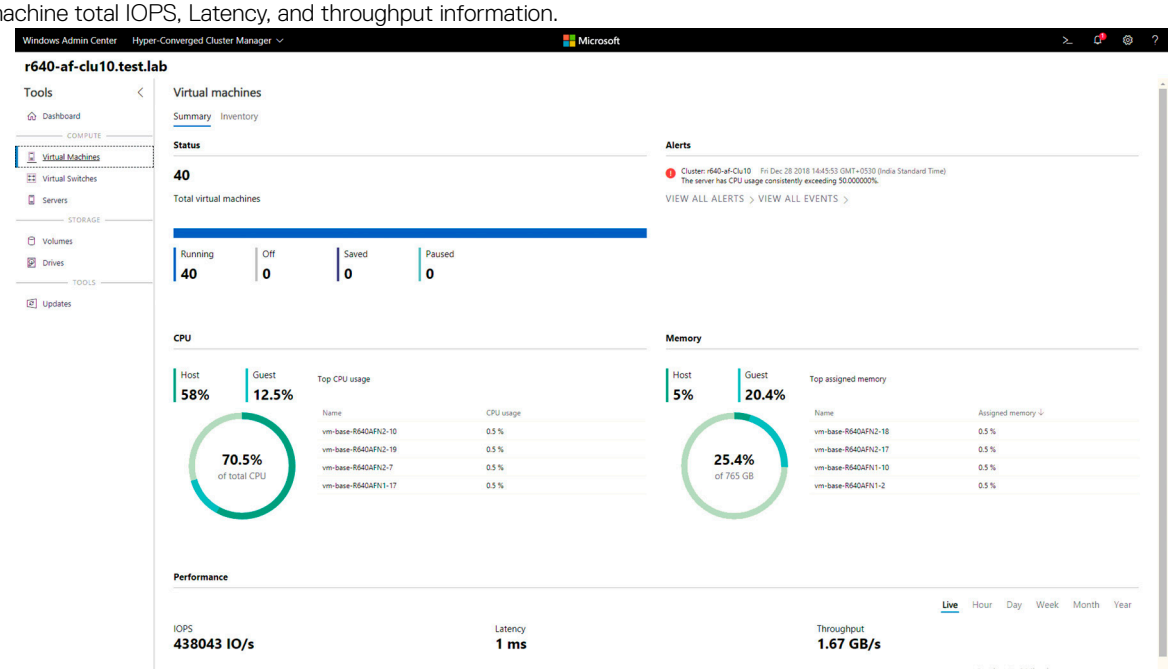

Figure 11. Virtual-machines Summary

The inventory tab provides the list of virtual machines hosted on the HCI cluster and provide access to manage virtual machines.

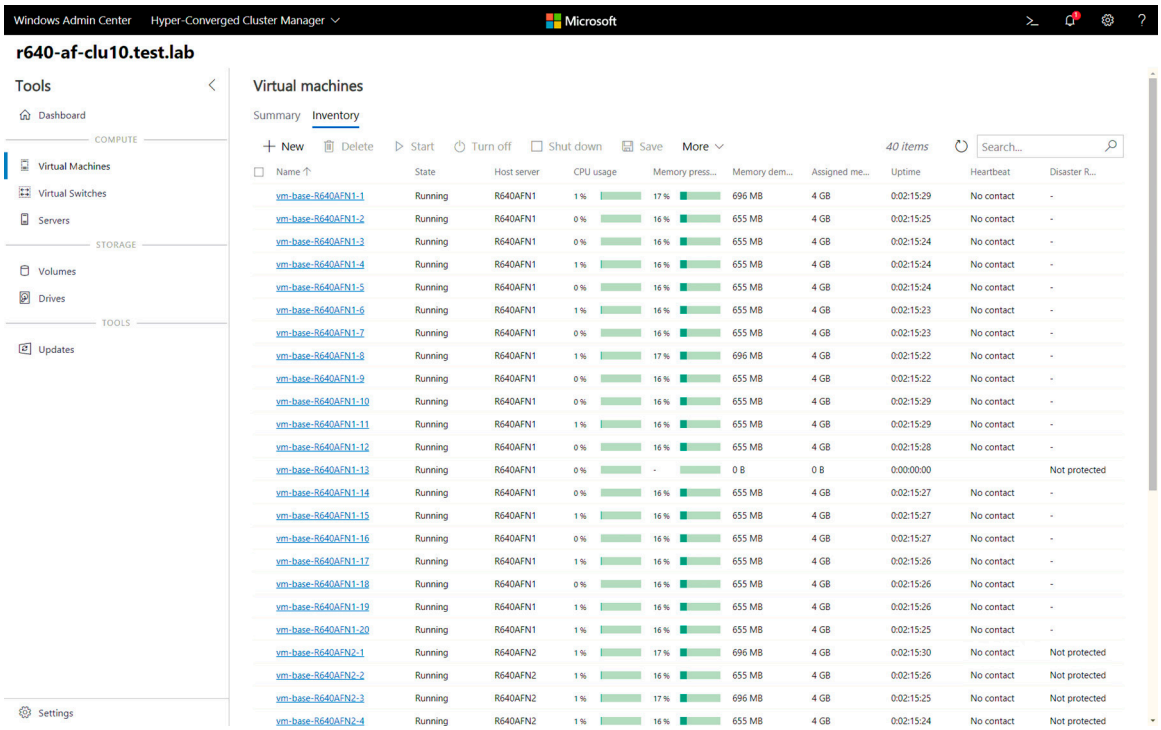

Figure 12. Virtual machines Inventory

You can also perform the following tasks from the WAC console:

- View list of virtual machines hosted on HCI cluster.
- View individual virtual machine state, Host server information, virtual machine uptime, CPU, memory utilization, and so on.
- Create new virtual machine.
- Modify virtual machine settings.
- Set up VM protection.
- Delete, start, turn off, shut down, save, delete saved state, pause, resume, reset, add new checkpoint, move, rename, and connect virtual machines.

Managing virtual switches

Virtual Switches tool in WAC enables you to manage Hyper-V virtual switches of the cluster nodes.

The virtual switches tool supports the following features:

- View existing virtual switches on the server
- Create new virtual switch
- Modify virtual switch properties
- Delete virtual switch

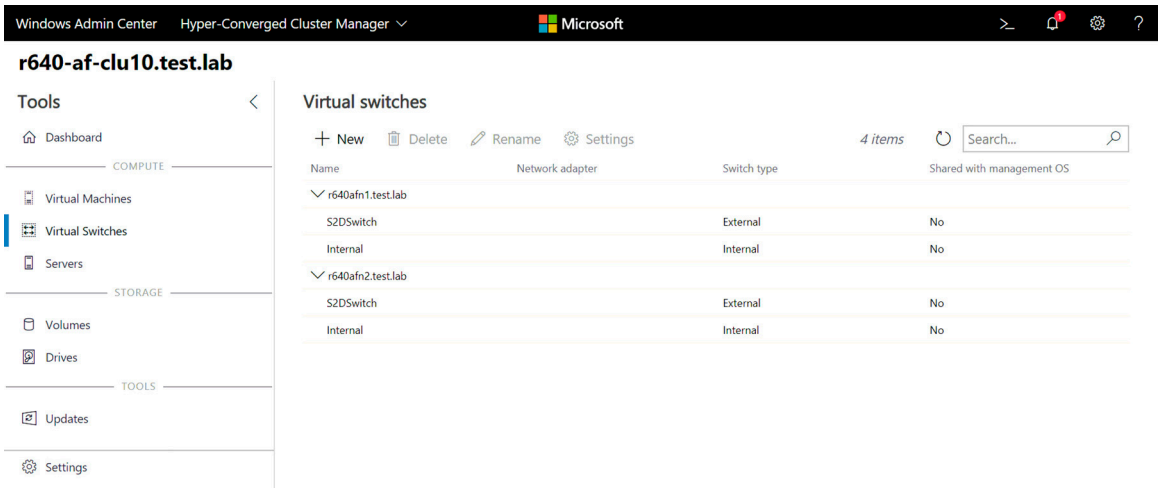

Figure 13. Virtual switches

Manage Windows updates

You can manage Windows updates on a cluster node. All the updates are performed in cluster-aware mode.

1 Click the HCI cluster and go to **Updates**.
On the right pane the **Available updates** and **Update history** tabs are displayed.
The Available updates tab provides information on the last updates and their status.

2 Click **Check Available Updates** to view the updates available for the cluster nodes.

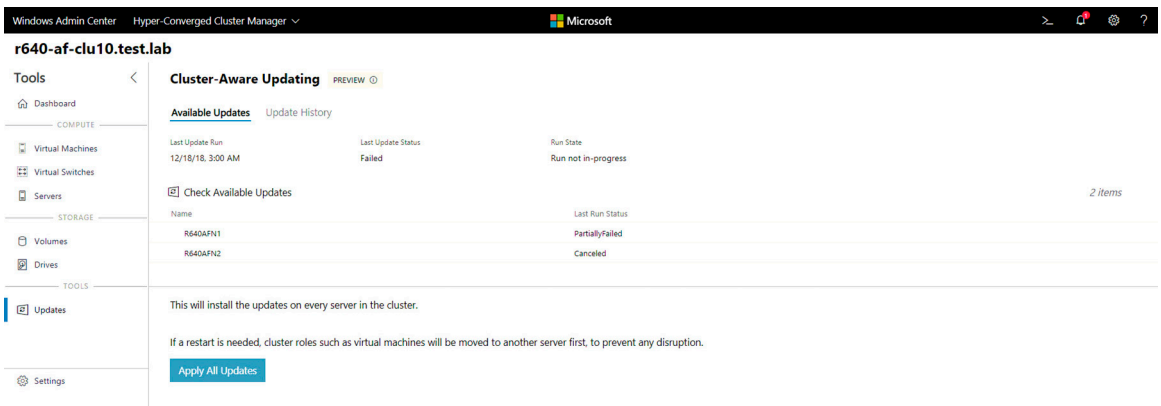

Figure 14. Available updates

3 Click **Apply All Updates** to install updates to the cluster nodes.
The updates are installed in a cluster-aware mode.

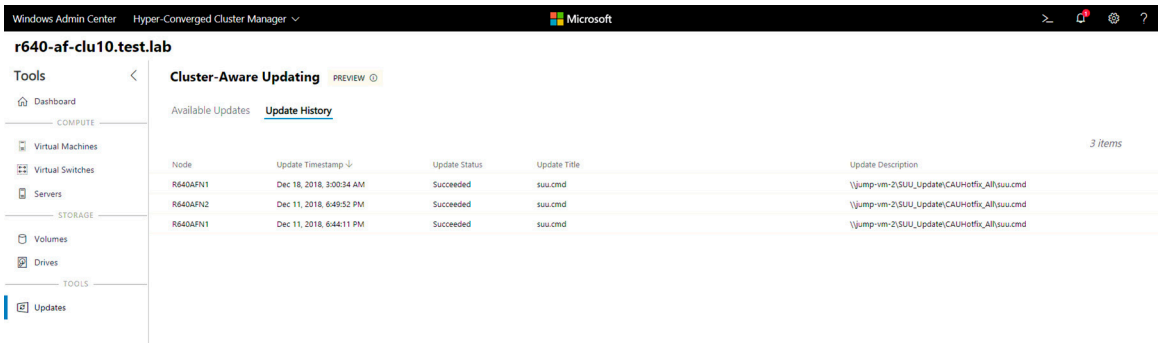

Figure 15. Update history

The **Update History** tab provides the information on the previous updates including cluster node name, last run time stamp, update status, update title, and update description.

4

Dell EMC Solutions for Azure Stack HCI Life Cycle Management

OpenManage Essentials (OME) 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. You can perform the following 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 Solutions for Azure Stack HCI

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.

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 WSMan protocol.

Access OME CLI prompt

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

Installing Hyper-V and failover cluster management tools

To manage the Azure Stack HCI 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

```xml
<?xml version="1.0" encoding="utf-8" ?>
<DiscoveryConfiguration Enabled="True">
  <NetMask>
    255.255.224.0
  </NetMask>
  <ICMPConfig>
    <Timeout>600</Timeout>
    <Retries>1</Retries>
  </ICMPConfig>
  <SNMPConfig Enabled="False">
    <SNMPV1V2CConfig Enabled="True">
      <GetCommunity>public</GetCommunity>
      <SetCommunity></SetCommunity>
    </SNMPV1V2CConfig>
    <SNMPV3Config Enabled="True">
      <SNMPV3Username>user1</SNMPV3Username>
      <SNMPV3AuthenticationProtocol>SHA1</SNMPV3AuthenticationProtocol>
      <SNMPV3EncryptionProtocol>AES</SNMPV3EncryptionProtocol>
    </SNMPV3Config>
    <Timeout>4</Timeout>
    <Retries>2</Retries>
  </SNMPConfig>
  <WMIConfig Enabled="False">
    <UserName>Administrator</UserName>
  </WMIConfig>
  <StoragePowerVaultConfig Enabled="False">

  </StoragePowerVaultConfig>
  <StorageEMCConfig Enabled="False">
    <UserName>Administrator</UserName>
    <Port>443</Port>
  </StorageEMCConfig>
  <WSManConfig Enabled="True">
    <UserName>root</UserName>
    <Timeout>4</Timeout>
    <Retries>6</Retries>
    <Port>443</Port>
    <SecureMode Enabled="True" SkipNameCheck="True" TrustedSite="True">
      <CertificateFile></CertificateFile>
    </SecureMode>
  </WSManConfig>
  <IPMIConfig Enabled="False">
    <UserName>Administrator</UserName>
    <KGKey></KGKey>
    <Timeout>5</Timeout>
    <Retries>2</Retries>
  </IPMIConfig>
  <SSHConfig Enabled="False">
    <UserName>Administrator</UserName>
    <Timeout>5</Timeout>
    <Retries>2</Retries>
    <Port>22</Port>
  </SSHConfig>
</DiscoveryConfiguration>
```

Figure 16. 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.

```
Name,SubnetMask,DisplayName,
172.18.129.4,255.255.224.0,Node1,
172.18.129.5,255.255.224.0,Node2,
172.18.129.6,255.255.224.0,Node3,
172.18.129.7,255.255.224.0,Node4,
```

Figure 17. 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 a discovery 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.

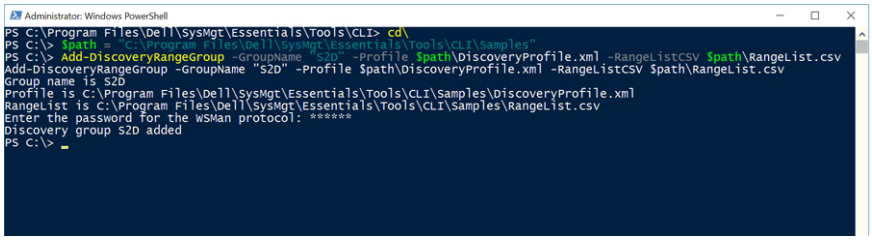

Figure 18. Add discovery range

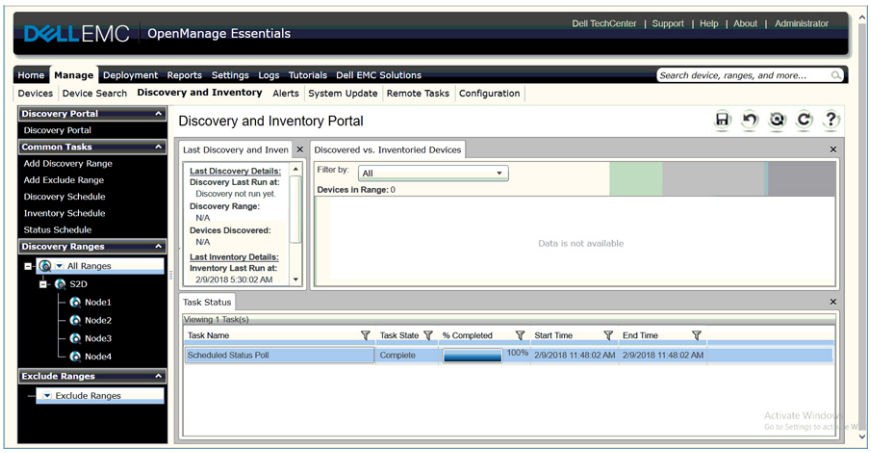

Figure 19. Discovery and inventory portal

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 is not automatically be submitted for discovery and inventory. To submit the range for discovery and inventory, run the `Set-RunDiscoveryInventory -GroupName "Group Name"` command.

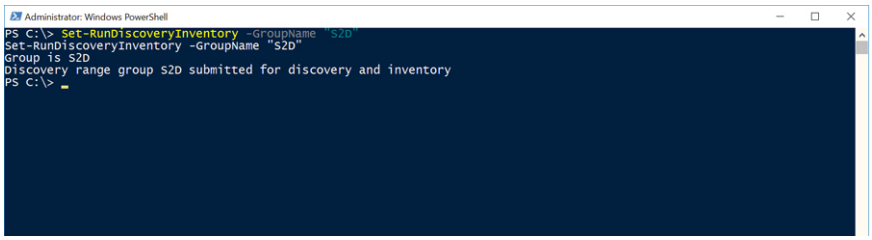

Figure 20. Run discovery inventory

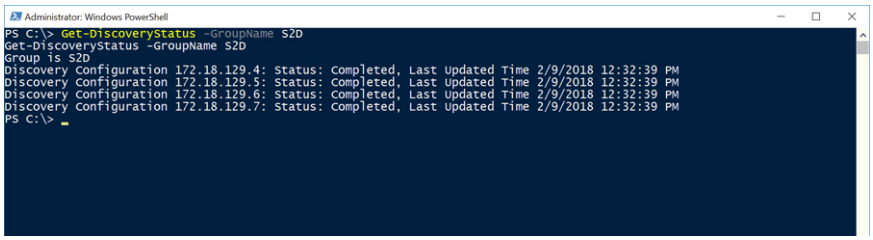

Figure 21. Discovery status

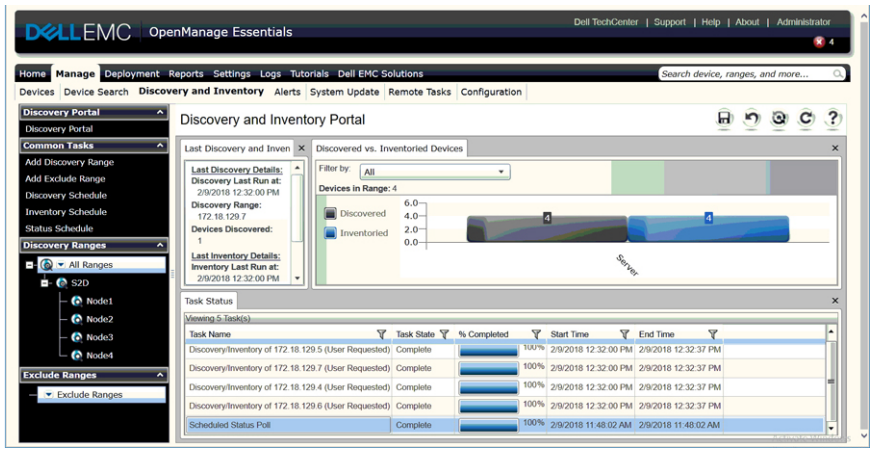

Figure 22. OME inventory status

Create custom group and add discovered devices to the custom group

Use the following command 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
```

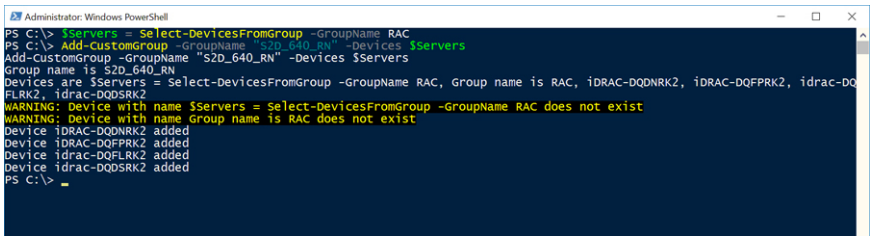

Figure 23. Creating custom group

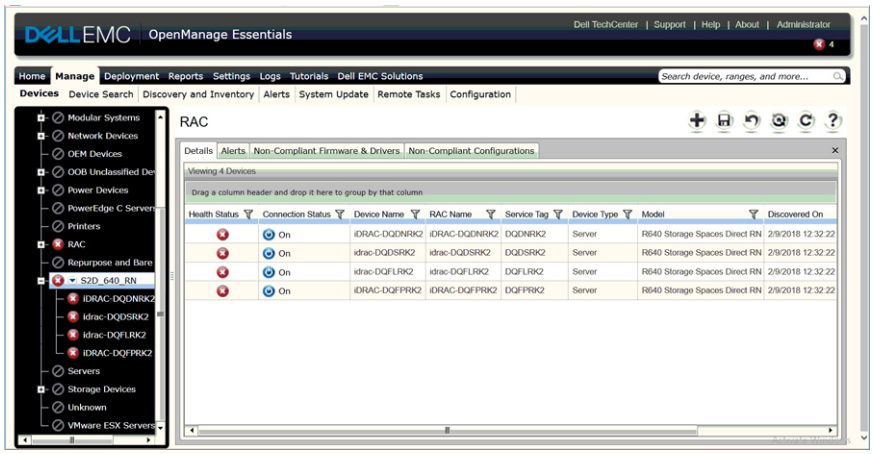

Figure 24. Custom group created

Obtain firmware catalog for Storage Spaces Direct Ready Nodes

For Storage Spaces Direct Ready Node, Dell EMC recommends that you use a Storage Spaces Direct specific catalog for a qualified set of firmware and BIOS. Get the latest catalog from one of the following options and copy to a location in the OME server.

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

Obtain Storage Spaces Direct Ready Node firmware catalog from FTP site

To download and extract the Storage Spaces Direct Ready Node catalog:

1. Download the WSSD catalog file from https://downloads.dell.com/catalog/WSSD-Catalog.xml.gz.
2. Extract the **WSSD-Catalog.xml.gz** file.
3. Copy the extracted **WSSD-Catalog.xml** file to a location in the OME server and associate the custom catalog to a device group using OME.

Obtain firmware catalog for Storage Spaces Direct Ready Nodes using Dell EMC Repository Manager

Generate the catalog along with the firmware and drivers using Dell EMC Repository Manager (DRM) and copy to a location in the OME server.

1 Install Dell EMC Repository Manager (DRM) version 3.0.1.423 or later.
2 On the DRM home page, click the **Dell EMC Repository Manager** drop-down list.
3 In the **Manage** section, click **Application Preferences**.
The **Preferences** window is displayed.
4 Click **Plug-ins**.
5 Select all the plug-ins and click **Update**.
A message is displayed about the successful completion of the update.
6 Click **Catalogs**.
7 Select all the catalogs and click **Update**.
8 Click **Close** to close the **Preferences** window.
9 On the home page, click **Add Repository**.
The **Add Repository** window is displayed.
10 Enter the **Repository name** and **Description**.
11 Select **Index Catalog-<version>** from the **Base Catalog** drop-down menu.
12 Select **Dell EMC Ready Solutions for Microsoft WSSD Catalog** from the **Catalog Group**.
13 Select the latest catalog from the **Catalogs** section.
14 Click **Save**.
The **Dell EMC Ready Solutions for Microsoft WSSD Catalog** is populated in the **Base Catalog** section.
15 In the **Manual** Repository Type, click **All systems in base catalog** and then click **Add**.
The repository is displayed on the repository dashboard available in the home page.
16 Select the repository and click **Export**.
The **Export Deployment Tools** window is displayed.
17 Select the location to export files and click **Export**.
The files are exported to the specified location.

Run the following PowerShell commands to copy the drivers to a different folder and install or update from the Storage Spaces Direct Ready Node.

```
$Drivers=Get-ChildItem -Path C:\ExportedPath\ -Include *Driver* -File -Recurse
```

```
Copy-Item -Path $Drivers -Destination C:\Q1\Drivers\
```

Associate custom catalog to device group using OME

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

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.
5 Click **Finish**.

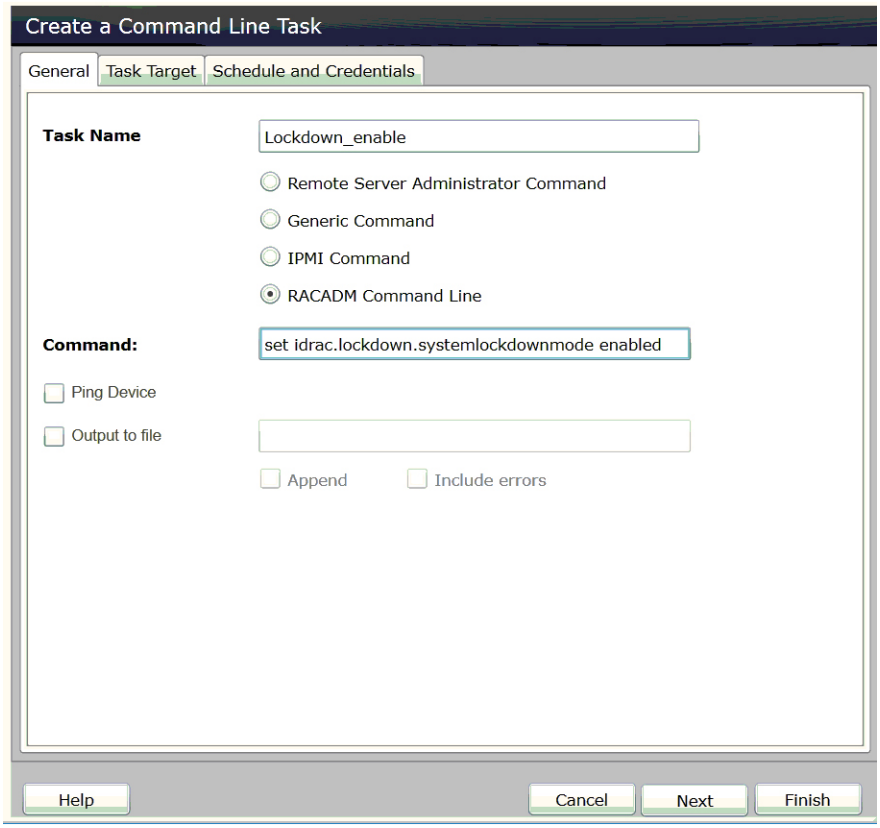

Figure 25. Enabling lockdown

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 networks are available using `Get-ClusterNetwork` command.
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 drives 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.

Placing 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). Verify that all the roles and virtual drives are drained properly and operational in other nodes after being moved:

   ```
   Suspend-ClusterNode -name “Hostname” -Drain
   ```

2. Place the target node in maintenance mode:

   ```
   Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_.FriendlyName -eq
   "<Hostname>"} | Enable-StorageMaintenanceMode
   ```

3. Run the `Get-Physical Disk` command, and ensure that the Operational Status value is in maintenance mode for the drives that belong to that server. This can also be done by running the following command and verifying that they all belong to the paused node:

   ```
   Get-Storagepool -IsPrimordial 0 |Get-PhysicalDisk | ? operationalstatus -eq 'In Maintenance
   Mode' |Get-StorageNode -PhysicallyConnected
   ```

4. Turn off the System Lockdown mode.

Enabling system lockdown prevents the node from making any updates 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 about System Lockdown, see Storage Spaces Direct Ready Node lockdown.

Update Storage Spaces Direct Ready Node using OME console

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

1. Go to **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.
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.

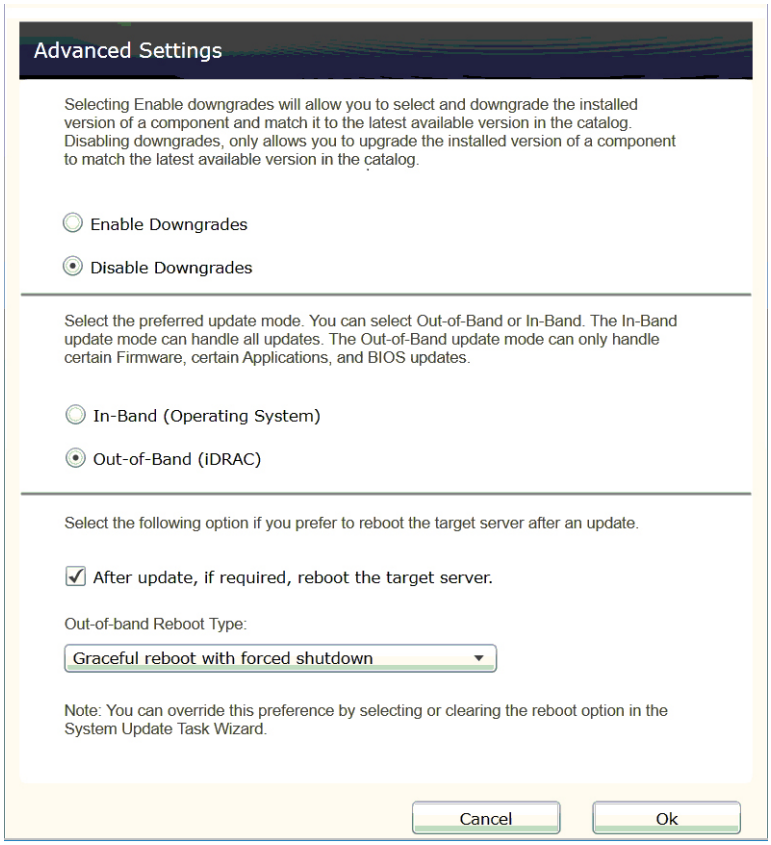

Figure 26. Advanced settings

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 displays Out-of-Band for all the available components. The **System Update** window is displayed.

7 Provide the task name.

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.

Update Storage Spaces Direct Ready Node 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. Generate the latest Storage Spaces Direct supported catalog file through DRM and copy the file to a network location before proceeding with the update process.

1 Log in to the iDRAC web interface.

2 Click **Maintenance** > **System Update**.

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 in the following figure:

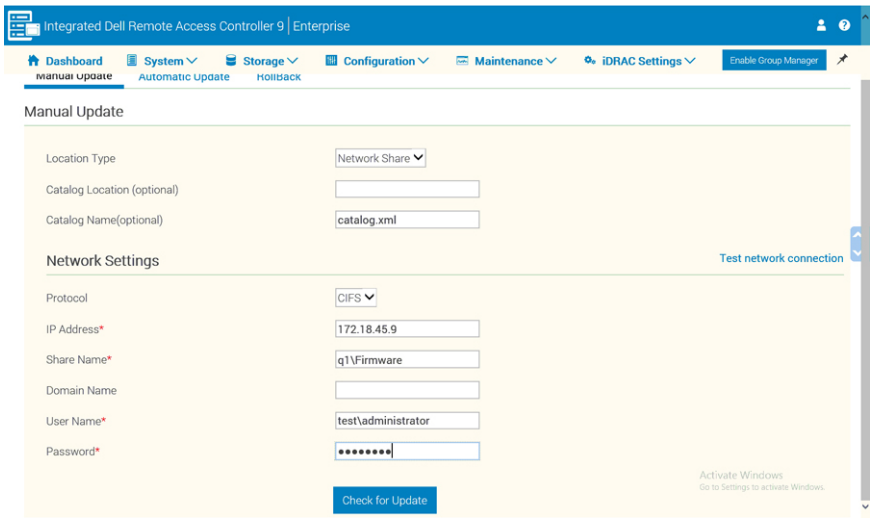

Figure 27. Check for updates

5 Click **Check for updates**.

A list of available updates is displayed as shown in the following figure.

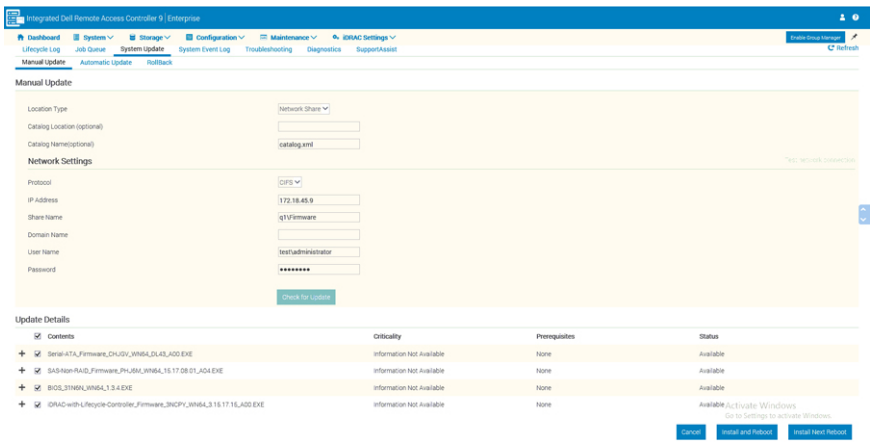

Figure 28. Select updates

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

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 list of compliant servers:

a Click **Manage** > **System Update**.
b In the **System Update** page, select the **Compliant Systems** tab.

2 To view the list of 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.

Updating Out of Box drivers

For certain system components, there may be a need to update the drivers to the latest, and Dell EMC supported versions as listed in the Supported Firmware and Software Matrix.

Run the following PowerShell command to retrieve the list of all driver versions installed on the local system:

```
Get-PnpDevice | Select-Object Name, @{l='DriverVersion';e={(Get-PnpDeviceProperty -InstanceId
$_.InstanceId -KeyName 'DEVPKEY_Device_DriverVersion').Data}} -Unique
```

After you identify the required driver version, you can download the driver installers from https://www.dell.com/support or by using the Dell EMC Repository Manager as described in the section Obtain firmware catalog for Storage Spaces Direct Ready Nodes using Dell EMC Repository Manager.

After the drivers are downloaded, copy the identified drivers to Storage Spaces Direct Ready Nodes from where you can manually execute the driver DUP files to install the OOB drivers.

Exit Storage Spaces Direct Ready Node from maintenance mode

After updating the Storage Spaces Direct Ready Node, exit the Storage maintenance mode and node maintenance mode by using the following commands:

```
Get-StorageFaultDomain -type StorageScaleUnit | Where-Object {$_.FriendlyName -eq "<Hostname>"}
| Disable-StorageMaintenanceMode
```

```
Resume-ClusterNode -Name “Hostname” -Failback Immediate
```

The Storage Spaces Direct operation to rebuild and re-balance the data to ensure load balancing is initiated.

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.

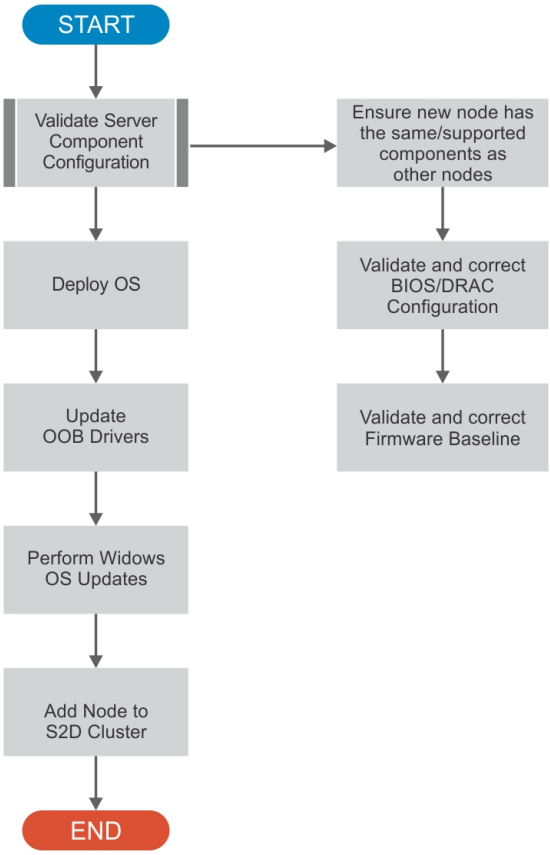

Figure 29. Expanding Ready Node cluster

Storage Spaces Direct Node Expansion

In an HCI cluster, adding server nodes increases the storage capacity, improves the overall storage performance of the cluster, and provides more compute resources to add virtual machines. Before adding new server nodes into an HCI cluster, complete the following requirements:

- Verify that the processor model, HBA 330, and NICs are of the same configuration as the current nodes on the cluster and PCIe slots.
- Ensure that all disk type and the amount in each node are same as to the node in use. Do not combine two different disk types in the same cluster or node. For example, you cannot combine SATA and SAS HDD/SSD drives in the same node or cluster. Following are the supported options to expand storage capacity of the cluster.

Table 1. Options to expand storage capacity of the cluster

Option 1	Option 2
Option 1 meeting the following conditions: – Drive is listed in the Support Matrix – Same drive manufacturer – Same capacity and endurance – Latest model – Latest firmware	Option 2 meeting the following conditions: – Drive is listed in the Support Matrix – Different drive manufacturer – Same capacity and endurance – Different model – Different firmware

- Ensure that the BIOS, drivers, firmware, and chipset are as listed in the support matrix.
- Apply the BIOS configuration to the node and configure iDRAC. For more information about configuring the node, see the Hyper-converged infrastructure deployment section in the *Dell EMC Solutions for Microsoft Azure Stack HCI* 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

Ensure that the following tasks are completed:

1 Pass cluster validation and SES device compliance tests.
2 Verify that the nodes are compliant with the firmware baseline.
3 Update the hardware timeout configuration for the Spaces port.
4 After the node configuration, update Microsoft Windows to bring the node to the same level as the cluster.

Adding server nodes manually

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

To manually add server nodes to the cluster, see https://technet.microsoft.com/windows-server-docs/storage/storage-spaces/add-nodes.

Storage Spaces Direct Storage expansion

In an HCI cluster, expanding storage by adding drives on the available slots on the cluster nodes adds storage capacity of the cluster and improves storage performance. Before the storage expansion, ensure that all disk types and the amount in each node is same and are equal to that of the node in use. Do not combine two different disk types in the same cluster or node. For example, you cannot combine SATA and SAS HDD/SSD drives in same node or cluster.

Following are the supported options to expand storage capacity of the cluster:

- Option 1: Expand the storage with the same drive manufacturer, capacity, endurance, latest model, latest firmware, and if it is available on the Ready Node support matrix.
- Option 2: Expand the storage with different drive manufacturer, model, firmware, same capacity, endurance, and if it is available on the Ready Node support matrix.

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 are claimed into the existing pool and Storage Spaces Direct starts the 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.

Extend volumes

Volumes that are 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/windows-server-docs/storage/storage-spaces/resize-volumes.

Perform 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 Solutions for Azure Stack HCI 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 Log in to the iDRAC web interface.

2 Go to **Storage** > **Controllers**.

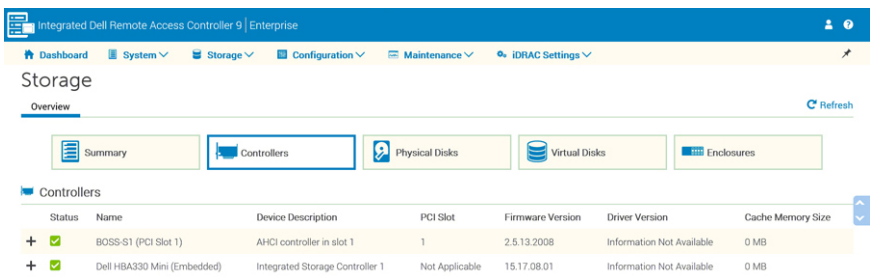

Figure 30. View controllers

3 Go to **Configuration** > **Storage Configuration** > **Virtual Disk Configuration** and then click **Create Virtual Disk**.

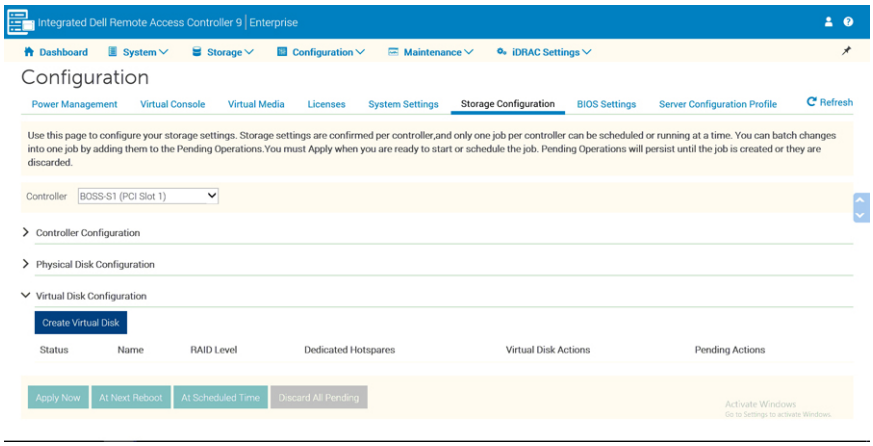

Figure 31. Creating virtual disk

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

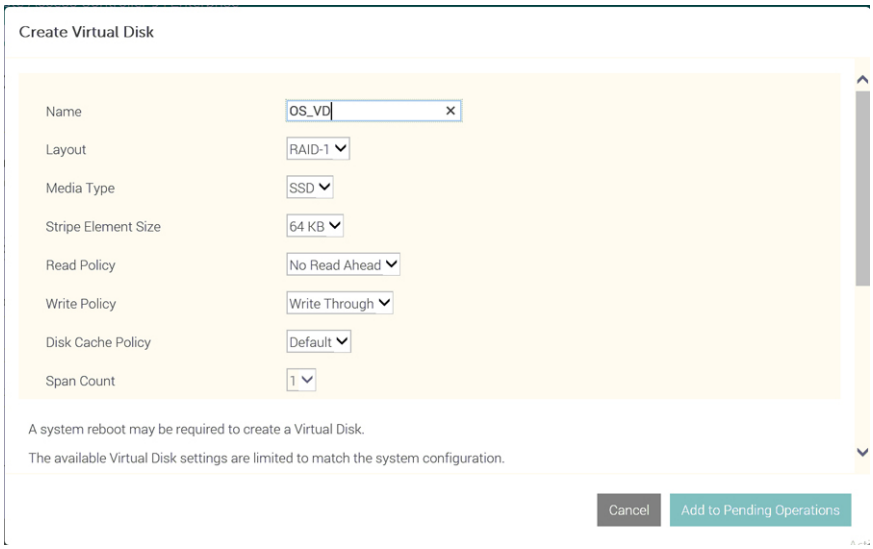

Figure 32. Provide virtual disk name

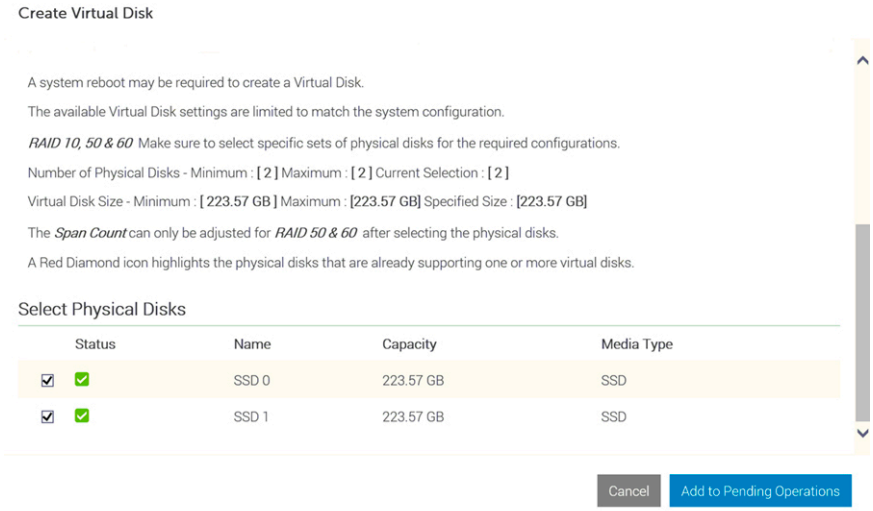

Figure 33. Set Physical Disks

5 Click **Add Pending Operations**.

6 Go to **Configuration** > **Storage Configuration** > **Virtual Disk Configuration**.

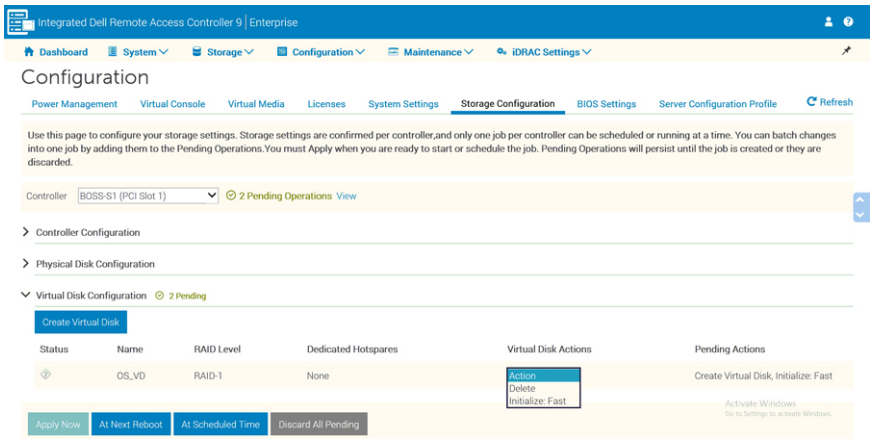

Figure 34. Initialize configuration

7 Select the VD and select **Initialize: Fast in Virtual Disk Actions**.

8 Reboot the server.

NOTE: The VD creation process may take several minutes to complete.

9 After the initialization is completed successfully, the VD health status is displayed.

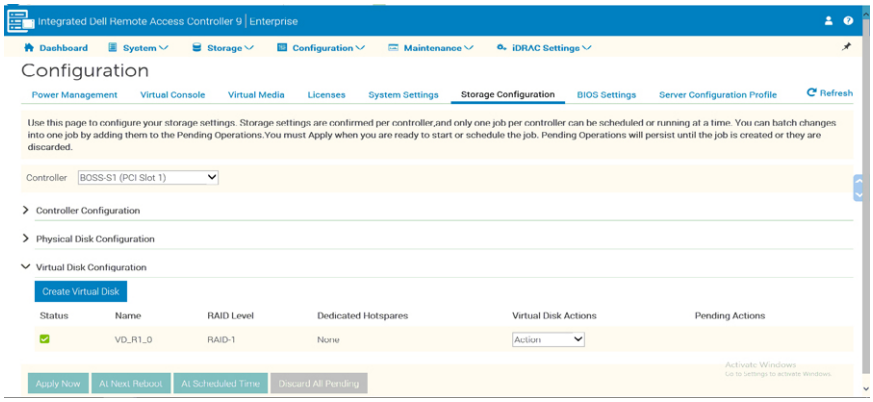

Figure 35. Virtual Disk Health status

Operating system recovery

This section provides an overview of steps involved in operating system recovery on the Dell EMC Solutions for Azure Stack HCI.

NOTE: Ensure that the RAID 1 VD created on the BOSS M.2 drives is reinitialized.

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.

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.

Factory Operating system recovery

For the factory installed OEM license of the operating system, Dell EMC recommends 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 triggers 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 hostname
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 hostname, AD domain join, QoS policy, RDMA configuration, configuring firewall, see the *Dell EMC Storage Spaces Direct Ready Node Deployment Guide* available at www.dell.com/wssdmanuals.

Some instructions for updating OOB drivers, configuring host networking, changing hostname, AD domain join, QoS policy, RDMA configuration, configuring firewall may be different for the 2-Node back-to-back connected infrastructure configuration. For more information, see the *Deployment Guide for 2-Node hyper-converged back-to-back connected infrastructure with R640 Storage Spaces Direct Ready Node* available in www.dell.com/wssdmanuals for deployment instructions specific to back-to-back connected infrastructure cluster.