

Dell EMC Solutions for Microsoft Azure Stack HCI—QuickStart Configurations for ROBO and Edge Environments

Integration and Deployment Guidance

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.

© 2018 -2019 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

1 Audience and scope.....	4
Assumptions.....	4
Known issues.....	4
2 QuickStart configurations for ROBO/Edge environments.....	5
R740xd Storage Spaces Direct Ready Node.....	6
R640 Storage Spaces Direct Ready Node.....	6
Dell EMC Networking S4112F-ON.....	7
Network connectivity.....	7
Switch configuration.....	8
3 Deployment guidance.....	9
Firmware baseline.....	10
QLogic NIC configuration.....	10
Install operating system.....	10
Manual OS deployment.....	11
Factory installed OS.....	11
Install roles and features.....	11
Change host name.....	11
Configure host networking.....	12
Update page file settings.....	13
Enabling jumbo frames.....	14
Domain join.....	14
Validate cluster creation readiness.....	14
Create cluster and enable Storage Spaces Direct.....	15
4 Post-cluster creation configuration.....	16
Remove host management network from Live Migration.....	16
Update hardware timeout for Spaces port.....	16
Configuring cluster witness.....	16
Additional resources.....	17

Audience and scope

This deployment guide focuses on deploying a Dell EMC Solution for Azure Stack HCI as a Quick Start hyper-converged infrastructure solution. This deployment guide includes an overview of the solution infrastructure, guidance on how to integrate the solution components, and instructions for preparing and deploying the solution infrastructure. This guide is applicable only to infrastructure built by using the validated and certified Dell EMC Microsoft Storage Spaces Direct Ready Nodes.

The audience for this document includes, but is not limited to, 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 or Windows Server 2019 Hyper-V and Storage Spaces Direct.

Customers who do not have Volume License agreements with Microsoft can order Dell EMC Microsoft Storage Spaces Direct Ready Nodes with the operating system pre-installed at the factory with OEM license or bare metal.

The Storage Spaces Direct Cluster deployment can be done in two ways:

- Dell EMC Services led: Certified deployment engineers can deploy the solution, which ensures accuracy and speed, reduced risk and downtime.
- Customer led: Customers can deploy the solution by referring to this deployment guide, provided they have the qualified level of technical expertise.

NOTE: Instructions in this deployment guide are applicable only to the generally available OS build of Windows Server 2016 with the latest applicable updates and Windows Server 2019 GA build with latest OS 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. Dell EMC recommends that you update the host OS with latest cumulative updates from Microsoft before starting the cluster creation and configuration tasks.

NOTE: While the instructions in this deployment guide can be used with other Dell EMC Microsoft Storage Spaces Direct Ready Nodes, R640 Storage Spaces Direct Ready Node has been used as an example for the deployment instructions.

Topics:

- [Assumptions](#)
- [Known issues](#)

Assumptions

This deployment 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.
- 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. See <https://www.dell.com/support/article/sln313305>.

QuickStart configurations for ROBO/Edge environments

The Dell EMC Quick Start for Azure Stack HCI is the latest offering in the portfolio of Dell EMC Solutions for Microsoft Azure Stack HCI. This quick start solution offers customers an easy to deploy and affordable starter infrastructure with best-in-class deployment services and support.




Figure 1. High level overview

This quick start solution is available in both hybrid and all-flash configurations, and each in three different sizes—small, medium, and large—and it is built on the R740xd Storage Spaces Direct Ready Node (for hybrid) and R640 Storage Spaces Direct Ready Node (for all-flash). This solution introduces a highly compact, redundant dual switch fabric in a 1U form factor. The Dell EMC Networking S4112F-ON switches support a dedicated network topology for both storage and management.

For more information about Dell EMC Solutions for Azure Stack HCI, see <https://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Microsoft-Storage-Spaces-Direct-Ready-Nodes-Solution-Overview-en.pdf>.

Topics:

- [R740xd Storage Spaces Direct Ready Node](#)

- [R640 Storage Spaces Direct Ready Node](#)
- [Dell EMC Networking S4112F-ON](#)

R740xd Storage Spaces Direct Ready Node

The R740xd Storage Spaces Direct Ready Node, based on Dell EMC PowerEdge R740xd server, is optimized for software-defined storage implementations that enable converged infrastructure (CI) and hyper-converged infrastructure (HCI) deployments such as the architecture recommended in this deployment guide. With 2 CPU sockets and a wide range of CPU options, this ready node provides capabilities to match your computational needs.

This 2U, rack mounted server provides high storage density with different drive options for a hybrid configuration. The following table provides an overview of the hybrid quick start configurations.

Table 1. Hybrid quick start configurations

Bundle	Small	Medium	Large
Server	R740xd Storage Spaces Direct Ready Node - 12 drive chassis	R740xd Storage Spaces Direct Ready Node - 12 drive chassis	R740xd Storage Spaces Direct Ready Node - 12 drive chassis
Processor	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 4110)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 5118)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 6148)
Memory	All Dell EMC Storage Spaces Direct Ready Node options (Default: 192 GB)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 384 GB)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 576 GB)
Storage raw capacity per node	8 TB SATA SSD: 2 x 960 GB MU SATA HDD: 4 x 2 TB SATA	32 TB SAS SSD: 4 x 960 GB MU SAS HDD: 8 x 4 TB NL-SAS	64 TB SATA SSD: 4 x 1.92 TB MU SATA HDD: 8 x 8 TB SATA
Cluster usable capacity**	7 TB	35 TB	69 TB
Network card	Dual Port 25 GbE QLogic 41262 rNDC: 2 x 10 GbE + 2 x 1 GbE Broadcom, QLogic, and Intel (default)	Dual Port 25 GbE QLogic 41262 rNDC: 2 x 10 GbE + 2 x 1 GbE Broadcom, QLogic, and Intel (default)	Dual Port 25 GbE QLogic 41262 rNDC: 2 x 10 GbE + 2 x 1 GbE Broadcom, QLogic, and Intel (default)

** Usable capacity calculated accounting for reserve and approximate overhead.

R640 Storage Spaces Direct Ready Node

R640 Storage Spaces Direct Ready Node, a 1U rack server based on PowerEdge R640 server, is optimized for software-defined storage implementations that enable converged infrastructure (CI) and hyper-converged infrastructure (HCI) implementations. This Ready Node supports up to 2 CPU sockets with a wide range of options in number of cores per CPU socket and 1.5 TB of memory when using DDR4 DIMMs.

The following table provides an overview of the all-flash quick start configurations.

Table 2. All-flash quick start configurations

Bundle	Small	Medium	Large
Server	R640 Storage Spaces Direct Ready Node - 10-drive chassis	R640 Storage Spaces Direct Ready Node - 10-drive chassis	R640 Storage Spaces Direct Ready Node - 10-drive chassis
Processor	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 4110)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 5118)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 2 x 6148)
Memory	All Dell EMC Storage Spaces Direct Ready Node options (Default: 192 GB)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 384 GB)	All Dell EMC Storage Spaces Direct Ready Node options (Default: 576 GB)
Storage raw capacity per node	3.84 TB SATA SSD: 4 x 960 GB MU SATA	15.36 TB SATA SSD: 8 x 1.92 TB MU SATA	38.4 TB SAS SSD: 10 x 3.84 TB MU SAS
Cluster usable capacity**	4 TB	17 TB	43 TB
Network card	Dual Port 25 GbE QLogic 4126 rNDC: 2 x 10 GbE + 2 x 1 GbE (Default: Intel)	Dual Port 25 GbE QLogic 41262 rNDC: 2 x 10 GbE + 2 x 1 GbE (Default: Intel)	Dual Port 25 GbE QLogic 41262 rNDC: 2 x 10 GbE + 2 x 1 GbE (Default: Intel)

** Usable capacity calculated accounting for reserve and approximate overhead.

Dell EMC Networking S4112F-ON

Dell EMC Networking S4112F-ON switch supports multiple link speeds from 10 Gbe, 25 Gbe to 100 Gbe, in SFP+ and SFP28 form factors, respectively. The switch when combined as a pair, provides a highly redundant, highly compact switch fabric in a 1U chassis form factor. The OS10 full feature operating system supports both Layer 2 and Layer 3 protocols. For more information about Dell EMC Networking S4112F-ON features, see the OS10 CLI guide at <https://support.dell.com>.

This compact switch configuration is ideal for deployments that support 4 Node cluster configurations where the footprint size is critical. In addition to providing the necessary throughput for a typical 4 node S2D cluster the switch fabric has additional capacity for expansion.

Network connectivity

The network fabric for this quick start solution, implements a non-converged topology for the in-band/out-of-band Management and Storage networks. The RDMA capable QLogic FastLinQ 41262 adapters are used in support of the storage traffic while the rNDC provides the 10 GbE SFP+ bandwidth used for the host management and VM network traffic. The out-of-band Management utilizes a 1000BASE-T SFP Copper RJ-45 transceiver into the 10 Gbe switch port as illustrated in figure 2.

The optimized network architecture of the switch fabric ensures redundancy of both the S2D storage and management networks. In this design, the storage networks are isolated to each respective switch (Storage 1 to TOR 1; Storage 2 to TOR 2) and as storage traffic (typically) never traverses the customer LAN networks, the VLT bandwidth has been optimized to fully support redundancy of the management network up to customer Datacenter uplink.

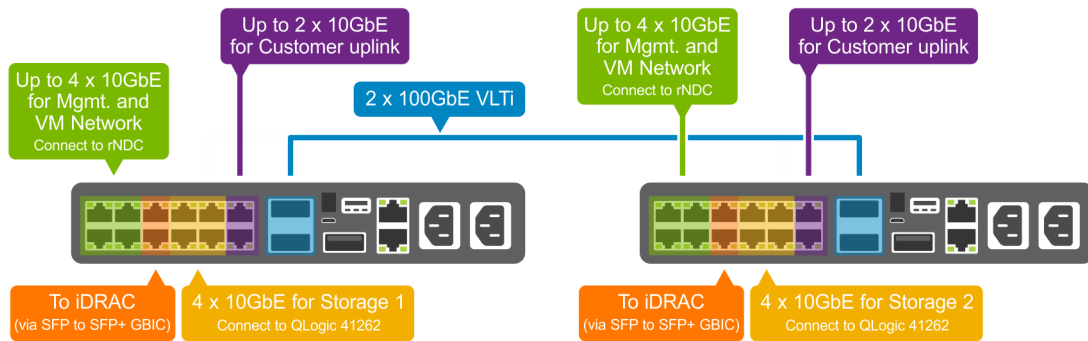


Figure 2. Network connectivity

① | **NOTE:** This topology supports both tagged and untagged VLANs for storage traffic.

Switch configuration

For sample switch configurations, see <https://community.emc.com/docs/DOC-70310>.

Deployment guidance

Do-It-Yourself (DIY) customers can follow the deployment guidance provided in the Dell EMC Solutions for Azure Stack HCI deployment guide for scalable infrastructure available at <https://www.dell.com/azurestackhcinanuals>, to implement this quick start solution. This deployment guide provides information on deployment prerequisites, deployment checklists, and detailed deployment and configuration steps.

Issues that arise during deployment are not covered as part of Dell ProSupport for Software or ProSupport Plus for Software. Support for deployment issues come under a separate paid services package. When you call in with a deployment issue, Dell Tech Support will route you to your Account Manager in Dell EMC Sales. The Account Manager will then help you in purchasing the onsite deployment services package.

The following sections provide a summary of the deployment steps and the PowerShell commands that need to run on the server nodes to complete the cluster configuration:

- 1 Configure Switch and validate connectivity to Nodes and DC infrastructure
- 2 Baseline firmware
- 3 Install Windows Server 2016 or Windows Server 2019 Data Center OS
- 4 Install out of the box drivers
- 5 Install required roles and features
- 6 Change host name
- 7 Configure host networking
- 8 Join cluster nodes to a domain
- 9 Validate cluster creation readiness
- 10 Create cluster and enable Storage Spaces Direct
- 11 Post-cluster creation configuration

Deployment guidance for management services such as Active Directory (AD) and Domain Name Services (DNS) is not a part of the deployment guidance. These services are needed for the cluster functionality and should already exist in the data center environment.

NOTE: This deployment guidance provides instructions and PowerShell commands for manually deploying an Azure Stack HCI cluster. For information about configuring host networking and creating Azure Stack HCI cluster using System Center Virtual Machine Manager (VMM), see <https://community.emc.com/docs/DOC-75346>.

Topics:

- [Firmware baseline](#)
- [QLogic NIC configuration](#)
- [Install operating system](#)
- [Install roles and features](#)
- [Change host name](#)
- [Configure host networking](#)
- [Domain join](#)
- [Validate cluster creation readiness](#)
- [Create cluster and enable Storage Spaces Direct](#)

Firmware baseline

This Dell EMC Ready Nodes for Storage Spaces Direct are validated and certified with certain firmware versions related to the components in the solution infrastructure. This matrix identifies the Dell EMC validated versions of component firmware, OS drivers, and network switch operating system and should be followed for ensuring that the solution infrastructure remains supported and delivers optimal performance.

This support matrix is available at <https://www.dell.com/azurestackhcimanuals> and is updated as new revisions of the software and firmware are validated.

The [firmware and driver update catalog](#) for Dell EMC Solutions for Azure Stack HCI provides a catalog of all validated versions of the ready node components. This catalog can be used with tools such as a OpenManage Essentials or OpenManage Enterprise or iDRAC to perform server component updates.

For the switch OS10 updates, see the Dell network download site at <https://www.force10networks.com/CSPortal20/Main/Login.aspx>

QLogic NIC configuration

QLogic FastLinQ 41262 network adapter supports both iWARP and RoCE for RDMA. This quick start solution is validated with iWarp for RDMA. Therefore, the network adapter must be configured manually to use iWARP for RDMA.

Perform the following steps for each port to configure the QLogic network adapters to use iWarp for RDMA.:

- 1 Press F2 during system boot to enter **System Setup**.
- 2 Click **System BIOS** and select **Device Settings**.
- 3 Select the QLogic network adapter from the list of adapters.
- 4 Click **Device Level Configuration** and ensure that **Virtualization Mode** is set to None.
- 5 Click **Back**, and then click **NIC Configuration**.
- 6 Select the following options in the **NIC Configuration** page:
 - Link Speed—SmartAN
 - NIC + RDMA Mode—Enabled
 - RDMA Operational Mode—iWARP
 - Boot Protocol—None
 - Virtual LAN Mode—Disabled
- 7 Click **Back** and click **Data Center Bridging (DCB) Settings**.
- 8 In the **Data Center Bridging (DCB) Settings** page, set **DCBX Protocol** to Disabled.
- 9 Click **Back**, and then click **Finish**. Click **Yes** to save the settings.
- 10 Click **Yes** to return to the Device Settings page.
- 11 Select the second port of the QLogic adapter and repeat the steps.
- 12 Click **Finish** to return to the System Setup page.
- 13 Click **Finish** to reboot the system.

Install operating system

There are two methods to deploy the operating system:

- Manual OS deployment - A manual method of install from OS deployment to cluster creation.
- Factory OS deployment - Dell EMC Ready Nodes can be ordered from the factory with Windows Server 2016 or Windows Server 2019 OS pre-installed.

 **NOTE:** The steps in the subsequent sections are applicable to either full OS or Server Core.

NOTE: For the PowerShell commands used in this section and subsequent sections that require network adapter name, use the `Get-NetAdapter` cmdlet to retrieve the right value for the associated physical network port. The network adapter names used in the commands in this guide are only shown as examples and may not represent the right naming convention for what is installed in the system.

Manual OS deployment

Dell Lifecycle Controller and Integrated Dell Remote Access Controller provide various options for [operating systems deployment](#). This includes manual or unattended installation by using the virtual media and OS deployment feature in [Lifecycle Controller](#).

The step-by-step procedure on how to deploy operating system is not within the scope of this guide.

The subsequent steps in this guide assume that the Windows Server 2016 Server or Windows 2019 Server Core Data Center edition deployment on the physical server is complete and that you have access to the iDRAC [virtual console of the physical server](#).

Factory installed OS

If the cluster nodes are shipped from Dell EMC factory with pre-installed Windows Server 2016 or Windows Server 2019 Data Center edition OEM license, the Out of box experience (OOBE) needs to be completed. This includes the following steps:

- Selecting Language and Locale Settings
- Accepting Microsoft and OEM EULA
- Setting up password for the Local Administrator account
- Updating the OS partition size and shrinking it as needed

Install roles and features

Windows Server 2016 or Windows Server 2019 Storage Spaces Direct hyper-converged Infrastructure (HCI) cluster deployment and configuration requires enabling the following operating system roles and features.

- Hyper-V service (this is not required if the OS is factory installed)
- Failover clustering
- Data center bridging (this is required only when implementing fully-converged network topology with RoCE for RDMA, and when implementing DCB for the fully-converged topology with iWarp for RDMA.)

These features can be enabled using the `Install-WindowsFeature` PowerShell cmdlet.

```
Install-WindowsFeature -Name Hyper-V, Failover-Clustering, Data-Center-Bridging -  
IncludeAllSubFeature -IncludeManagementTools -Verbose
```

NOTE: Hyper-V role installation requires a reboot of the system. As the subsequent steps require a reboot as well, they are all combined into a single reboot.

Change host name

By default, the OS deployment assigns a random name as the host computer name. For easier identification and uniform configuration, Dell EMC recommends that you change the host name to something that is easily identifiable and relevant. This can be done by using the `Rename-Computer` cmdlet.

```
Rename-Computer -NewName WSSDQS01 -Restart
```

NOTE: This command induces an automatic restart at the end of the rename operation.

Configure host networking

A non-converged network is recommended, which must be configured at the host OS level as well. The following figure depicts the Switch Embedded Team (SET) configuration needed for the management and storage traffic on the physical adapters.

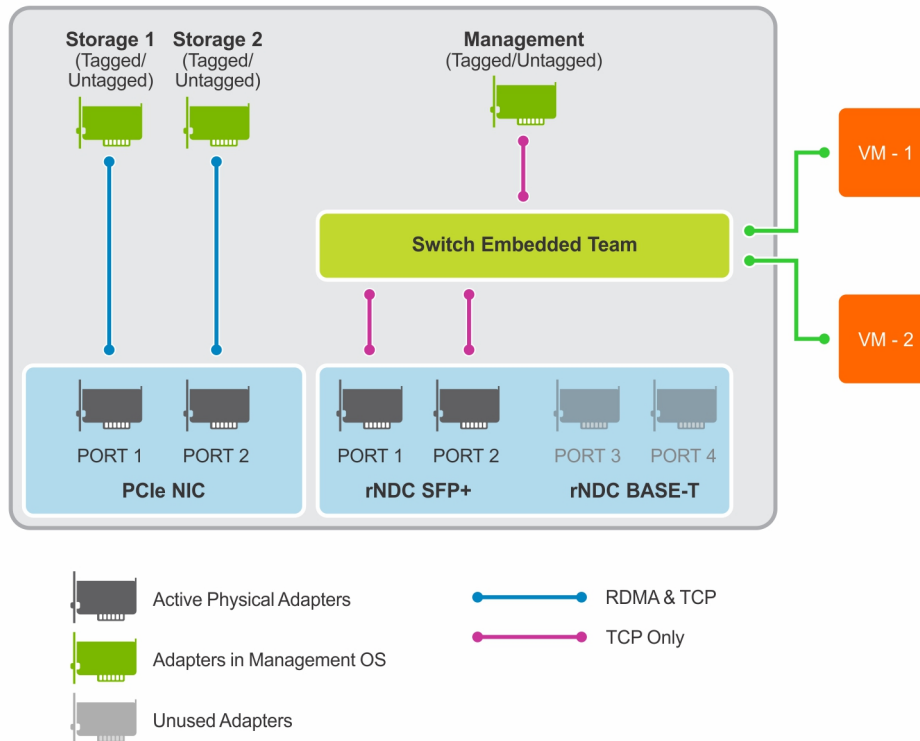


Figure 3. Host networking configuration

In this configuration, one VM switch (Switch Embedded Teaming) is implemented for the management traffic connected to the SFP+ ports of the Intel rNDC and the storage traffic uses the 2 x 25 GbE ports from the QLogic 41262 network adapter.

- ① **NOTE: All PowerShell commands in this section must be run at the local console to ensure that there are no failures due to network disconnections during configuration.**
- ① **NOTE: The default load balancing algorithm for Windows Server 2016 SET is Dynamic whereas the same in Windows Server 2019 is Hyper-V Port.**

Dell EMC recommends a non-converged network implementation and using physical network adapters for the storage traffic rather than implementing a Switch Embedded team (SET). However, in a non-converged configuration option, if VM adapters need to have the RDMA capability, SET configuration will be necessary for storage adapters. For information on this configuration choice, see <https://community.emc.com/docs/DOC-73779>.

#Create Management/VM Network VM switch

```
New-VMSwitch -Name MgmtSwitch -AllowManagementOS 0 -NetAdapterName 'NIC1','NIC2' -
MinimumBandwidthMode Weight -Verbose
```

#Add management virtual adapter and configure tagged VLAN

```
Add-VMNetworkAdapter -ManagementOS -Name 'Management' -SwitchName MgmtSwitch -Passthru |
SetVMNetworkAdapterVlan -Access -VlanId 102 -Verbose
```

Configure Tagged VLAN on the physical adapters used for storage traffic.

```
Set-NetAdapter -Name 'SLOT 1 PORT 1' -VlanID 103
Set-NetAdapter -Name 'SLOT 1 PORT 2' -VlanID 104
```

NOTE: The number of storage LANs can be equal to the number of storage adapters or 1.

NOTE: The preceding steps are not necessary if you plan to use untagged storage VLANs.

#Host Management Adapter

```
New-NetIPAddress -InterfaceAlias 'vEthernet (Management)' -IPAddress 172.16.102.51
DefaultGateway 172.16.102.1 -PrefixLength 29 -AddressFamily IPv4 -Verbose
```

#DNS server address

```
Set-DnsClientServerAddress -InterfaceAlias 'vEthernet (Management)' -ServerAddresses
172.16.102.202
```

#Storage 1 Adapter

```
New-NetIPAddress -InterfaceAlias 'SLOT1 PORT 1' -IPAddress 172.16.103.51 -PrefixLength 29 -
AddressFamily IPv4 -Verbose
```

#Storage 2 Adapter

```
New-NetIPAddress -InterfaceAlias 'SLOT1 PORT 2' -IPAddress 172.16.104.51 -PrefixLength 29 -
AddressFamily IPv4 -Verbose
```

In the pre-deployment configuration, the QLogic 41262 NICs get configured to use iWARP for RDMA. However, the driver in the OS defaults to RoCE v2 for RDMA. This needs to be changed by using the `Set-NetAdapterAdvancedProperty` cmdlet.

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 2' -DisplayName 'NetworkDirect Technology' -
DisplayValue 'iWarp'
```

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 2' -DisplayName 'NetworkDirect Technology' -
DisplayValue 'iWarp'
```

NOTE: For QLogic 41262 driver versions earlier than 8.37.37.0, the display name of this attribute is 'RDMA Mode'.

It is important to ensure that the physical adapters used for storage traffic have RDMA enabled. Perform the following steps to enable RDMA for Live Migration traffic:

- 1 Run the following command to enable RDMA on the storage virtual adapters in the host OS. The argument to the Name parameter can be retrieved by using the `Get-NetAdapter` cmdlet:

```
Enable-NetAdapterRDMA -Name 'SLOT 1 PORT 1', 'SLOT 1 PORT 2'
```

- 2 Verify the RDMA configuration for the storage adapters by using the following command:

```
Get-SmbClientNetworkInterface | Where-Object { $_.FriendlyName -Like '*SLOT 1 PORT*' }
```

- 3 Run the following command to enable RDMA for Live Migration traffic:

```
Set-VMHost -VirtualMachineMigrationPerformanceOption SMB
```

Update page file settings

To be able to capture the active memory dump when a Blue Screen Of Death (BSOD) occurs, sufficient size should be allocated for the page file. Dell EMC recommends that this should be at least 40 GB + size of CSV block cache.

The Cluster CSV block cache size value can be determined by using the following command:

```
$blockCacheMB = (Get-Cluster).BlockCacheSize
```

NOTE: On Windows Server 2016, the default block cache size is 0. On Windows Server 2019, the block cache is set to a value 1024 (1 GB).

To update the page file settings, run the following command:

```
Set-WmiInstance -Class Win32_PageFileSetting -Arguments @{
    Name = 'C:\Pagefile.sys'
    InitialSize = [Math]::Round(40960 + $blockCacheMB)
    MaximumSize = [Math]::Round(40960 + $blockCacheMB)
} -Verbose
```

Enabling jumbo frames

Enabling jumbo frames on the storage network could improve the overall read/write performance of the Storage Spaces Direct cluster. To take full advantage of jumbo frames technology, the entire data path must be set to use jumbo frames. Jumbo frames must be configured end-to-end from virtual adapters to the switch ports where the physical server ports are connected. The lowest value within the data path determines the maximum frame size used for that path.

Run the following PowerShell commands to enable jumbo frames:

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 1' -RegistryKeyword "*JumboPacket" -
Registryvalue 9014

Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 2' -RegistryKeyword "*JumboPacket" -
Registryvalue 9014
```

To set jumbo frame on the switch, run the following commands on the switch:

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/1
OS10(config-if-eth1/1/1)# mtu 9216
OS10(config-if-eth1/1/1)# show configuration
!
interface ethernet1/1/1
 no shutdown
 switchport access vlan 1
 mtu 9216
OS10(config-if-eth1/1/1)
```

Domain join

The cluster nodes must be a part of an Active Directory domain before the cluster creation actions can be performed. This task can be performed by using the *Add-Computer* cmdlet.

```
$credential = Get-Credential
Add-Computer -DomainName S2dlab.local -Credential $credential -Restart
```

- ① **NOTE:** This command induces an automatic restart at the end of domain join operation. This command needs to be run on each host that will be a part of the Storage Spaces Direct cluster.
- ① **NOTE:** Optionally, you may want to add all newly created computer objects from the HCI cluster deployment to a different Organizational Unit (OU) in the AD directory Services. In this case, *-OUPath* parameter can be used along with the *Add-Computer* cmdlet.

Validate cluster creation readiness

Before creating a host cluster, Dell EMC recommends ensuring that the nodes that would be a part of the cluster are configured as needed and are ready for the cluster creation. This can be done using the *Test-Cluster* cmdlet.

The commands in this section need to be executed on only one node in the infrastructure.

Before creating the host cluster, execute the `Get-PhysicalDisk` command on all cluster nodes and verify the output to ensure that all disks are in healthy state, and there are equal number of disks per node.

Validate that the nodes have homogeneous hardware configuration.

```
Test-Cluster -Node WSSDQS01, WSSDQS02, WSSDQS03, WSSDQS04 -Include 'Storage Spaces Direct',  
'Inventory', 'Network', 'System Configuration'
```

The `Test-Cluster` cmdlet generates an HTML report of all validations performed and includes a summary of the validation. Review this report before creating a cluster.

Create cluster and enable Storage Spaces Direct

Cluster creation can be done using the `New-Cluster` cmdlet.

```
New-Cluster -Name WSSDQSCluster -Node WSSDQS01, WSSDQS02, WSSDQS03, WSSDQS04 -StaticAddress  
172.16.102.55 -NoStorage -IgnoreNetwork 172.16.103.0/29, 172.16.104.0/29 -Verbose
```

In the above command, the `-StaticAddress` parameter is used to specify an IP address for the cluster in the same IP subnet as the host management network. The `-NoStorage` switch parameter specifies that the cluster needs to be created without any shared storage.

The `New-Cluster` cmdlet generates an HTML report of all configurations performed and includes a summary of the validation. Review this report before enabling Storage Spaces Direct.

After the cluster creation is complete, the `Enable-ClusterS2D` cmdlet can be used to create the storage spaces direct configuration on the cluster.

This command should not be run in a remote session. Use the local console session for executing this cmdlet.

```
Enable-ClusterS2D -Verbose
```

The `Enable-ClusterS2D` cmdlet generates an HTML report of all configurations performed and includes a summary of the validation. Review this report. This report is typically stored in the local temporary folder on the node where the `Enable-ClusterS2D` cmdlet was run. The path to the cluster report gets shown in the verbose output of the command.

This cmdlet, at the end of the operation, discovers and claims all the available disks into an auto-created storage pool.

The cluster creation can be verified by using any of the following commands:

```
Get-ClusterS2D  
Get-StoragePool  
Get-StorageSubSystem -FriendlyName *Cluster* | Get-StorageHealthReport
```

Post-cluster creation configuration

The following steps must be performed to optimize cluster operations.

Topics:

- [Remove host management network from Live Migration](#)
- [Update hardware timeout for Spaces port](#)
- [Configuring cluster witness](#)
- [Additional resources](#)

Remove host management network from Live Migration

After the cluster creation is complete, Live Migration, by default, is configured to use all available networks. It is recommended to disable Live migration on host management network, or in other words, exclude host management network from Live Migration settings.

This can be done by using the following PowerShell commands:

```
$clusterResourceType = Get-ClusterResourceType -Name 'Virtual Machine'

$hostNetworkID = Get-ClusterNetwork | Where-Object { $_.Address -eq '172.16.102.0' } | Select -ExpandProperty ID

Set-ClusterParameter -InputObject $clusterResourceType -Name MigrationExcludeNetworks -Value $hostNetworkID
```

In the above command, 172.16.102.0 represents the host management subnet.

Update hardware timeout for Spaces port

Run the following PowerShell commands on every node in the cluster to update the configuration in the Windows registry:

```
Set-ItemProperty -Path HKLM:\SYSTEM\CurrentControlSet\Services\spaceport\Parameters -Name HwTimeout -Value 0x00002710 -Verbose

Restart-Computer -Force
```

NOTE: This command induces a reboot of the node at the end of the registry update. This update should be done on all Storage Spaces Direct Nodes being deployed using this Guide immediately after initial deployment. Update one node at a time and wait until each node rejoins the cluster.

Configuring cluster witness

Cluster Witness is an optional configuration step for a 4-node cluster and helps maintain a cluster or storage quorum when there is a node or network communication failure where nodes continue to operate but can no longer communicate between one another. Cluster witness can either be a file share or a cloud-based witness.

For information about configuring a file share witness, see <https://techcommunity.microsoft.com/t5/Failover-Clustering/New-File-Share-Witness-Feature-in-Windows-Server-2019/ba-p/372149>.

For information about configuring a cloud-based witness, see <https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness>.

After the host cluster creation and enabling Storage Spaces Direct steps are complete, Dell EMC recommends that you perform the following steps:

- 1 Post deployment verification—This is recommended to ensure that the infrastructure is functional and ready for operations.
- 2 OS license activation—By default, the OS is installed in the evaluation mode. This needs to be activated immediately after the OS installation.

The OS license activation step is not required if the OS is factory installed.

For more information about these steps, see the Operations Guide at <https://www.dell.com/azurestackhcimanuals>.

Additional resources

[Dell EMC Solutions for Microsoft Azure Stack HCI Overview](#)

[Support Matrix for Dell EMC Solutions for Azure Stack HCI](#)

[Firmware and Driver Update Catalog for Dell EMC Solutions for Azure Stack HCI](#)