

Dell EMC Solutions for Microsoft Azure Stack HCI

Deployment Guide for scalable hyper-converged infrastructure with PowerEdge R440, R740xd, R740xd2, and R640 Storage Spaces Direct Ready Nodes

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.....	5
Assumptions.....	5
Known issues.....	5
2 Virtualization infrastructure with Dell EMC Solutions for Azure Stack HCI.....	6
R740xd Storage Spaces Direct Ready Node.....	8
R740xd2 Storage Spaces Direct Ready Node.....	9
R640 Storage Spaces Direct Ready Node.....	10
PowerEdge R440.....	11
3 Integrating solution components.....	13
Network connectivity in Dell EMC Solutions for Azure Stack HCI.....	14
Port mapping.....	17
4 Virtualized infrastructure deployment by using Dell EMC Ready Nodes.....	20
Deployment prerequisites.....	20
Software versions.....	21
Dell EMC validated firmware matrix.....	21
Deployment checklist.....	21
Management environment checklist.....	21
Network configuration checklist.....	22
Host OS network checklist.....	23
Pre-deployment configuration.....	24
Network switch configuration.....	24
iDRAC and BIOS configuration.....	25
QLogic NIC configuration.....	27
Firmware baselining.....	31
Hyper-converged infrastructure deployment.....	31
OS deployment.....	31
Install roles and features.....	32
Update Out of Box drivers.....	32
Change host name.....	33
Configure firewall.....	33
Cluster deployment and configuration.....	34
RDMA configuration.....	39
QoS Policy configuration.....	40
Remove host management network from Live Migration.....	41
Update hardware timeout for Spaces port.....	41
Update page file settings.....	42
Enabling jumbo frames.....	42
Configuring cluster witness.....	42
Recommended next steps.....	43

5 Deployment services.....	44
6 Additional resources.....	45
A Firewall port requirements.....	46
B Sample deployment checklists.....	48

Audience and scope

This deployment guide focuses on deploying a Dell EMC Solutions for Azure Stack HCI. 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 using the validated and certified Dell EMC Storage Spaces Direct Ready Nodes and Windows Server 2016 and Windows Server 2019.

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 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 ensure accuracy, speed, reduced risk and downtime.
- Customer led: Customers can follow 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 Azure Stack HCI cluster creation and configuration tasks.

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
- Dell EMC Networking switches and concepts such as Data Center Bridging (DCB) and Virtual Link Trunking (VLT)
- Deploying and configuring Windows Server 2016 or 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>.

Virtualization infrastructure with Dell EMC Solutions for Azure Stack HCI

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 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.

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 on available configurations, see [Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview](#).

NOTE: For the two-node cluster deployment, it is mandatory that a cluster witness is configured. See the section on configuring cluster witness for available options and other references to deployment instructions.

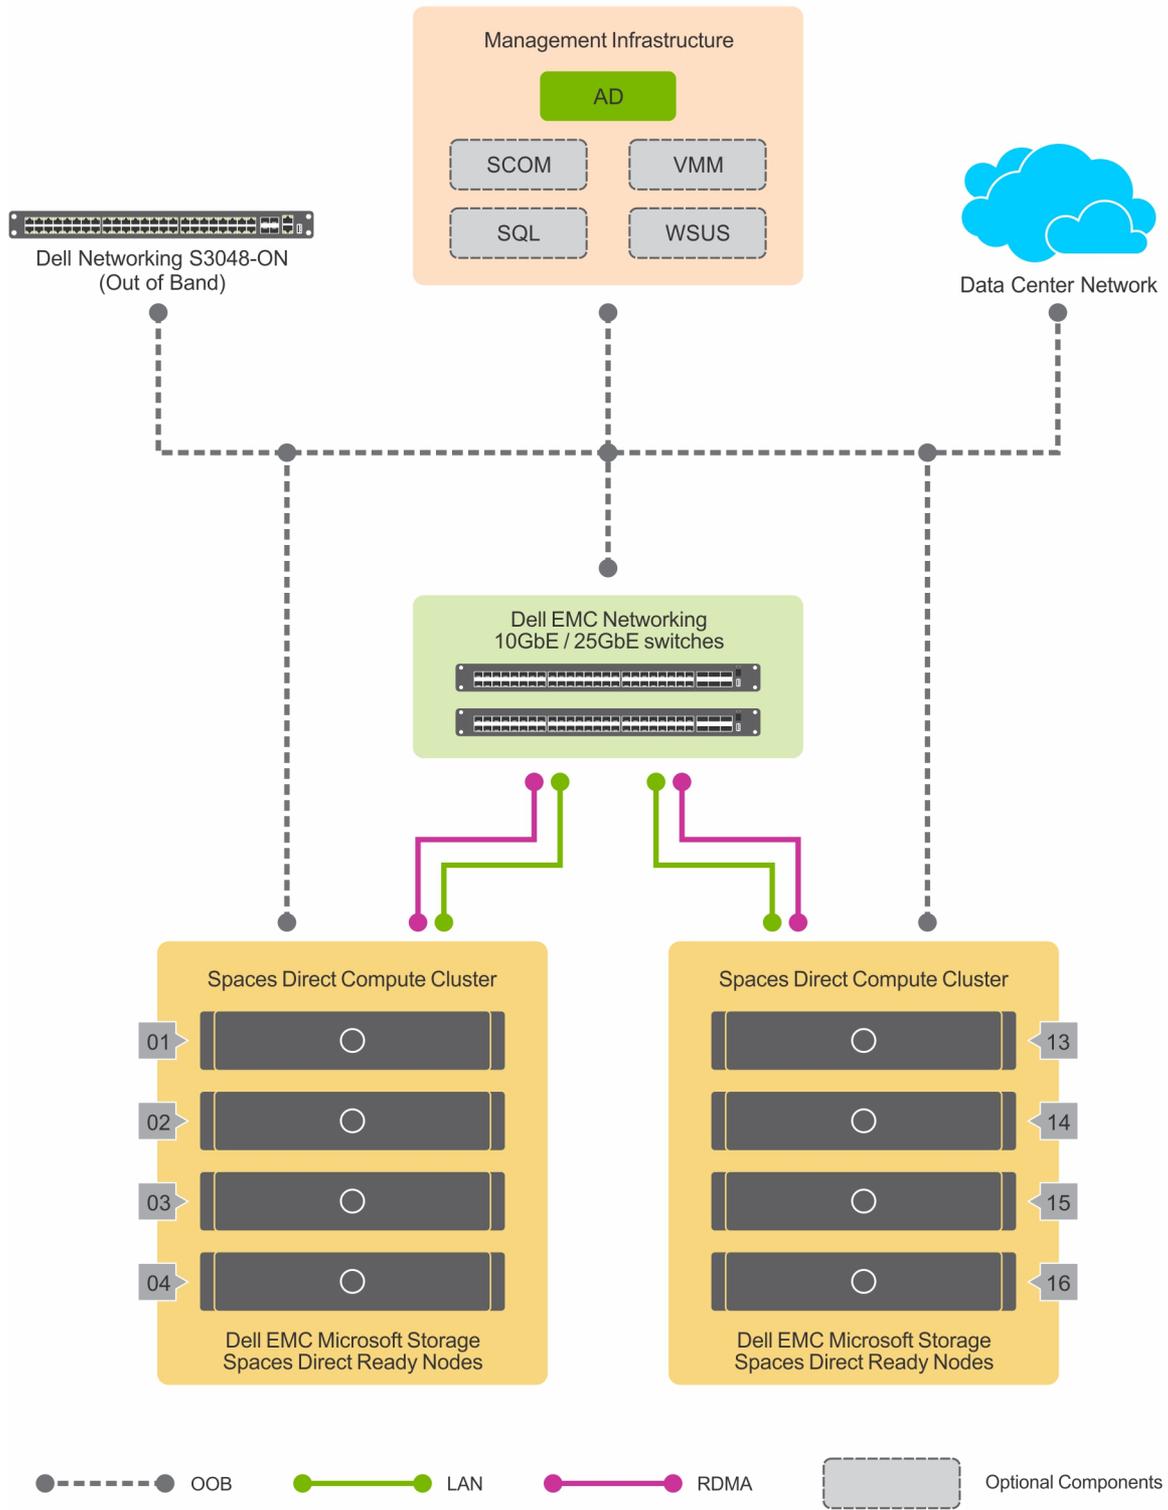


Figure 1. Hyper-converged virtualized solution using Dell EMC Ready Nodes

Dell EMC Solutions for Azure Stack HCI do not include management infrastructure components such as a cluster for hosting management VMs and services such as Active Directory (AD), Domain Name Service (DNS), Windows Server Update Services (WSUS), and System Center components such as Virtual Machine Manager (VMM) and Operations Manager (OM). Therefore, the instructions in this guide do

not include deployment of any of these services and/or components and assume that at least an Active Directory domain controller and a DNS server are available in the existing management infrastructure.

This deployment guidance provides instructions and PowerShell commands for deploying an Azure Stack HCI cluster manually. 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>.

The subsequent sections provide an overview of the hardware and software components in the virtualized solution based on Dell EMC Ready Nodes.

Topics:

- [R740xd Storage Spaces Direct Ready Node](#)
- [R740xd2 Storage Spaces Direct Ready Node](#)
- [R640 Storage Spaces Direct Ready Node](#)
- [PowerEdge R440](#)

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 lists the configuration options that are available.

Table 1. System configuration

Configuration	Chassis	Drive layout
Hybrid	R740xd Storage Spaces Direct Ready Node (18-drives) - 3.5" drive form factor	Up to 6 x SSD in the front bay Up to 12 x HDD (6 in the front bay, 4 in the internal bay, and 2 in the rear flex bay)
	R740xd Storage Spaces Direct Ready Node (12-drives) - 3.5" drive form factor	2-4 x SSD in the front bay (cache) 4-8 x HDD in the front bay (capacity)
All-Flash	R740xd Storage Spaces Direct Ready Node (24-drive) - 2.5" drive form factor	4 -24 x Mixed-use or Read-intensive SSD
All-NVMe	R740xd Storage Spaces Direct Ready Node (24 drive with 12 x NVMe capable)	4-12 x PCIe Mixed-use NVMe SSD
NVMe Add-In-Card (AIC) + HDD	R740xd Storage Spaces Direct Ready Node (12-drives) - 3.5" drive form factor with two NVMe AIC	2 x NVMe PCIe Add In Card (AIC) in PCIe Slot 1,7 (cache) 12 x HDD in the front bay (capacity)
NVMe (Small Form Factor) SFF + HDD	R740xd Storage Spaces Direct Ready Node (24-drives) - 2.5" drive form factor with up to 4 NVMe SFF	2 or 4 x 2.5" PCIe NVMe SFF drives in the front bay (cache) Up to 20 x 2.5" HDD in the front bay (capacity)

Configuration	Chassis	Drive layout
NVMe SFF + SSD	R740xd Storage Spaces Direct Ready Node (24-drives) - 2.5" drive form factor with up to 4 NVMe SFF	2 or 4 x 2.5" PCIe NVMe SFF drives in the front bay (cache) Up to 20 x 2.5" SSD in the front bay (capacity)

NOTE: The NVMe AIC are not pre-installed in the Ready Node. These cards must be installed before proceeding to the component integration section. See the slot numbers mentioned in the preceding table for slot priority of the NVMe AIC.

The following table provides an overview of the 740xd Storage Spaces Direct Ready Nodes in the solution.

Table 2. System components

Component	Specification
NIC	Intel Ethernet 10G 4P X710/I350 rNDC and one of the following: <ul style="list-style-type: none"> up to two Mellanox ConnectX-4 LX 25 GbE SFP add-in adapter card up to two QLogic FastLinQ 41262 Dual Port 25 GbE SFP28 add-in adapter card
Storage adapter	HBA 330
Boot device	BOSS S.1 with 2 x BOSS M.2 devices in RAID 1
Drives	See Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview
LAN switch	(2) Dell EMC Networking 10 GbE or 25 GbE switches
OOB switch	(1) S3048

R740xd2 Storage Spaces Direct Ready Node

The R740xd2 Storage Spaces Direct Ready Node provides dense storage capacity configurations 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 guidance. This server offers a wide range of drive capacity configurations to enable up to 364 TB (using 24 x 3.5" SAS/SATA HDD) of raw storage capacity. With two processor sockets and a wide range of processor options, the R740xd2 Storage Spaces Direct Ready Node offers the capabilities that match not just the storage needs but also the computing needs.

This 2U rack-mounted server is available in hybrid drive configurations. The following table lists the configuration options that are available.

Table 3. System configuration

Item	Configuration
Chassis	12 x 3.5" front bay 1 + 12 x 3.5" front bay 2 with butterfly riser Config 4
Processor	Intel Xeon E5-3600 family
Memory	Minimum Memory: 96 GB Maximum Memory: 384 GB
Storage Controller	Internal HBA330
Storage - OS Boot	BOSS with Dual M.2 240 or 480 GB (RAID 1)
Storage - SSD (cache)	Up to 8 x SAS (WI/MU)/SATA MU SSDs (Front Bay 2 only)
Storage - HDD (capacity)	Up to 20 x 3.5" NL-SAS/SATA HDD (Front Bay 1 and 2)

Item	Configuration
Network cards	Add-in-Card: 1 or 2 Mellanox Connect-X4 LX (Dual Port 25 GbE) or 1 or 2 x QLogic QL41262 (Dual Port 25 GbE) On-Board: Broadcom 5720 Dual Port 1 GbE Mezz Card: Broadcom 57416 Dual Port 10 GbE SFP+
iDRAC	iDRAC 9

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 R640 Storage Spaces Direct Ready Node is available in different chassis configurations that offer several internal storage choices.

The following table lists the system configuration options for this Ready Node.

Table 4. System configuration

Component	Chassis	Drive layout
Hybrid	R640 Storage Spaces Direct Ready Node (12-drives) - 2.5" drive form factor	2-4 x SSD (cache) 4-10 x HDD (capacity)
	R640 Storage Spaces Direct Ready Node (10-drives) - 2.5" drive form factor	2-4 x SSD (cache) 4-8 x HDD (capacity)
All-Flash	R640 Storage Spaces Direct Ready Node (12-drives) - 2.5" drive form factor	12 x SSD
	R640 Storage Spaces Direct Ready Node (10-drives) - 2.5" drive form factor	10 x SSD

NOTE: The 12-drive chassis configuration in R640 Storage Spaces Direct Ready Node is available only with one PCIe slot and that is used by the BOSS M.2 device for OS RAID. In this chassis configuration, network connectivity is provided by using a Mellanox ConnectX 4 LX or QLogic FastLinQ 41262 rNDC.

The following table provides an overview of the R640 Storage Spaces Ready Node in the solution.

Table 5. System components

Component	Specification
NIC	For 12-drive chassis: Mellanox ConnectX 4 LX or QLogic FastLinQ 41262 rNDC For 10-drive chassis: Intel Ethernet 10G 4P X710/I350 rNDC and one of the following: <ul style="list-style-type: none"> up to two Mellanox Connectx-4 LX 25 GbE SFP add-in adapter card

Component	Specification
	<ul style="list-style-type: none"> up to two QLogic FastLinQ 41262 25 GbE SFP28 add-in adapter card
Storage adapter	HBA 330
Boot device	BOSS S.1 with 2 x BOSS M.2 devices in RAID 1
Drives	See Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview
LAN switch	(2) Dell EMC Networking 10 GbE or 25 GbE switches
OOB switch	(1) S3048

PowerEdge R440

Dell EMC PowerEdge R440 is a 1U, two 2-socket rack server that supports up to 1 TB of physical memory. The PowerEdge R440 server with the 10 drive chassis configuration offers multiple hybrid and all-flash storage options for Microsoft Storage Spaces Direct based hyper-converged infrastructure.

The following table lists the system configuration options for PowerEdge R440.

Table 6. System configuration

Component	Chassis	Drive layout
Hybrid	PowerEdge R440 (10 drives) - 2.5" form factor	2 x SSD (cache) 4-8 x HDD (capacity)
All-Flash	PowerEdge R440 (10 drives) - 2.5" form factor	4-10 x SSD (capacity)

The following table provides an overview of the R440 Storage Spaces Ready Node in the solution.

Table 7. System components

Component	Specification
NIC	On-Board Broadcom 5720 Dual Port 1 GB LOM and Broadcom 57416 Dual Port 10 GbE SFP+ Network LOM Mezz Card and the following: (1) QLogic FastLinQ 41262 (Dual Port 25 GbE)
Storage adapter	HBA 330
Boot device	BOSS S.1 with 2 x BOSS M.2 devices in RAID 1
Drives	See Dell EMC Ready Nodes for Microsoft Storage Spaces Direct with Hyper-V Solution Overview
LAN switch	(2) Dell EMC Networking 10 GbE or 25 GbE switches
OOB switch	(1) S3048

NOTE: PowerEdge R440 as a cluster node in Storage Spaces Direct infrastructure is validated with QLogic 41262 DP 25 GbE adapter only for converged or non-converged traffic.

The redundant switch configuration shown in Figure 1 provides high availability (HA). If you are connecting the nodes to two separate network switches and implementing switch embedded teaming (SET), both switches require access to all subnets so that failover can occur. When using RoCE for RDMA, DCB must be configured as it provides enhancements to the Ethernet protocol, which improves the

functionality of data center networks. To take advantage of the Mellanox RDMA over Converged Ethernet (RoCE) network adapters, Priority Flow Control (PFC), and Enhanced Transmission Selection (ETS) are required. PFC and ETS are configured on all nodes and all network switches interconnecting the nodes.

Integrating solution components

This section provides recommendations on server and network switch placement in the racks and port mapping on the TOR and OOB switches. The section on Network Architecture and Configuration provides details on configuring the TOR and OOB switches.

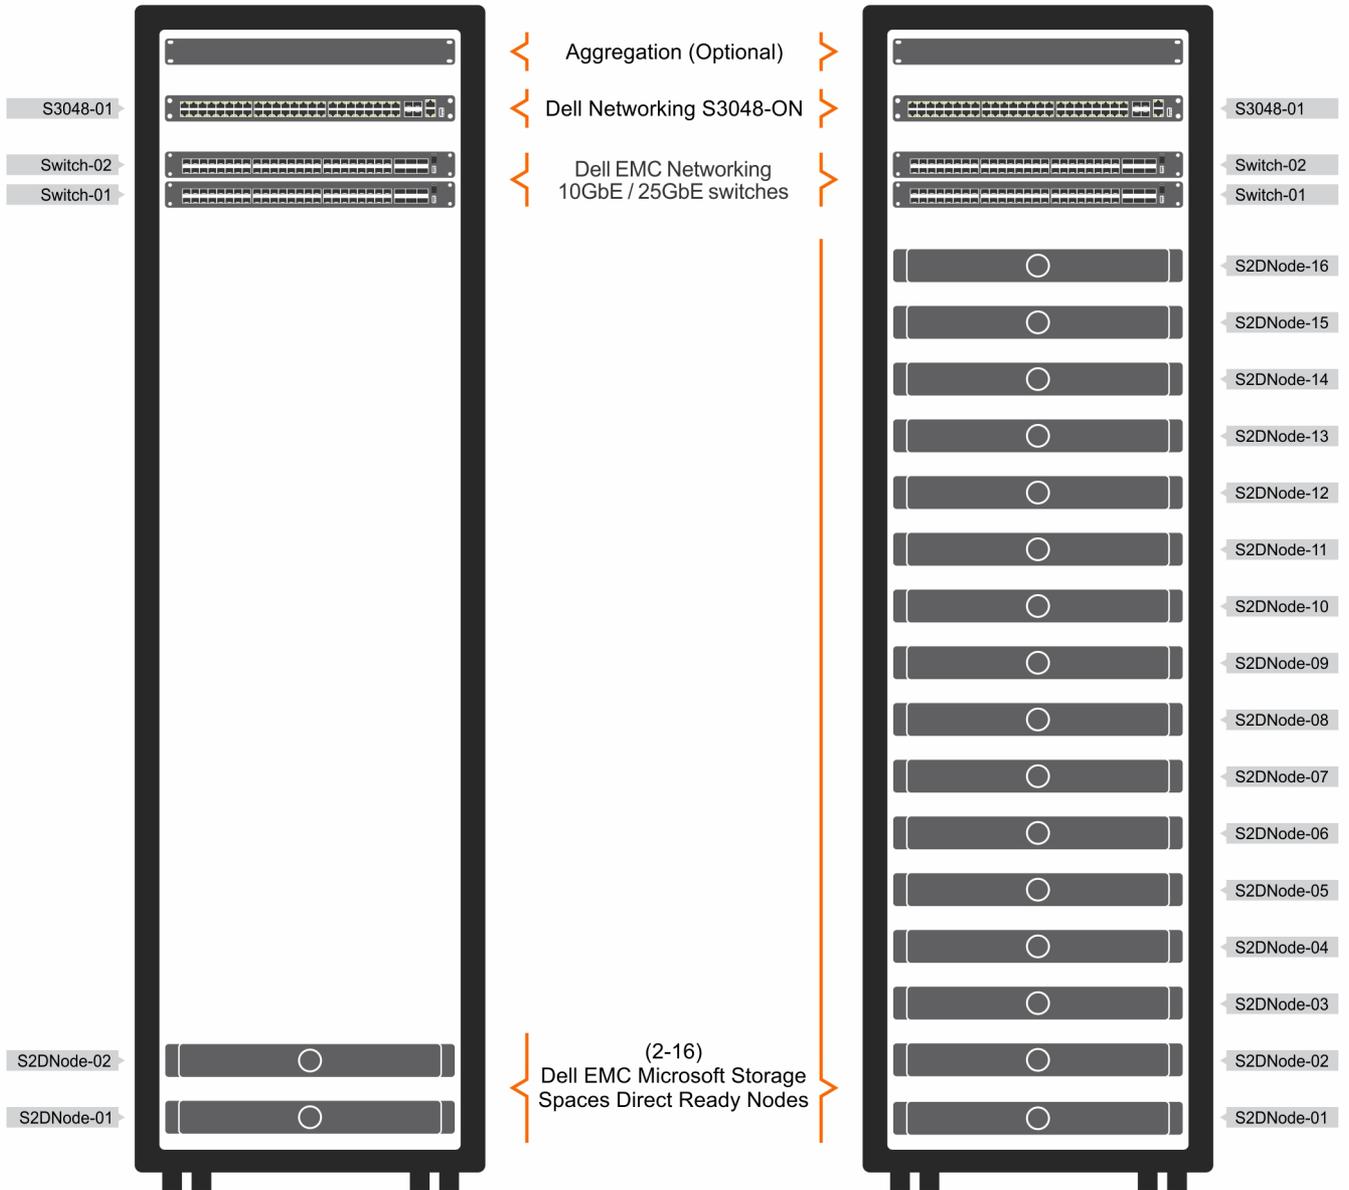


Figure 2. Rack elevation for the 2 node and 16 node HCI configurations

- ① **NOTE:** Dell EMC does not support expansion to a larger cluster size from a 2-node cluster. A 3-node cluster provides fault-tolerance only for simultaneous failure of a single node and a single drive. If the deployment requires future expansion and better fault-tolerance, consider starting with a 4-node cluster at a minimum.
- ① **NOTE:** The 2 node configuration built on Storage Spaces Direct Ready Node has been validated with back-to-back connections for storage connectivity. For more information about the back-to-back connected deployment model, see <https://www.dell.com/azurestackhcimanuals>.

Topics:

- [Network connectivity in Dell EMC Solutions for Azure Stack HCI](#)
- [Port mapping](#)

Network connectivity in Dell EMC Solutions for Azure Stack HCI

There are two physical network topologies used with the Azure Stack HCI Solution—fully-converged and non-converged. Both topologies provide either two 10 GbE/40 GbE or two 25 GbE/100 GbE top-of-rack (TOR VLT) switches, and a single 1 GbE out-of-band (OOB) management switch.

Fully converged

In this configuration option, each Storage Spaces Direct Ready Node (with the exception of R640 Storage Spaces Direct Ready Node with 12-drive chassis configuration which supports only one network adapter) uses up to two Mellanox ConnectX-4 LX or QLogic FastLinQ 41262 25 GbE network adapters for both host management and storage traffic. Each port from each network adapter on each server is connected to a different switch in the redundant top-of-rack (TOR) network fabric.

The following figure illustrates the connection.

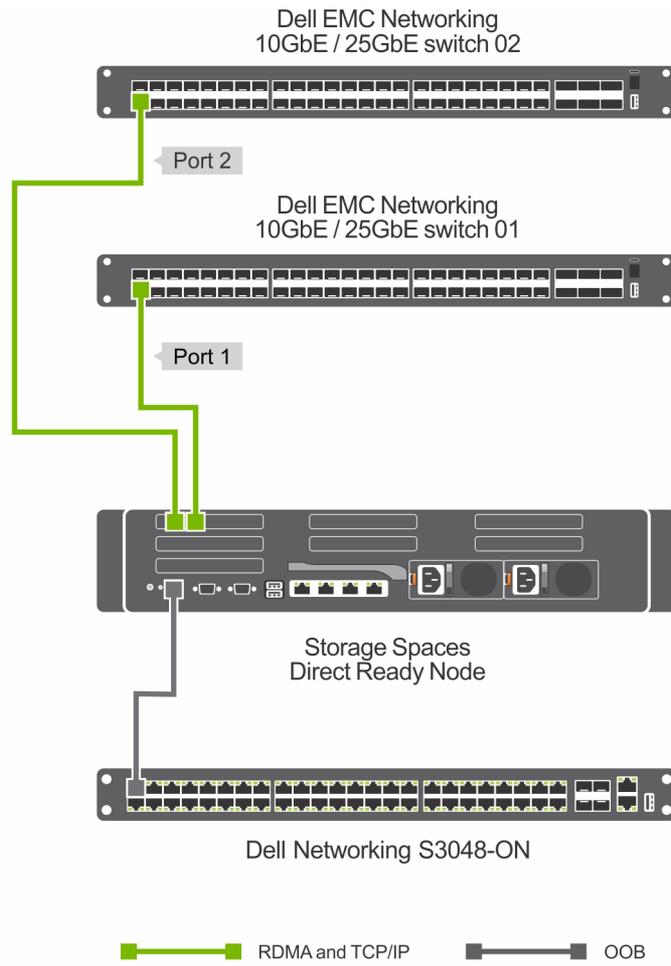


Figure 3. Network and iDRAC connection in a fully-converged network configuration

Non-converged network

In this second network connectivity option, each Storage Spaces Direct Ready Node (with the exception of R640 Storage Spaces Direct Ready Node with 12-drive chassis configuration which supports only one network adapter) uses up to two Mellanox ConnectX-4 LX or QLogic FastLinQ 41262 25 GbE network adapters for storage traffic and the 10 GbE SFP+ or the 1 GbE BASE-T ports on the on-board Intel Ethernet 10G 4P X710/I350, or Broadcom 57412 DP 10 Gb SFP+ + 5720 DP 1 Gb, or QLogic 57800 DP 10 Gb SFP+ DP 1 Gb BASE-T rNDC for host management and VM traffic.

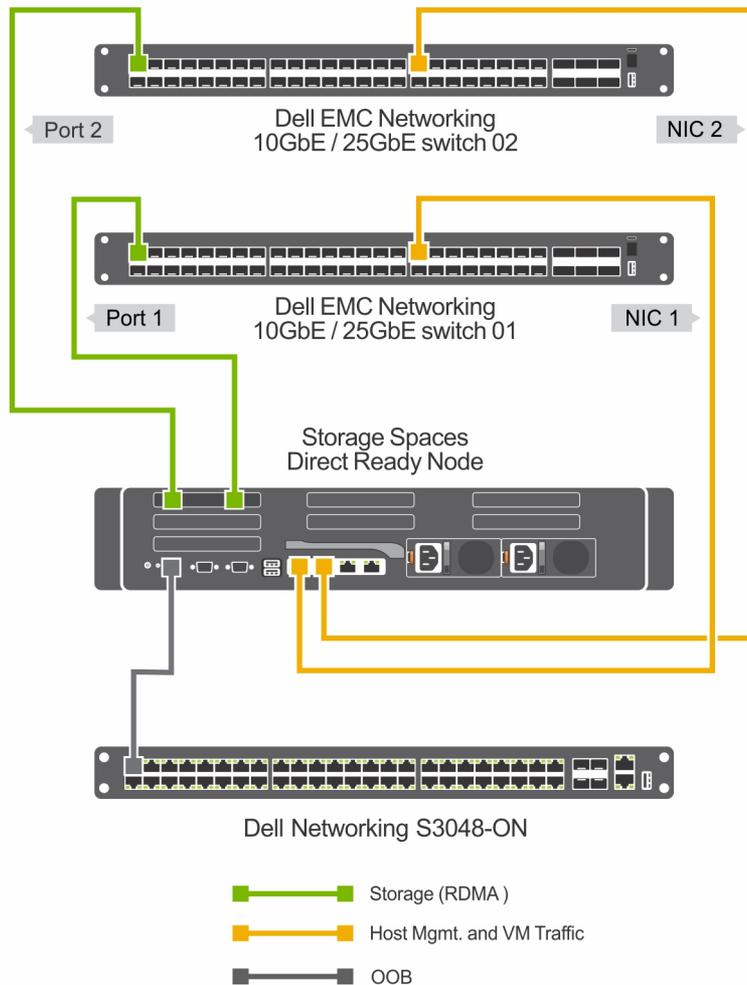


Figure 4. Network and iDRAC connection in a non-converged network configuration

The network connectivity that uses Mellanox ConnectX-4 LX and QLogic FastLinQ 41262 network adapters in the server provides Remote Direct Memory Access (RDMA) for storage traffic. RDMA enables significantly increased throughput and lower latency by performing direct memory transfers between servers. Storage Spaces Direct uses SMB for all intra-node communication and uses RDMA with SMB Direct to enhance performance of the overall infrastructure. RDMA with RoCE configuration on Mellanox ConnectX-4 LX network adapters requires Data Center Bridging (DCB) and Quality of Service (QoS) configured in the host operating system as well as the top of rack network fabric switch in a fully-converged topology.

For All-NVMe configurations that implement iWARP for RDMA using the QLogic FastLinQ 41262 network adapters, implementing DCB in the host OS and TOR network switches is recommended in the fully-converged network topology.

The redundant switch configuration shown in Figure 1 provides high availability (HA). If you are connecting the nodes to two separate network switches and implementing switch embedded teaming (SET), both switches require access to all subnets so that failover can occur. When using RoCE for RDMA, DCB must be configured as it provides enhancements to the Ethernet protocol, which improves the functionality of data center networks. To take advantage of the Mellanox RDMA over Converged Ethernet (RoCE) network adapters, Priority Flow Control (PFC), and Enhanced Transmission Selection (ETS) are required. PFC and ETS are configured on all nodes and all network switches interconnecting the nodes.

When deploying the fully-converged topology, the Node switch ports must be in Trunk Mode and configured for all Management, and Storage Networks VLANs throughout both VLT switches.

In the non-converged topology, the management network can be connected to the same TOR switches as the storage traffic or it can be connected to a completely different network fabric.

In the dual-NIC configuration with either Mellanox or QLogic adapters, for increased throughput, all four network ports can be connected for storage traffic. This may be necessary for all-flash or all-NVMe configurations.

Each Storage Spaces Direct Ready Node has a dedicated network connection from the Integrated Dell Remote Access Controller (iDRAC) to the Dell Networking S3048-ON switch configured for OOB management.

NOTE: The R640 Storage Spaces Direct Ready Node in a 12-drive chassis configuration can be integrated using the fully converged network topology only.

When using the non-converged network topology, the host network can be configured in different ways. For detailed information about the host network configuration options, see <https://community.emc.com/docs/DOC-73779>. For Storage Spaces Direct, Dell EMC recommends a non-converged network topology with storage traffic on the physical adapters and the host management/virtual traffic implemented as virtual adapters connected to a Switch Embedded Team created using the onboard 10 GbE or 1 GbE ports.

NOTE: In the non-converged network topology, it is not necessary to configure DCB/PFC in the top-of-rack switches or QoS in the host OS.

NOTE: Both the topologies support tagged and untagged VLANs for storage traffic.

NOTE: For a list of all supported network topology configurations, see <https://community.emc.com/docs/DOC-73779>.

NOTE: For sample configurations of the TOR switches with DCB and other configurations, see <https://community.emc.com/docs/DOC-70310>

NOTE: For a list of supported switches within Dell EMC Solutions for Azure Stack HCI, see the support matrix available at <https://www.dell.com/azurestackhcimanuals>.

Port mapping

This section provides recommendations to switch port mapping on a Storage Spaces Direct Ready Node server. The illustrations provided in the section describe how the server NIC ports for both storage and out of band management traffic are connected to the TOR switches.

To make network configuration consistent and easy to understand, it is recommended to follow a consistent mapping across the TOR and OOB switch. The following figures illustrate this mapping.

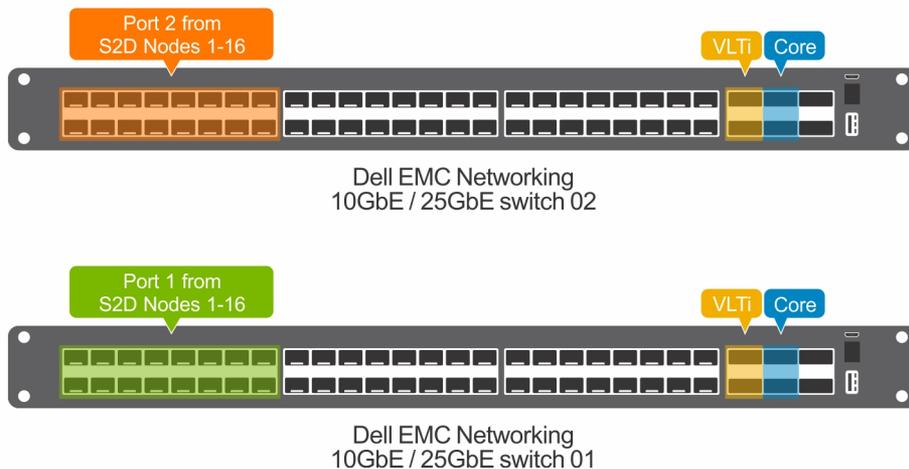


Figure 5. Fully-converged port mapping for Ready Nodes

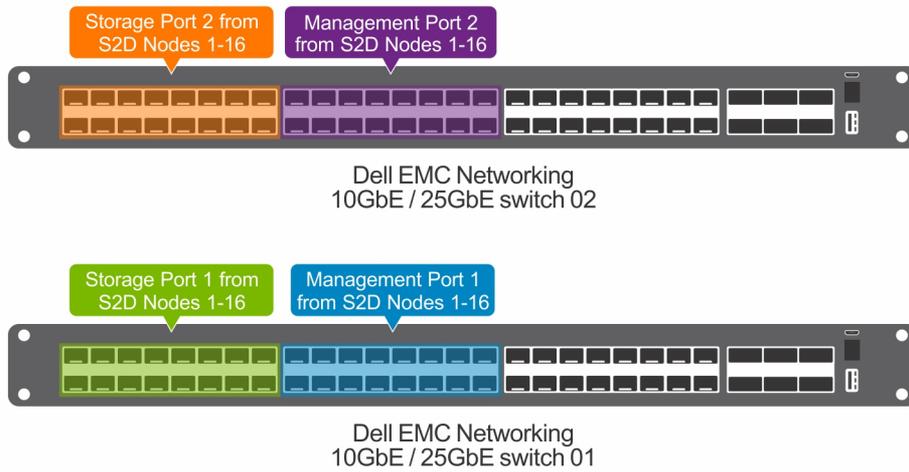


Figure 6. Non-converged TOR port mapping for Ready Nodes

As shown in the figure, ports from the Mellanox and QLogic adapters from each server are mapped to ports 1-16 on TOR1 and TOR2 respectively. The 40 GbE ports, port 49 and 50 from each TOR are connected together in a Virtual Link Trunking Interconnect (VLT) for inter-switch connectivity.

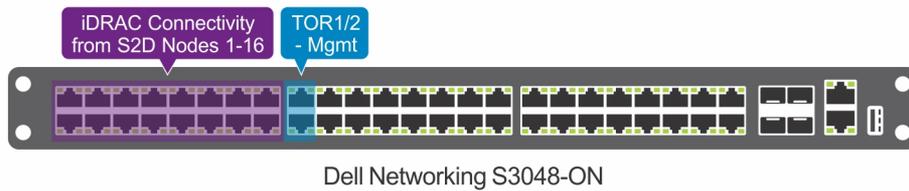


Figure 7. OOB port mapping for Ready Nodes

In another optional deployment scenario for the OOB and TOR network fabric, the OOB switch can be uplinked to the TOR switches by using the 10 GbE ports. Doing so enables connectivity to the OOB interfaces of the cluster components through the TOR switch connections. The following figure illustrates this.

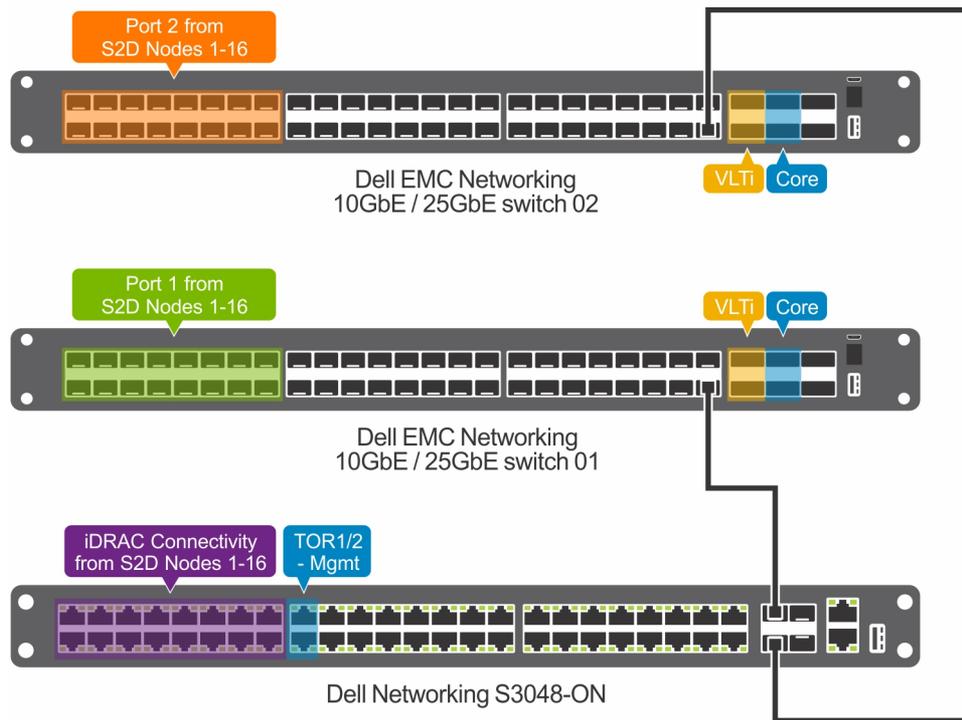


Figure 8. Fully-converged network connectivity with Dell Networking S3048 (OOB switch) connecting to Dell Networking S4048 10 GbE/25 GbE switches

Virtualized infrastructure deployment by using Dell EMC Ready Nodes

Dell EMC Microsoft Storage Spaces Direct Ready Nodes described in this guide can be deployed in two different ways.

- Manual OS deployment—A manual method of installation starting from OS deployment to cluster creation.
- Factory OS deployment—R740xd, R740xd2, and R640 Storage Spaces Direct Ready Nodes and the PowerEdge R440 servers can be ordered from the factory with Windows Server 2016 or Windows Server 2019 pre-installed.

Each of the above mentioned methods have certain deployment prerequisites and pre-deployment configuration that need to be performed including the network switch configuration.

The subsequent sections of this guide describe the deployment prerequisites for each of these methods and provide details on the supported [software and firmware versions](#).

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 the Windows Server 2019 GA build. 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: Each task in this deployment guide requires one or more PowerShell commands to be run. Dell EMC recommends using these commands to complete the deployment tasks instead of Graphical User Interface (GUI) as there may be scenarios where GUI may not work as expected. For example, the cluster validation GUI wizard within the Failover Cluster Manager fails intermittently due to a known issue in the Microsoft code.

Topics:

- [Deployment prerequisites](#)
- [Deployment checklist](#)
- [Pre-deployment configuration](#)
- [Hyper-converged infrastructure deployment](#)
- [Recommended next steps](#)

Deployment prerequisites

This hyper-converged virtualized solution based on Dell EMC Ready Nodes deployment assumes that the management services required for the OS deployment and cluster configuration are present in the existing infrastructure where the Storage Spaces Direct cluster deployment is being done.

The following table describes different management services, their purpose, and applicability in each of the deployment methods.

Table 8. Management services

Management Service	Purpose	Deployment - Required/Optional
Active Directory	User authentication	Required
Domain Name Service	Name resolution	Required

Management Service	Purpose	Deployment - Required/Optional
Windows Software Update Service (WSUS)	Provides local source for Windows Updates	Optional
MS SQL Server	Provides database backend for VMM and SCOM	Optional

Software versions

The following table lists the software versions required for the Dell EMC Ready Nodes deployment. The software requirements, like the management services, vary between the deployment methods described above.

Table 9. Software versions

Components	Version
Operating System	Windows Server 2016 Data Center or Windows Server 2019 Data Center
Active Directory Forest/domain functional level	Windows Server 2008 R2 or later

Dell EMC validated firmware matrix

This Dell EMC Ready Nodes for Storage Spaces Direct is validated and certified with certain firmware versions related to the components in the solution infrastructure. This matrix identifies the Dell EMC validated versions of software and firmware 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 gets updated when new revisions of the software and firmware are validated.

Deployment checklist

This section provides a checklist for configuration settings that need to be applied during pre-deployment configuration and deployment of operating system and other software configuration post OS deployment. For example, network switch configuration requires VLAN ID configuration and IP address space used with each VLAN.

Fill in the checklists provided in the subsequent sections before proceeding to the pre-deployment configuration.

The section [Sample deployment checklists](#) provides completed examples of these checklists for reference.

Management environment checklist

This Dell EMC Ready Nodes deployment is a brownfield deployment and therefore requires information such as active directory domain FQDN, DNS server addresses, and so on.

The following table captures the necessary inputs as a checklist.

Table 10. Management environment checklist

Item	Value
AD Domain FQDN	
Domain Administrator or equivalent credentials	
DNS Server addresses	
WSUS Server FQDN (optional)	

Network configuration checklist

Before starting the deployment, identify the IP scope and VLAN information for various traffic classes in the solution infrastructure. The *minimum IP addresses needed* column in the following table can be used to identify the correct scope. The value shown in this column is based on the number of components that require the specified traffic class used in this solution. Ensure that the IP scope selected for the traffic class meets the minimum IP addresses requirement.

The IP Scope and VLAN ID information provided in the below table are only examples, and you must choose these values based on existing data center architecture.

Consult with the customer network engineering team for VLAN ID and VLAN IP addresses applicable to your solution.

Table 11. Network configuration

Traffic Class	Purpose	Minimum IP addresses needed	VLAN ID	Tagged/ Untagged	Recommended Subnet mask	VLAN IP Addresses
Out of band	Required for OOB management of server nodes and TOR switches	19		Untagged	/27	
Host Management	Management of cluster and cluster nodes	17		Tagged/ Untagged	/27	TOR1: TOR2:
Storage 1	SMB traffic	16		Tagged/ Untagged	/27	TOR1: TOR2:
Storage 2	SMB Traffic	16		Tagged/ Untagged	/27	TOR1: TOR2:

TOR and OOB switch configuration may also require configuring settings such as hostnames, IP routes, DCB priority settings, and so on. The following table captures this as a checklist.

Table 12. Network configuration checklist

Item	Value
OOB Switch hostname	
TOR1 Switch hostname	

Item	Value
TOR2 Switch hostname	
Enable password	
Additional user/password	
IP route on OOB (if needed)	
IP route on TOR1 / TOR2 (if needed)	
DCB Bandwidth for SMB traffic	

Host OS network checklist

Dell EMC recommends to have consistent host naming and IP addressing across multiple nodes in the virtualized cluster deployment. The host OS network configuration includes naming for the virtual switches and adapters, and assigning host names and IP addresses.

The following table provides the checklist for capturing the host OS network switch and adapter details.

Table 13. Host OS network switch and adapter details

Item	Value
Management Virtual Switch (for non-converged)	
Storage Virtual Switch (for non-converged)	
Virtual Switch (for fully-converged)	
Management Adapter	
Storage 1 Adapter	
Storage 2 Adapter	

For the host OS configuration in any deployment method, static IP address assignment is recommended for all networks. The following table provides the checklist for capturing the details of the host OS network switch and adapter.

Table 14. Host OS network checklist

	Host Name	Management IP	Storage1 IP	Storage2 IP	OOB IP	OOB Host name
Node 1						
Node 2						
Node 3						
Node 4						
Node 5						
Node 6						
Node 7						
Node 8						
Node 9						
Node 10						

	Host Name	Management IP	Storage1 IP	Storage2 IP	OOB IP	OOB Host name
Node 11						
Node 12						
Node 13						
Node 14						
Node 15						
Node 16						

Pre-deployment configuration

This section describes the pre-deployment configuration that must be performed before deploying the hyper-converged virtualized solution based on Dell EMC Microsoft Storage Spaces Direct Ready Nodes.

Network switch configuration

When considering the fully-converged network topology of the Storage Spaces Direct solution, network resiliency is a critical option that is achieved from both a physical and logical standpoint. The illustrations below are examples of the Dell EMC/Microsoft network design.

NOTE: Management network redundancy is a combination of either iDRAC or OS DNS/IP resolution.

Dell EMC recommends that you deploy a network topology that supports a dual control plane while sharing a single data plane. The Dell EMC proprietary technology is referred to as Virtual Link Trunk (VLT). This technology provides network resiliency for data I/O.

When configuring switch VLT redundant technology, VRRP provides a virtual floating IP address that any node can reference as a gateway. If a switch fails, the virtual IP is transferred to a peer switch.

Virtual Router Redundancy Protocol (VRRP) is an active/standby First Hop Redundancy Protocol (FHRP). When used with VLT peers, VRRP becomes an active/active protocol. The VRRP virtual MAC address is present as local destination address in the FIB table of both the VLT peers. Using this method, the backup VRRP router forwards intercepted frames whose destination MAC address matches the VRRP virtual MAC address.

There are three basic types of network required for a standard Storage Spaces Direct deployment—OOB Management, Host Management, and Storage.

Depending on the number of network ports (2 or 4) used within the storage configuration, there will be 2 or 4 fault domains. This guidance provides all network topology examples with 2 storage fault domains.

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

Table 15. Solution Network VLANs

VLAN NETWORK TYPE	SUBNET MASK	HOST IP ADDRESSES	VLAN ID TAG/UNTAG
Out of band Management		19	Untagged
Host Management		17	Tagged or Untagged
Storage 1 (Fault Domain 1)		16	Tagged or Untagged
Storage 2 (Fault Domain 2)		16	Tagged or Untagged

iDRAC and BIOS configuration

The R740xd, R740xd2, and R640 Storage Spaces Direct Ready Nodes are pre-configured at the factory for optimized system BIOS and iDRAC settings. This eliminates the need for customers to manually configure these settings to a recommended baseline.

NOTE: PowerEdge R440 servers are not factory configured with the Microsoft Storage Spaces Direct optimized BIOS and iDRAC configuration settings. For these servers, you must update the BIOS and iDRAC configuration settings before deployment. The list of all optimized configuration settings is available at <https://www.dell.com/support/article/us/en/19/sln313842>

The integrated Dell Remote Access Controller (iDRAC) in Dell EMC Microsoft Storage Spaces Direct Ready Nodes can be configured to obtain an IP address from DHCP or can be assigned a static IP address. When the OOB network in the environment does not provide DHCP IP addresses, an IPv4 address must be statically configured on each iDRAC network interface manually. This can be done by accessing the physical server console by using KVM or other means.

Perform the following steps to configure the IPv4 address for iDRAC. This method can be used for configuring any additional BIOS settings.

- 1 Press F12 during the system boot.

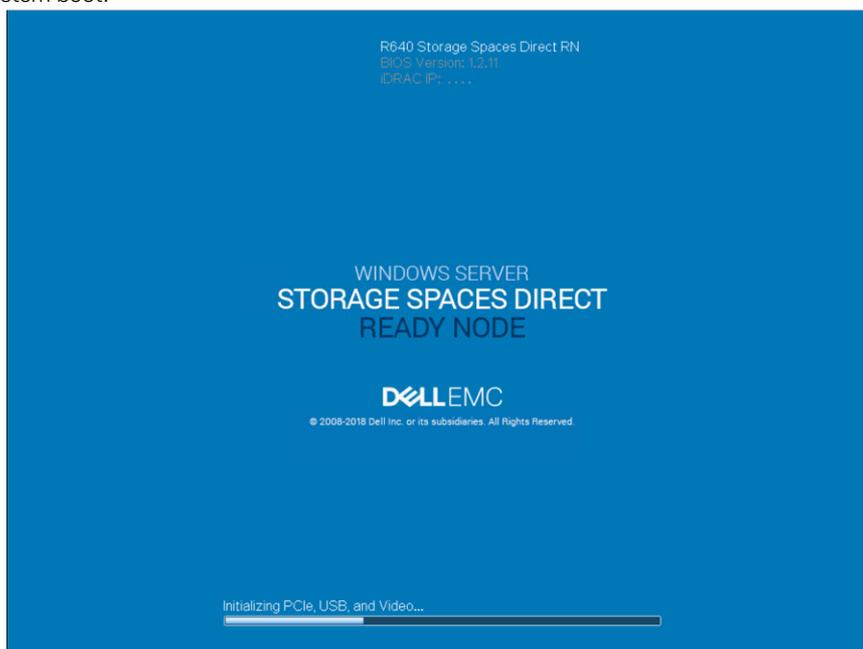


Figure 9. Enter iDRAC

NOTE: PowerEdge R440 servers do not have the customized BIOS splash screen.

- 2 Select **iDRAC Settings**.

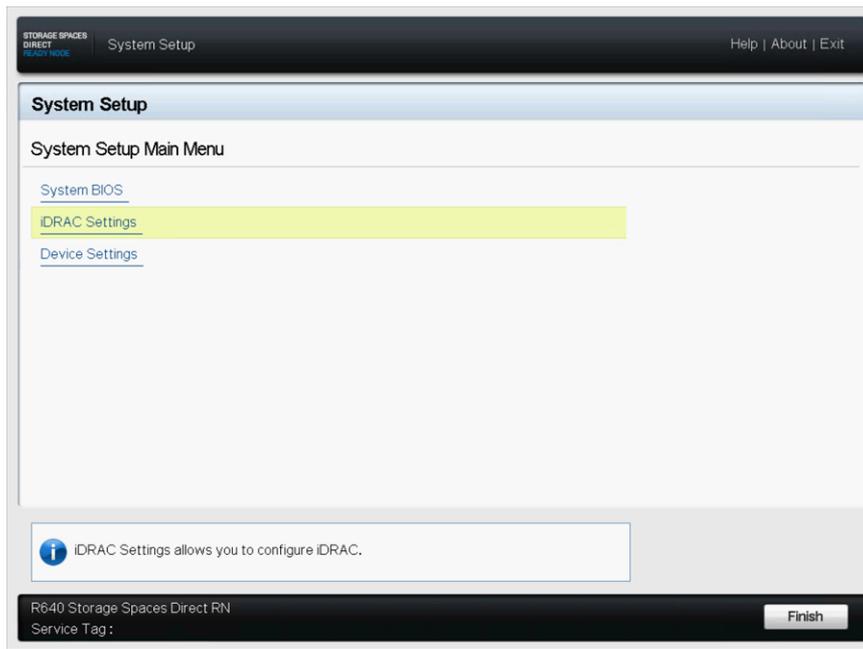


Figure 10. System Setup Main Menu

- 3 Select **Network**.

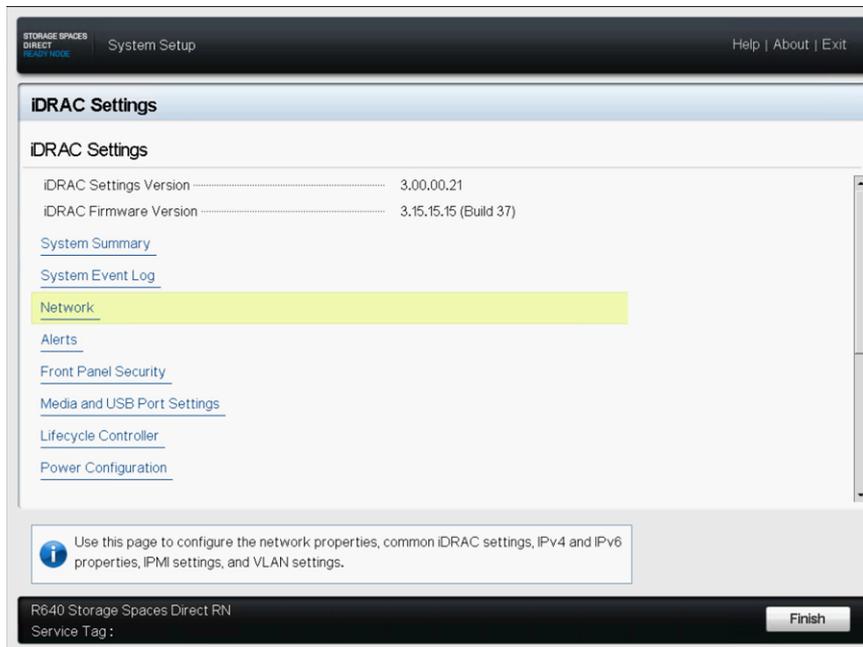


Figure 11. iDRAC Settings

- 4 In IPV4 Settings, against Enable IPv4, select **Enabled**.



Figure 12. Network Settings

- 5 Enter the static IPv4 address details.
- 6 Click **Back** and then click **Finish** to return to the System Setup page.

QLogic NIC configuration

QLogic FastLinQ 41262 network adapter supports both iWARP and RoCE for RDMA. The Storage Spaces Direct Ready Nodes are validated with only iWARP for RDMA when using the QLogic network adapters and therefore based on the network configuration chosen, you must configure the adapter manually to enable iWARP for RDMA.

Perform the following steps for each port to configure the QLogic network adapters:

- 1 Press F2 during system boot to enter **System Setup**.
- 2 Click **System BIOS** and select **Device Settings**.

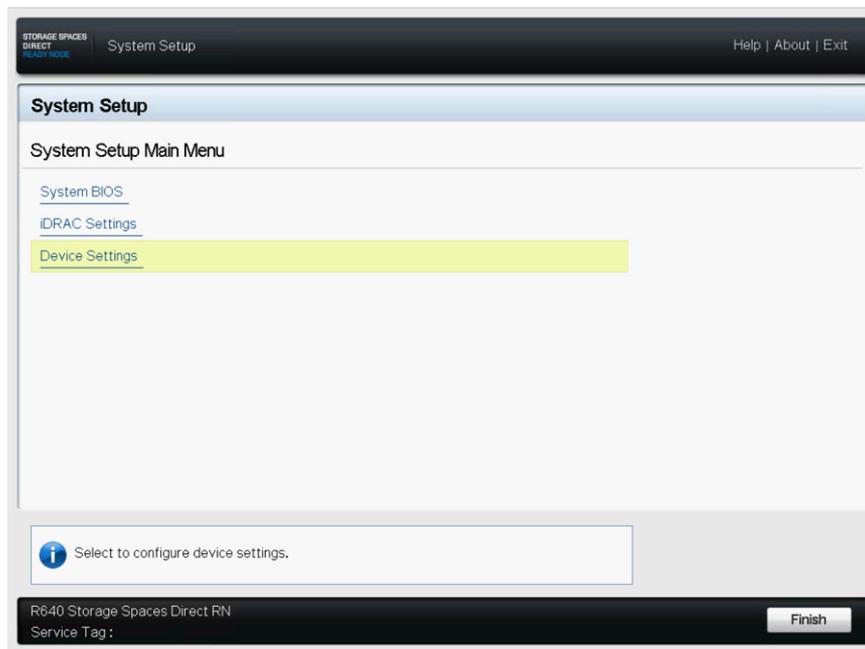


Figure 13. Device settings

- 3 Select the QLogic network adapter from the list of adapters.

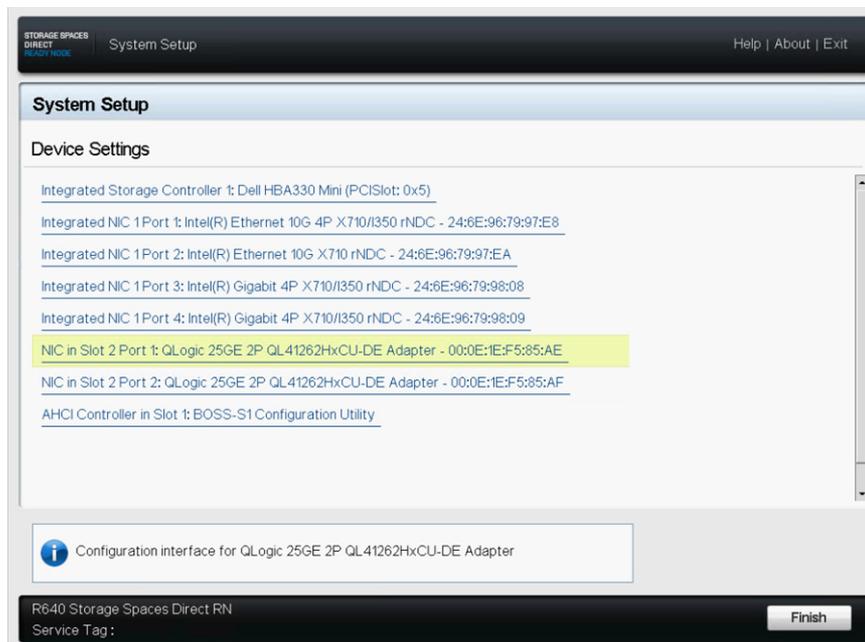


Figure 14. System Setup

- 4 Click **Device Level Configuration** and ensure that **Virtualization Mode** is set to None.

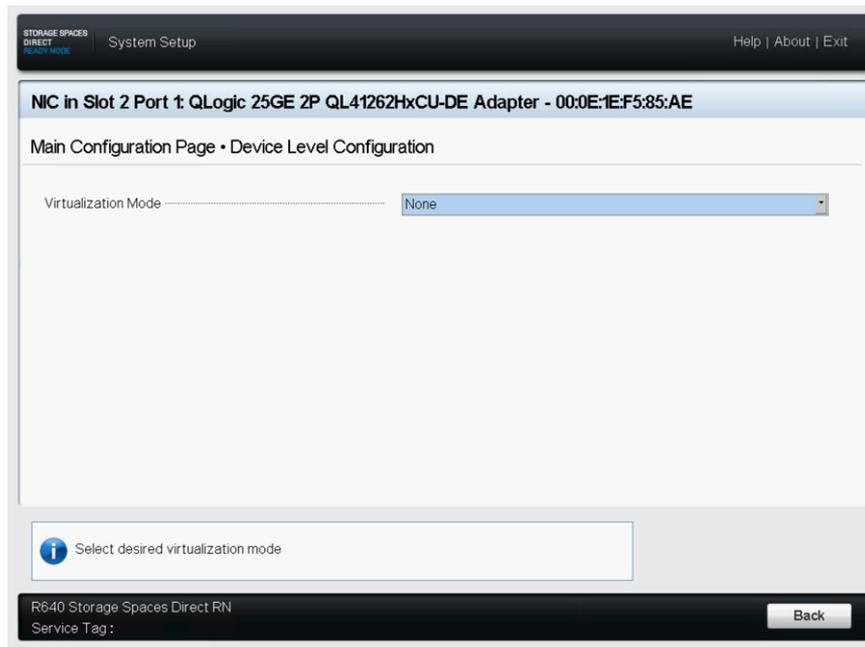


Figure 15. Virtualization mode

- 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

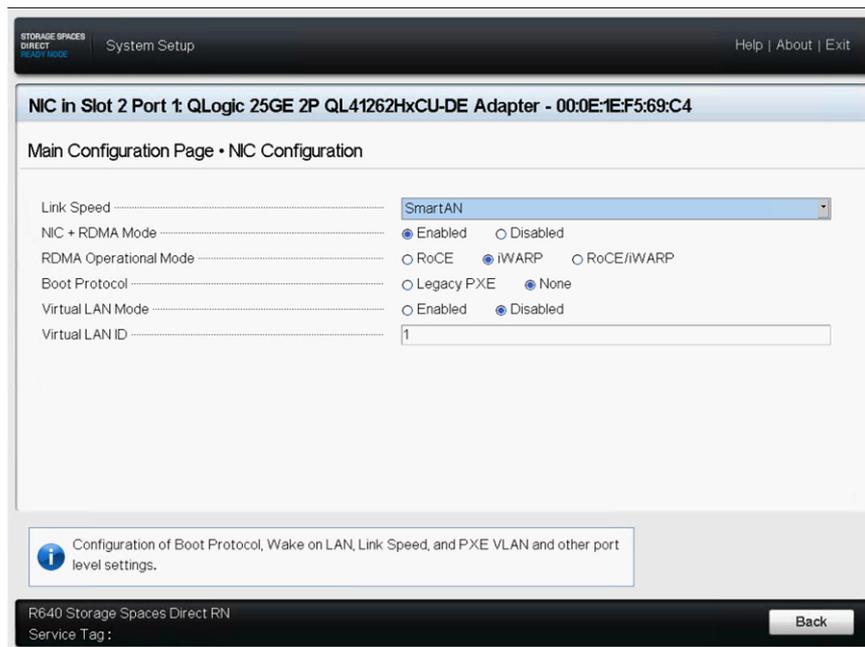


Figure 16. NIC configuration options

- 7 Click **Back** and click **Data Center Bridging (DCB) Settings**.
- 8 In the **Data Center Bridging (DCB) Settings** page, set **DCBX Protocol** to Disabled.

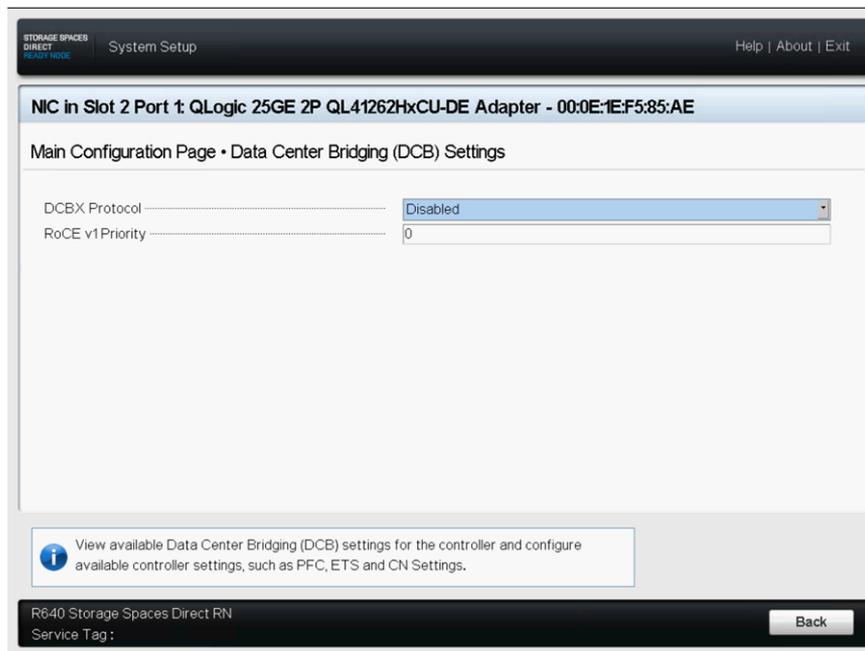


Figure 17. DCB settings

- 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.

Firmware baselining

The Dell EMC Ready Node has a supported firmware matrix and the nodes in the cluster must comply with the [firmware matrix](#). It is important to ensure that each server has the right firmware revisions for components used within the server.

This can be verified by using the iDRAC [system inventory feature](#) or by using a scripted method such as the RACADM command line interface.

Hyper-converged infrastructure deployment

This section describes the steps involved in installing OS on the bare metal servers and deploying the hyper-converged infrastructure (HCI) with Storage Spaces Direct. PowerShell commands are provided to configure cluster deployment from the command line.

Unless mentioned otherwise, the following steps should be performed on each physical node in the infrastructure that will be a part of Storage Spaces Direct HCI.

OS deployment

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 2019 OS pre-installed.

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

NOTE: The command output shown in the subsequent sections may show only Mellanox ConnectX-4 LX adapters as physical adapters. This is only shown as an example.

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

The factory installed OEM OS is pre-activated and the Hyper-V role is pre-deployed. Therefore, after the OOB steps are complete, the post OS deployment steps described in section [Install Roles and Features](#) should be performed to complete the cluster deployment and Storage Spaces Direct configuration.

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.

Update Out of Box drivers

For certain system components, there may be a need to update the inbox version of the driver to most up-to-date and Dell EMC supported versions listed in the [Supported Firmware and Software Matrix](#).

The following PowerShell command can be used to retrieve a list of all driver versions currently installed on the local system:

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

Before configuring host networking, ensure that the Out of Box (OOB) drivers are updated. After the required driver version is identified, you can download the driver installers from support.dell.com or by using the Dell EMC Solutions For Azure Stack HCI Update Catalog.

NOTE: QLogic FastLinQ adapter does not have an inbox driver in Windows Server 2016. This must be installed before attempting host network configuration.

After the drivers are downloaded, attach a folder containing these driver DUPs to the system as virtual media image. Perform the following steps to attach the virtual media image folder:

- 1 Click **Virtual Media** in the iDRAC virtual console menu.
- 2 Click **Create Image**.

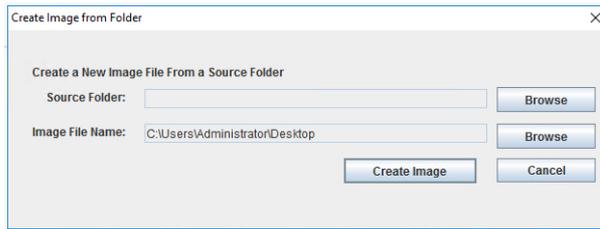


Figure 18. Create new image

- 3 Click **Browse** and select the folder where the driver DUP files are stored and change the name of the image, if required.

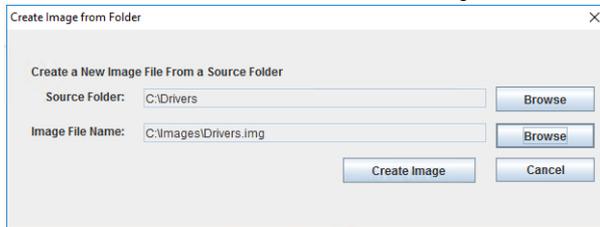


Figure 19. Virtual media image folder

- 4 Click **Create Image**.
- 5 Click **Finish**.
- 6 From the **Virtual Media** menu, click **Connect Virtual Media**.
- 7 From the **Virtual media** menu, click **Map Removable Disk** and click **Browse** to select the image created in the step 4.

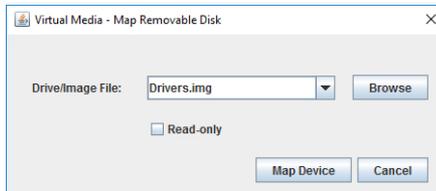


Figure 20. Map device

After the image is mapped, it appears as a drive in the host OS and the driver DUP files can be executed to install the OOB drivers.

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 S2DNode01 -Restart
```

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

Configure firewall

For the cluster operations post-deployment and optional monitoring configuration, certain firewall rules have to be enabled on the cluster nodes. For a complete list of ports or firewall rules that need to be enabled, see [Appendix A](#).

For configuring firewall rules at the command prompt, see [technet.microsoft.com/en-us/library/jj554906\(v=wps.630\).aspx](https://technet.microsoft.com/en-us/library/jj554906(v=wps.630).aspx).

Cluster deployment and configuration

The following sections provide instructions on cluster deployment and Storage Spaces Direct configuration.

Configure host networking

This section focuses on configuring the host networking such as VM switches, VM network adapters, and other QoS and RDMA configurations.

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.

VM switch and adapter configuration

The Dell EMC Solutions for Azure Stack HCI can be configured to implement a fully-converged or non-converged network configuration for storage and management connectivity. 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 RDMA capability, SET configuration is necessary for storage adapters.

NOTE: For specifics of configuration such as VM switch name, adapter names and VLAN IDs, see the [Deployment checklist](#).

You must configure the host OS network based on the physical connectivity of the network (fully converged or non-converged).

The following sections describe the two different host networking configurations and provide the commands for configuration.

NOTE: The following configuration instructions assume that only two physical ports are used for storage in a non-converged or for the VM switch in a fully-converged configuration. For configuration instructions that use four physical ports for storage or a fully-converged SET, see <https://community.emc.com/docs/DOC-73779>.

Fully converged network

The following figure illustrates the fully converged configuration in a Storage Spaces Direct cluster node.

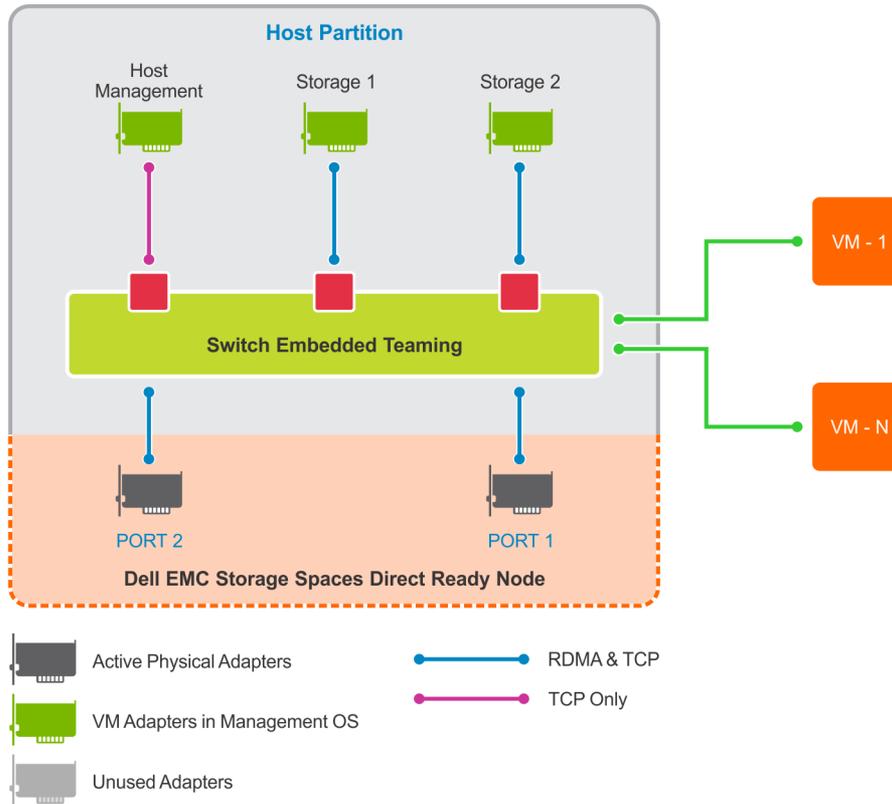


Figure 21. Fully Converged VM Switch and Adapter configuration

The following PowerShell commands can be used to configure the VM switch as a SET team, and configure the VM network adapters in the host OS for management, storage 1, and storage 2 traffic.

NOTE: The default load balancing algorithm for Windows Server 2016 SET is Dynamic and in Windows Server 2019 it is Hyper-V port. Dell EMC recommends that you always use the default load balancing algorithm.

Perform the following steps to configure the OS network:

- Run the following command to create a VM switch in the SET configuration by using the physical network ports from the Mellanox or QLogic network adapters in the system:

```
New-VMSwitch -Name S2DSwitch -AllowManagementOS 0 -NetAdapterName 'SLOT 1 PORT 1','SLOT 1 PORT 2' -MinimumBandwidthMode Weight -Verbose
```

In the above example, the arguments to the *NetAdapterName* parameter represents the physical NIC ports that need to be a part of the SET configuration. These interface names can be retrieved using the *Get-NetAdapter* cmdlet.

NOTE: The minimum Bandwidth Mode set to Weight can be used to shape the VM Network traffic, and it is not used for host OS network adapters. Setting the minimum Bandwidth Mode is optional.

Name	InterfaceDescription	ifIndex	Status
MacAddress	LinkSpeed		
NIC2	Intel(R) Ethernet 10G X710 rNDC	7	Disconnected
24-6E-96-52-CC-A4	10 Gbps		
NIC4	Intel(R) I350 Gigabit Network Conn...#2	6	Disconnected
24-6E-96-52-CC-C3	0 bps		

NIC3	Intel(R) I350 Gigabit Network Connec...	2	Disconnected
24-6E-96-52-CC-C2	0 bps		
NIC1	Intel(R) Ethernet 10G 4P X710/I350 rNDC	4	Disconnected
24-6E-96-52-CC-A2	10 Gbps		
SLOT 1 Port 2	Mellanox ConnectX-4 Lx Ethernet Ad...#2	10	Up
24-8A-07-59-4C-69	10 Gbps		
SLOT 1 Port 1	Mellanox ConnectX-4 Lx Ethernet Adapter	8	Up
24-8A-07-59-4C-68	10 Gbps		

The argument "0" to *AllowManagementOS* parameter prevents creation of a VM network adapter in the host operating system.

This command creates a SET with Switch Independent teaming mode and Dynamic Load Balancing algorithm settings.

- 2 **NOTE: For the host management network, do not run the `Set-VMNetworkAdapterVlan` command if DHCP is used (with untagged VLANs) for management.**

Run the following command to create and configure the host management network adapter:

```
Add-VMNetworkAdapter -ManagementOS -Name 'Management' -SwitchName S2DSwitch -Passthru | Set-VMNetworkAdapterVlan -Access -VlanId 102 -Verbose
```

- 3 Run the following command to add the host OS VM network adapters for Storage 1 and Storage 2 traffic and configure the VLAN IDs:

```
Add-VMNetworkAdapter -ManagementOS -Name 'Storage1' -SwitchName S2DSwitch -Passthru | Set-VMNetworkAdapterVlan -Access -VlanId 103 -Verbose
```

```
Add-VMNetworkAdapter -ManagementOS -Name 'Storage2' -SwitchName S2DSwitch -Passthru | Set-VMNetworkAdapterVlan -Access -VlanId 104 -Verbose
```

- NOTE: The number of storage VLANs can be equal to the number of storage adapters or 1.**

- 4 Run the following command to configure IP addresses:

```
#Host Management Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Management)' -IPAddress 172.16.102.51
DefaultGateway 172.16.102.1 -PrefixLength 27 -AddressFamily IPv4 -Verbose
#Storage 1 Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Storage1)' -IPAddress 172.16.103.51 -
PrefixLength
27 -AddressFamily IPv4 -Verbose
#Storage 2 Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Storage2)' -IPAddress 172.16.104.51 -
PrefixLength
27 -AddressFamily IPv4 -Verbose
```

Non-converged network

The following figure provides a high-level overview of VM switch and management adapter configuration in the Storage Spaces Direct Ready Node in a non-converged network configuration. This approach includes physical links for the storage traffic and a SET for host management and VM traffic.

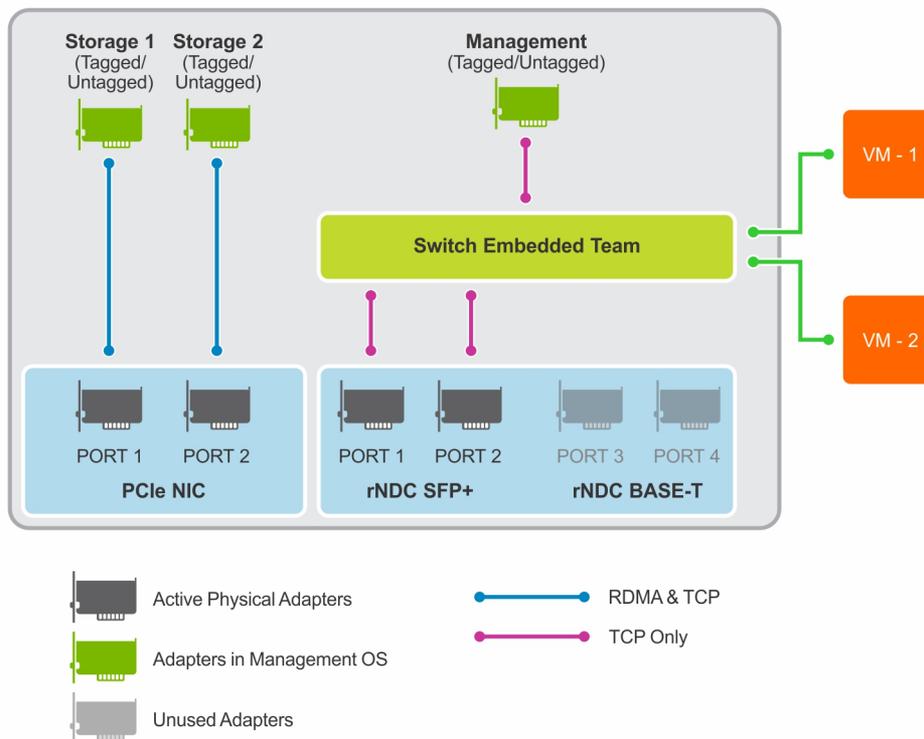


Figure 22. Non-converged network configuration

NOTE: Non-converged network configuration can implement a SET VM switch for storage as well. For instructions to configure SET for storage, see <https://community.emc.com/docs/DOC-73779>.

NOTE: The default load balancing algorithm for Windows Server 2016 SET is Dynamic and in Windows Server 2019 it is Hyper-V port.

The physical NIC ports from the Mellanox or QLogic network adapters are used for storage traffic and the SFP+ ports or the 1Gb BASE-T ports from the rNDC are used to create a VM switch in SET configuration. The host management and VM adapters get attached to this VM switch.

1 Run the following command to create a VM switch in the SET configuration by using the 10 GbE SFP+ network ports from rNDC:

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

NOTE: For the host management network, do not run the `Set-VMNetworkAdapterVlan` command if DHCP is used (with untagged VLANs) for management.

Run the following command to create and configure the host management network adapter:

```
Add-VMNetworkAdapter -ManagementOS -Name 'Management' -SwitchName Management -Passthru | Set-
VMNetworkAdapterVlan -Access -VlanId 102 -Verbose
```

3 Run the following command to configure VLAN IDs, if required, 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 VLANs can be equal to the number of storage adapters or 1.

4 Run the following command to configure IP addresses:

```
#Host Management Adapter New-NetIPAddress -InterfaceAlias 'vEthernet (Management)' -
IPAddress 172.16.102.51 DefaultGateway 172.16.102.1 -PrefixLength 27 -AddressFamily IPv4 -
Verbose
#DNS server address
Set-DnsClientServerAddress -InterfaceAlias 'vEthernet (Management)' -ServerAddresses
```

```

172.16.102.202
#Storage 1 Adapter
New-NetIPAddress -InterfaceAlias 'SLOT 1 PORT 1' -IPAddress 172.16.103.51 -PrefixLength
27 -AddressFamily IPv4 -Verbose
#Storage 2 Adapter
New-NetIPAddress -InterfaceAlias 'SLOT 1 PORT 2' -IPAddress 172.16.104.51 -PrefixLength
27 -AddressFamily IPv4 -Verbose

```

- ① **NOTE: Static IP assignment on the management adapter is not required if a DHCP server is being used for management network.**

In this configuration, default gateway and DNS configuration are required only for the host management network.

AD 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.

See the [Deployment checklist](#) for the domain administrator or equivalent credentials needed for the domain join.

- ① **NOTE: Connecting to AD directory services by using the host management network may require routing to the AD network. Ensure that this is in place before proceeding to domain join.**

```

$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 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.**

Create host cluster

Before creating a host cluster, Dell EMC recommends to ensure 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.

- ① **NOTE: The commands in this section need to be executed on only one node in the infrastructure.**
- ① **NOTE: 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.**
- ① **NOTE: Validate that the nodes have homogeneous hardware configuration.**

```

Test-Cluster -Node S2Dnode01, S2Dnode02, S2dNode03, S2dNode04 -Include 'Storage Spaces Direct',
'Inventory', 'Network', 'System Configuration'

```

- ① **NOTE: 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.**

Use the *New-Cluster* cmdlet to create the host cluster as follows:

```

New-Cluster -Name S2DSystem -Node S2Dnode01, S2Dnode02, S2dNode03, S2dNode04 -StaticAddress
172.16.102.55 -NoStorage -IgnoreNetwork 172.16.103.0/27, 172.16.104.0/27 -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.

- ① **NOTE: 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.**

Configuring 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 a 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
```

Change RDMA mode on QLogic NICs—iWARP only

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 1' -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'.

ⓘ | NOTE: This change is required only for Ready Nodes with QLogic 41262 adapters used for storage traffic.

RDMA configuration

Storage Spaces Direct hyper-converged Infrastructure with Mellanox Connect-X 4 LX and QLogic FastLinQ 41262 adapters use converged network design and RDMA for storage traffic. Hence, it is important to ensure that the physical and virtual 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 when storage traffic is going through a SET in the host OS. The argument to the Name parameter can be retrieved by using the `Get-NetAdapter` cmdlet:

```
Enable-NetAdapterRDMA -Name 'vEthernet (Storage1)', 'vEthernet (Storage2)'
```

If physical network ports are used for storage traffic, run the following command:

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

- Run the following command to map the storage virtual adapters in the host OS to the physical NIC ports to ensure that the storage traffic uses these affinity rules, and traffic spreads across evenly. This is not required if the storage traffic is using the physical network ports.

```
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName "Storage1" -ManagementOS -
PhysicalNetAdapterName "SLOT 1 PORT 1"
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName "Storage2" -ManagementOS -
PhysicalNetAdapterName "SLOT 1 PORT 2"
```

For each storage network adapter, the "RDMA Capable" value must be "True".

NOTE: With the inbox Mellanox drivers, this value appears as "False". This can be fixed by installing the updated drivers from Dell EMC. See the [Support Matrix](#) for details on the supported versions of firmware and drivers.

- Verify the RDMA configuration for the virtual storage adapters (in fully converged topology) by using the following command:

```
Get-SmbClientNetworkInterface | Where-Object { $_.FriendlyName -Like "*Storage*" }
```

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

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

QoS Policy configuration

NOTE: For configurations that use QLogic 41262 network adapters with iWARP for RDMA, and NVMe drives for storage, configuring DCB in the OS and TOR switches is recommended to ensure optimal storage performance. For other disk configurations such as the Hybrid and/or all-flash SSD configurations, configuring DCB for iWARP is not necessary. All Ready Node configurations using Mellanox adapters in fully-converged network topology should implement the steps mentioned in this section.

The Dell EMC Microsoft Storage Spaces Direct Ready nodes configured with Mellanox ConnectX-4 LX adapters require DCB configuration on the TOR switches and Quality of Service (QoS) configuration in the host OS. It is important to ensure that the QoS policies are configured to prioritize the SMB traffic related to the storage adapters. The QoS configuration in the host OS should match the QoS configuration performed in the [Network switch configuration section](#).

The following table provides an overview of the QoS priorities and the required state for the set of priorities.

NOTE: QoS configuration is needed only for fully-converged network topology. For non-converged network, the storage traffic uses dedicated NIC ports and therefore DCB/PFC and QoS configuration is not needed.

Table 16. QoS priorities

QoS Priority	State
0,1,2,4,6,7	Disabled
3	Enabled
5	Enabled

Perform the following steps to configure QoS in the host OS:

- Run the following command to create a new QoS policy with a match condition set to 445. This displays the TCP port dedicated for Server Message Block traffic:

```
New-NetQosPolicy -Name 'SMB' -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3
New-NetQosPolicy 'Cluster' -Cluster -PriorityValue8021Action 5
```

The arguments '3' and '6' to the *PriorityValue8021Action* parameter indicate the IEEE 802.1p values for SMB and cluster traffic and they should match the priority with the **Enabled** state as mentioned in the [table](#).

- 2 Run the following command to map the IEEE 802.1p priority enabled in the system to a traffic class:

```
New-NetQoSTrafficClass -Name 'SMB' -Priority 3 -BandwidthPercentage 50 -Algorithm ETS
New-NetQoSTrafficClass -Name 'Cluster' -Priority 5 -BandwidthPercentage 1 -Algorithm ET
```

The above commands specify that the transmission algorithm used is Enhanced Transmission Selection (ETS) and the traffic classes SMB and Cluster get bandwidth percentage of 50 and 1 respectively.

NOTE: The argument to `-BandwidthPercentage` shown in this step is an example only. This can be and should be modified based on the infrastructure requirements and type of network configuration.

- 3 Run the following command to configure flow control for the priorities mentioned in the [table](#):

```
Enable-NetQoSFlowControl -Priority 3,5
Disable-NetQoSFlowControl -Priority 0,1,2,4,6,7
```

- 4 Run the following command to enable QoS for the Mellanox network adapter ports. The argument for the *InterfaceAlias* can be retrieved by using the *Get-NetAdapter* cmdlet.

```
Enable-NetAdapterQos -InterfaceAlias 'SLOT 1 PORT 1','SLOT 1 PORT 2'
```

- 5 Run the following command to disable DCBX Willing mode in the operating system:

```
Set-NetQoSDCbxSetting -Willing $false
```

- 6 Run the following command to set Mellanox NIC DCBX mode to **Host in charge**:

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 1' -DisplayName 'Dcbxmode' -DisplayValue 'Host in charge'

Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 2' -DisplayName 'Dcbxmode' -DisplayValue 'Host in charge'
```

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

NOTE: For performance optimization and reliability, Dell EMC recommends that you update the hardware timeout configuration for the 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
```

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.

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 specifically on the interfaces supporting the Storage network, helps improve the Storage Spaces Direct cluster overall read/write performance. It is important to note that an end-to-end configuration of jumbo frames is required in order to take advantage of the technology. In addition, considerations need to be made when configuring the technology as support for jumbo frame sizes vary between software, NIC, and switch vendors. The lowest value within the data path determines the maximum frame size used for that path.

For the storage network adapters in the host OS, this can be done by using the `Set-NetworkAdapterAdvancedProperty` cmdlet.

NOTE: Network adapters from different vendors support different jumbo packet sizes. The configured value must be consistent across the host OS and network switch configuration.

For information on configuring jumbo frames at the switch port level, see [Sample Switch Configurations](#).

Configuring cluster witness

Microsoft recommends configuring a cluster witness for a four node Storage Spaces Direct cluster. A cluster witness must be configured for a 2-node cluster.

Cluster witness configuration 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.

NOTE: If you choose to configure a file share witness, it should exist outside the 2-node cluster.

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://technet.microsoft.com/en-us/windows-server-docs/failover-clustering/deploy-cloud-witness>.

Recommended next steps

Before proceeding to the next steps and operational management of the cluster, generate a cluster validation report to ensure that all configuration is in order. This can be done using the `test-Cluster` cmdlet.

```
Test-Cluster -Node S2DNode01, S2DNode02, S2DNode03, S2DNode04 -Include 'System Configuration',  
'Inventory', 'Network', 'Storage Spaces Direct'
```

This command generates an HTML report with a list of all the tests that were performed and completed without errors.

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

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

 **NOTE:** 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/azurestackhcinmanuals>.

Deployment services

Issues that arise during installation and configuration are not covered even if you have purchased Dell ProSupport or ProSupport Plus, as support for installation and configuration issues come under a separate paid services package. When you call in with a installation and configuration 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.

Additional resources

- [iDRAC documentation](#)
- [Supported firmware and software matrix](#)
- [Storage Spaces Direct overview](#)

Firewall port requirements

Table 17. Firewall port requirements

Source	Target	Protocol	Port	Comment
Any	Domain Controllers	TCP/UDP	53	DNS
		TCP/UDP	88	Kerberos
		UDP	123	NTP
		TCP	135	RPC, EMP
		UDP	137	NetLogon, NetBIOS Name Resolution
		UDP	138	DFSN, NetLogon, NetBIOS, Datagram Service
		TCP	139	DFSN, NetBIOS Session Service, NetLogon
		TCP/UDP	389	LDAP
		TCP/UDP	445	SMB, CIFS, SMB2, DFSN, LSARPC, NbtSS, NetLogonR, SAMR, SrvSvc
		TCP/UDP	464	Kerberos change/set password
		TCP	636	LDAP (SSL)
		TCP	3268	Global Catalog
		TCP	3269	Global Catalog (SSL)
		TCP	5722	RPC, DFSR (SYSVOL)
		TCP	9389	SOAP
		TCP	1025:5000	RPC, DCOM, EPM, DRSUAPI, NetLogon, SamR, FRS (2003)
UDP	1025:5000	DCOM, RPC, EPM (2003)		
TCP	49152:65535	RPC, DCOM, EPM, DRSUAPI, NetLogonR, SamR, FRS (2008)		
UDP	49152:65535	DCOM, RPC, EPM (2008)		
Local Subnet	All Hosts and VMs	UDP	137:138	Allow Name/Share Resolution

Source	Target	Protocol	Port	Comment
		TCP	139	Allow Name/Share Resolution
Any	Console VM	TCP	3389	Remote Desktop
WSUS (on VMM VM)	Any	TCP	80	SWUS Updates (HTTP)
		TCP	443	SWUS Updates (HTTPS)

Sample deployment checklists

Table 18. Sample checklist

Fields	Values
AD Domain FQDN	hci.lab
Domain Administrator or equivalent credentials	Username: hci\administrator Password: <DO NOT WRITE IT DOWN>
DNS Server addresses	dns.s2dlab.local
WSUS Server FQDN (if needed)	wsus.s2dlab.local

Table 19. Sample checklist

Traffic Class	Purpose	Minimum IP addresses needed	VLAN ID	Tagged / Untagged	Subnet mask	VLAN IP Addresses
Out of band	Required for OOB management of server nodes and TOR switches	18	100	Untagged	/24	OOB: 172.16:100.1
Host Management	Management of cluster and cluster nodes	17	102	Tagged	/24	TOR1: NA TOR2: NA
Storage 1	SMB traffic	16	103	Tagged	/24	TOR1: NA TOR2: NA
Storage 2	SMB Traffic	16	104	Tagged	/24	TOR1: NA TOR2: NA

Table 20. Sample checklist

Fields	Values
OOB Switch hostname	S2D-OOB
TOR1 Switch hostname	S2D-TOR1
TOR2 Switch hostname	S2D-TOR2
Enable password	<DO NOT WRITE IT DOWN>

Fields	Values
Additional user/password	NA
IP route on OOB (if needed)	NA
IP route on TOR1 / TOR2 (if needed)	NA
DCB Bandwidth for SMB traffic	50%

Table 21. Sample checklist

Fields	Values
Virtual Switch	S2DSwitch
Management Adapter	Management
Storage 1 Adapter	Storage1
Storage 2 Adapter	Storage2

Table 22. Sample checklist

	Host Name	Management IP	Storage1 IP	Storage2 IP	OOB IP	OOB Host name
Node 1	S2DNode01	172.16.102.51	172.16.103.51	172.16.104.51	172.16.100.51	S2D-DRAC-1
Node 2	S2DNode02	172.16.102.52	172.16.103.52	172.16.104.52	172.16.100.52	S2D-DRAC-2
Node 3	S2DNode03	172.16.102.53	172.16.103.53	172.16.104.53	172.16.100.53	S2D-DRAC-3
Node 4	S2DNode04	172.16.102.54	172.16.103.54	172.16.104.54	172.16.100.54	S2D-DRAC-4
Node 5	S2DNode05	172.16.102.55	172.16.103.55	172.16.104.55	172.16.100.55	S2D-DRAC-5
Node 6	S2DNode06	172.16.102.56	172.16.103.56	172.16.104.56	172.16.100.56	S2D-DRAC-6
Node 7	S2DNode07	172.16.102.57	172.16.103.57	172.16.104.57	172.16.100.57	S2D-DRAC-7
Node 8	S2DNode08	172.16.102.58	172.16.103.58	172.16.104.58	172.16.100.58	S2D-DRAC-8
Node 9	S2DNode09	172.16.102.59	172.16.103.59	172.16.104.59	172.16.100.59	S2D-DRAC-9
Node 10	S2DNode10	172.16.102.60	172.16.103.60	172.16.104.60	172.16.100.60	S2D-DRAC-10
Node 11	S2DNode11	172.16.102.61	172.16.103.61	172.16.104.61	172.16.100.61	S2D-DRAC-11
Node 12	S2DNode12	172.16.102.62	172.16.103.62	172.16.104.62	172.16.100.62	S2D-DRAC-12
Node 13	S2DNode13	172.16.102.63	172.16.103.63	172.16.104.63	172.16.100.63	S2D-DRAC-13
Node 14	S2DNode14	172.16.102.64	172.16.103.64	172.16.104.64	172.16.100.64	S2D-DRAC-14

	Host Name	Management IP	Storage1 IP	Storage2 IP	OOB IP	OOB Host name
Node 15	S2DNode15	172.16.102.65	172.16.103.65	172.16.104.65	172.16.100.65	S2D-DRAC-15
Node 16	S2DNode16	172.16.102.66	172.16.103.66	172.16.104.66	172.16.100.66	S2D-DRAC-16