

Dell EMC Microsoft Storage Spaces Direct Ready Node

Deployment Guide for scalable hyper-converged infrastructure with R740xd 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.

Copyright © 2018 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 Ready Nodes.....	6
R740xd Storage Spaces Direct Ready Node.....	8
R640 Storage Spaces Direct Ready Node.....	9
Dell EMC Networking S4112F-ON.....	10
Dell EMC Networking S4128-ON.....	11
Dell EMC Networking S4148-ON.....	11
Dell Networking S4048-ON.....	11
Dell EMC Networking S3048-ON.....	11
Dell EMC Networking S5048F.....	11
Windows Server 2016.....	11
Storage Spaces Direct.....	12
System Center Virtual Machine Manager 2016.....	12
System Center Operations Manager 2016.....	12
3 Integrating solution components.....	14
Network Connectivity in R740xd Storage Spaces Direct Ready Nodes.....	15
Network Connectivity in R640 Storage Spaces Direct Ready Nodes.....	18
Port mapping.....	21
4 Virtualized infrastructure deployment by using Dell EMC Ready Nodes.....	24
Deployment prerequisites.....	24
Software versions.....	25
Dell EMC validated firmware matrix.....	25
Deployment checklist.....	25
Management environment checklist.....	26
Network configuration checklist.....	26
Host OS network checklist.....	27
Pre-deployment configuration.....	28
Network switch configuration.....	28
iDRAC and BIOS configuration.....	29
QLogic NIC configuration.....	32
Firmware baselining.....	35
Hyper-converged infrastructure deployment.....	35
Recommended next steps.....	49
5 Deployment services.....	51
6 Additional resources.....	52

A Firewall port requirements.....	53
B Sample deployment checklists.....	55
C VMM preparation.....	58
Create VMM host group.....	59
Add Run As Account for cluster creation.....	59
VMM logical network details.....	60
Create and configure logical and VM networks—fully converged.....	60
Create and configure logical and VM networks—dedicated management.....	61
Create IP address pools.....	63
Create an uplink port profile.....	64
Create a logical switch.....	64

Audience and scope

This deployment guide focuses on deploying a scalable hyper-converged infrastructure solution on Dell EMC Microsoft Storage Spaces Direct Ready Nodes. This deployment guide includes an overview of the solution infrastructure, guidance to integrate the solution components, and instructions on 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 Hyper-V and Storage Spaces Direct.

Dell EMC Microsoft Storage Spaces Direct Ready Nodes can be ordered with the operating System pre-installed from the factory with OEM license or bare metal for customers who already have Volume License agreements with Microsoft.

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

- Dell EMC Services led: Certified deployment engineers ensure accuracy, speed, reduce 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. These instructions are not validated with Windows Server version 1709. Storage Spaces Direct Ready nodes do not support the Windows Server Semi-Annual Channel release.

Topics:

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

Assumptions

This 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 and virtual link trunk
- Deploying and configuring Windows Server 2016 Hyper-V infrastructure
- Deploying and configuring System Center products such as SCOM, if the deployment involves onboarding Storage Spaces Direct cluster into the existing System Center

Known issues

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

Virtualization infrastructure with Dell EMC Ready Nodes

Dell EMC Microsoft Storage Spaces Direct Ready Nodes encompasses various configurations of R740xd and R640 Storage Spaces Direct Ready Node to power the primary compute cluster deployed as 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 along side the redundant top of rack switches, a separate out of band network, and an existing management infrastructure in the data center.

This hyper-converged virtualized solution based on Dell EMC 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

These Dell EMC Microsoft Storage Spaces Direct Ready Nodes 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 is available in the existing management infrastructure.

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](#)
- [R640 Storage Spaces Direct Ready Node](#)
- [Dell EMC Networking S4112F-ON](#)
- [Dell EMC Networking S4128-ON](#)
- [Dell EMC Networking S4148-ON](#)
- [Dell Networking S4048-ON](#)
- [Dell EMC Networking S3048-ON](#)
- [Dell EMC Networking S5048F](#)
- [Windows Server 2016](#)
- [System Center Virtual Machine Manager 2016](#)
- [System Center Operations Manager 2016](#)

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)

Configuration	Chassis	Drive layout
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)
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: - 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

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 3. 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)

Component	Chassis	Drive layout
		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 4. 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: - up to two Mellanox Connectx-4 LX 25 GbE SFP add-in adapter card - 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

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.

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.

Dell EMC Networking S4112F-ON

Dell EMC Networking S4112F-ON is a 10 GbE L2, and L3 capable switch with 12x10 GbE and 3x100 GbE ports. This switch has the port density that is optimized for small and medium-sized infrastructure and software-defined storage infrastructure.

Dell EMC Networking S4128-ON

Dell EMC Networking S4128-ON is a 10 GbE, L2, and L3 capable network switch optimized for hyper-converged deployment. This switch provides up to 2 x 10/25/40/50/100G ports for uplink or data center core connectivity.

Dell EMC Networking S4148-ON

Dell EMC Networking S4148-ON is a 10 GbE, L2, and L3 capable network switch optimized for hyper-converged deployment. This switch provides up to 4x10/25/40/50/100G ports for uplink or data center core connectivity.

Dell Networking S4048-ON

Dell Networking S4048-ON is a 10 GbE, L2, and L3 capable network switch for storage, cluster, and client traffic. This switch provides up to 6 x 40 GbE ports for uplink or data center core connectivity.

For validated operating system versions on the Dell Networking switches, see [Supported Firmware and Software Matrix](#).

Dell EMC Networking S3048-ON

Dell Networking S3048-ON is a 1000 BASE-T, L2, and L3 capable switch that provides 48 ports that support 10 MbE/100 MbE/1 GbE and four 10 GbE SFP+ uplinks. In this Ready Nodes configuration, S3048-ON is deployed to support the Out Of Band (OOB) connectivity between the Ready Nodes and the Dell Networking S4048-ON switches.

For a sample configuration of the Top-Of-Rack (TOR) switches with DCB and other required configurations, see [Sample Switch Configurations](#).

Dell EMC Networking S5048F

The Dell EMC Networking S5048F is a 48 port 25 GbE switch that features 6 ports of 100 GbE and ONIE for zero-touch installation of alternate network operating systems.

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.

Windows Server 2016

Windows Server 2016 is the latest server operating system offered by Microsoft. This operating system serves as the foundation for this entire deployment, because every node uses Windows Server 2016 and its multiple capabilities.

For a list of new features and changes, see <https://technet.microsoft.com/en-us/windows-server-docs/get-started/what-s-new-in-windows-server-2016>

Storage Spaces Direct

One of the new features of Windows Server 2016 is Storage Spaces Direct. This new storage feature uses the local disks within the nodes to create highly available software defined storage.

There are two different deployment options possible for Storage Spaces Direct. In the first method, the storage and compute cluster are kept separate. This deployment option is known as converged or disaggregated deployment. This method allows for scaling of storage and compute clusters in a manner that is independent of each other. The second deployment option, known as hyper-converged deployment, enables running the virtualization services directly on top of the servers hosting storage spaces direct. This ensures that there is no need to configure and maintain file services separately in a different cluster and therefore reduces the need for additional physical servers.

This deployment guide focuses on the hyper-converged deployment option of the storage spaces direct. For more information about Storage Spaces Direct and these deployment options, see <https://technet.microsoft.com/en-us/windows-server-docs/storage/storage-spaces/storage-spaces-direct-overview>.

System Center Virtual Machine Manager 2016

Virtual Machine Manager serves as a central point of management for virtualization environments, cloud deployments and more. This deployment guide includes instructions on using VMM to deploy a host cluster and enabling Storage Spaces Direct.

For more information about VMM, see [https://technet.microsoft.com/en-us/library/gg610610\(v=sc.12\).aspx](https://technet.microsoft.com/en-us/library/gg610610(v=sc.12).aspx)

System Center Operations Manager 2016

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

Dell EMC Ready Nodes can be monitored using the operations manager management packs associated with the bare metal hardware and Storage Spaces Direct cluster. The cluster operations section in this guide describes how to deploy these management packs and configure monitoring for the bare metal nodes and the Storage Spaces Direct cluster performance monitoring.

For more information, see [https://technet.microsoft.com/en-us/library/hh205987\(v=sc.12\).aspx](https://technet.microsoft.com/en-us/library/hh205987(v=sc.12).aspx).

This guide provides the deployment instructions for implementing the Storage Spaces Direct hyper-converged cluster built using the Dell EMC Ready Nodes. The following flowchart provides an overview of these instructions and the flow of deployment instructions.




Figure 2. Deployment overview

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 3. 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/wssdmanuals>.

Topics:

- [Network Connectivity in R740xd Storage Spaces Direct Ready Nodes](#)
- [Network Connectivity in R640 Storage Spaces Direct Ready Nodes](#)
- [Port mapping](#)

Network Connectivity in R740xd Storage Spaces Direct Ready Nodes

The host networking on each R740xd Storage Spaces Direct Ready Node can be configured in two different ways.

Fully converged

In this configuration option, each R740xd Storage Spaces Direct Ready Node 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 this 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 4. Network and iDRAC connection

Dedicated management network

In this second network connectivity option, each R740xd Storage Spaces Direct Ready Node 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 5. Dedicated Host Management connectivity using the 10 GbE SFP+ rNDC ports in R740xd Storage Spaces Direct Ready Node




Figure 6. Dedicated Host Management connectivity using the 1 Gb BASE-T rNDC ports in R740xd Storage Spaces Direct Ready Node

The dedicated 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, only one port (Port 1) per NIC is used to provide redundancy at the network adapter level.

Each R740xd 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.

Network Connectivity in R640 Storage Spaces Direct Ready Nodes

Based on the chassis configuration, the type of network adapter in the R640 Storage Spaces Direct Ready Node differs. In the 12-drive chassis, there is only a Mellanox Connect-X 4 LX rNDC or a QLogic FastLinQ 41262 rNDC available for host management and storage traffic. For the 10 drive chassis variant, up to two Mellanox Connect-X 4 LX or a QLogic FastLinQ 41262 25 GbE SFP28 adapter card can be used.

The 10-drive variant of the R640 Storage Spaces Direct Ready Node also includes an on-board Intel Ethernet 10G 4P X710/I350 rNDC. Similar to R740xd Storage Spaces Direct Ready Node, the 10 drive chassis can have two different network connectivity options—fully converged and dedicated host management network.

Fully converged network in R640 10-drive chassis

In this configuration option, each R640 Storage Spaces Direct Ready Node 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 the network adapter on each server is connected to a different switch in the redundant top-of-rack (TOR) network fabric.




Figure 7. Fully converged network connectivity in R640 Storage Spaces Direct Ready Node (10-drive chassis)

Dedicated management network in R640 10-drive chassis

In this second network connectivity option, each R640 Storage Spaces Direct Ready Node 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 10Gb SFP+ DP 1 Gb BASE-T rNDC for host management and VM traffic.

The following figures illustrate the network connectivity in these chassis configurations.




Figure 8. Dedicated host management connectivity in R640 Storage Spaces Direct Ready Node (10 drive chassis)

Fully converged network connectivity in 12-drive chassis

The R640 Storage Spaces Direct Ready Node configuration supports only one network adapter that is an rNDC. Therefore, this chassis configuration supports only fully converged network design.




Figure 9. Fully converged network connectivity in R640 Storage Spaces Direct Ready Node (12-drive chassis)

In the dual-NIC configuration, with either Mellanox or QLogic adapters, only one port (Port 1) per NIC is used to provide redundancy at the network adapter level.

Each R640 Storage Spaces Direct Ready Node server has a dedicated network connection from the Integrated Dell Remote Access Controller (iDRAC) to the Dell EMC Networking S3048-ON switch configured for OOB management.

Port mapping

This section provides recommendations to switch port mapping on a standard 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 10. 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 Trunk (VLT) for inter-switch connectivity.



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 11. Network connectivity with Dell Networking S3048 (OOB Switch) connecting to Dell Networking S4048

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 and R640 Storage Spaces Direct Ready Nodes can be ordered from the factory with Windows Server 2016 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 also provide details on the supported [software and firmware versions](#).

- NOTE:** Instructions in this deployment guide are applicable only to Windows Server 2016 generally available OS build with latest applicable updates. These instructions are not validated with Windows Server version 1709. S2D Ready nodes do not support the Windows Server Semi-Annual Channel release.
- 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](#)
- [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 5. Management services

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

Windows Software Update Service (WSUS)	Provides local source for Windows Updates	Optional
MS SQL Server	Provides database backend for VMM and SCOM	Optional
System Center VMM	Provides virtualization host and VM management	Optional
System Center Operations Manager	Provides monitoring and reporting services for compute infrastructure	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 6. Software versions

Components	Version
Operating System	Windows Server 2016 Data Center Core
Virtual Machine Manager (where applicable)	SCVMM 2016 with UR 4
Operations Manager (where applicable)	SCOM 2016 with UR 4
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/wssdmanuals> 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 7. Management environment checklist

AD Domain FQDN	
Domain Administrator or equivalent credentials	
VMM Server FQDN (optional)	
VMM Administrator credentials (optional)	
DNS Server addresses	
SCOM Server FQDN (optional)	
SCOM Administrator credentials (optional)	
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 8. Network configuration

Traffic Class	Purpose	Minimum IP addresses needed	VLAN ID	Tagged/Untagged	IP address space	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	/26	TOR1: TOR2:

Storage 1	SMB traffic	16		Tagged	/27	TOR1: TOR2:
Storage 2	SMB Traffic	16		Tagged	/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 9. Network configuration checklist

OOB Switch hostname	
TOR1 Switch hostname	
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 (RoCE only)	

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 10. Host OS network switch and adapter details

Virtual Switch/Logical Switch Name	
Management Adapter/ Logical & VM Network Name	
Storage 1 Adapter/ Logical and VM Network Name	
Storage 2 Adapter/ Logical and VM Network Name	
Uplink Port Profile Name (VMM Only)	
Management IP Pool range (VMM only)	
Storage1 IP pool range (VMM only)	
Storage2 IP pool range (VMM only)	

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 11. 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						
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 hyper-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 deploying 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 one switch fails, the virtual IP is transferred to the peer switch.

There are five basic networks required for a standard Storage Spaces Direct deployment—Switch Management, OOB Management, Host Management, Storage1, and Storage2.

Storage typically has two fault domains, referenced here as Storage 1 and Storage 2.

For sample switch configurations for Dell EMC Networking S5058F-ON, S4148-ON, and S4048-ON, see <>

Table 12. Solution Network VLANs

VLAN NETWORK TYPE	MINIMAL NETWORK MASK	HOST IP ADDRESSES	VLAN ID TAG/UNTAG
Out of band Management	/27 (255.255.255.224)	19	Untagged
Host Management	/26 (255.255.255.192)	17	Untagged
Storage 1 (Fault Domain 1)	/27 (255.255.255.224)	16	Tagged
Storage 2 (Fault Domain 2)	/27 (255.255.255.224)	16	Tagged

iDRAC and BIOS configuration

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

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 and/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 12. Enter iDRAC

- 2 Select **iDRAC Settings**.




Figure 13. System Setup Main Menu

- 3 Select **Network**.




Figure 14. iDRAC Settings

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




Figure 15. Network Settings

- 5 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 and therefore based on the network configuration chosen, you must configure the adapter manually to either enable iWARP or RoCE 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**.
- 3 Select the QLogic network adapter from the list of adapters.




Figure 16. Device Settings




Figure 17. System Setup

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




Figure 18. 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 19. 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 20. 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 HCI 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 as relevant, 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 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 the [Unified Server Configurator](#) (USC).

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 Core Data Center edition deployment on the physical server is complete and that you have access to the [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 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 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 for RoCE over RDMA only)

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 'DEVPKKEY_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**.

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 21. 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 22. 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 23. 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 B](#).

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).

VMM based cluster deployment and configuration

When System Center Virtual Machine Manager is present in the existing data center environment, it is recommended to use VMM networking and deployment artifacts to configure host networking and create and configure a Storage Spaces Direct cluster. The following subsections describe the steps involved in performing these configuration tasks and provide PowerShell commands for each configuration task.

If VMM based deployment of cluster is chosen, configuration tasks in the section *Manual Cluster Deployment and Configuration* should be ignored.

- ❗ **NOTE:** Ensure that the prerequisites (listed in the section [VMM preparation](#)) for Storage Spaces Direct cluster creation and configuration by using VMM are met before proceeding to the next section.

Configure Host Management VLAN

To be able to join the active directory domain, the physical network adapter must have the tagged VLAN ID from the host management subnet. This can be configured using the following PowerShell command.

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 1' -DisplayName 'VLAN ID' -DisplayValue 102
```

The above command will set the VLAN ID property of the physical network adapter port identified by SLOT 1 PORT 1 as 102.

- ❗ **NOTE:** This command must be run on each host that will be a part of the Storage Spaces Direct cluster.
- ❗ **NOTE:** You can see the name of the network adapter using the `Get-NetAdapter` cmdlet.
- ❗ **NOTE:** This step should not be performed if you have DHCP server for management network and it uses untagged VLANs.

Assign Host Management IP address

If there is no DHCP server for management network, the network adapter port identified as SLOT 1 PORT 1 should be configured with an IP v4 address to be able to communicate with the Active Directory domain controllers and join the domain. Also, the name server configuration must be complete to ensure the domain name resolution works as expected during domain join operation.

- ❗ **NOTE:** Note: The static or DHCP IP assigned to the physical NIC port should be from the same IP range and pool as the Host Management IP Pool configured in VMM. See the VMM Preparation section and the checklist for more information.

```
New-NetIPAddress -InterfaceAlias 'SLOT 1 PORT 1' -IPAddress 172.16.102.51 -DefaultGateway 172.16.102.1 -PrefixLength 24
```

```
Set-DnsClientServerAddress -InterfaceAlias 'SLOT 1 PORT 1' -ServerAddresses 172.16.102.202
```

The above command sets the IPv4 addresses of the physical network port identified as SLOT 1 PORT 1 and sets the DNS server address on the same interface for name resolution.

- ❗ **NOTE:** This command needs to run on each host that will be a part of the Storage Spaces Direct cluster. For IP address assignment for each host, refer to the deployment checklist.

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.

- ❗ **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
```

- ❗ **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.
- ❗ **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.

Add physical hosts to VMM

The VM hosts that are domain joined can be added to the VMM host group.

- ❶ **NOTE:** All commands in this section use VMM jobs. Therefore, after executing any set of commands, monitor the VMM job queue and identify any job failures.
- ❶ **NOTE:** Commands in this section and subsequent section require the SCVMM PowerShell commands and therefore should be either run directly on the VMM server or with a system that has the SCVMM PowerShell cmdlets installed and has access to the SCVMM server. If running on a remote system with SCVMM PowerShell module, the logged on account on the system must have right authorization perform VMM actions. Also, these commands assume that SCVMM server has network connectivity to the physical hosts over the host management network.

```
$runAsAccount = Get-SCRunAsAccount -Name 'Administrator'
$hostGroup = Get-SCVMHostGroup -Name 'SpacesDirectHosts'
Add-SCVMHost -ComputerName "s2dnode01.hci.lab" -RunAsynchronously -VMHostGroup $hostGroup -
Credential $runAsAccount
Add-SCVMHost -ComputerName "s2dnode02.hci.lab" -RunAsynchronously -VMHostGroup $hostGroup -
Credential $runAsAccount
Add-SCVMHost -ComputerName "s2dnode03.hci.lab" -RunAsynchronously -VMHostGroup $hostGroup -
Credential $runAsAccount
Add-SCVMHost -ComputerName "s2dnode04.hci.lab" -RunAsynchronously -VMHostGroup $hostGroup -
Credential $runAsAccount
```

- ❶ **NOTE:** This brings all hosts into the VMM for virtualization management.

Configure fully-converged logical switch on Hyper-V Hosts

As the next step in preparation for the cluster creation and Storage Spaces Direct configuration, the logical switch needs to be deployed with the appropriate VM Network and subnet configurations. The following set of PowerShell commands can be used to perform this configuration.

- ❶ **NOTE:** The following PowerShell commands implement a fully-converged logical switch on the Hyper-V hosts. If a dedicated management switch is needed, skip this section and proceed to the section [Configure fully configured logical switch on Hyper-V hosts with dedicated management](#).

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

#FQDN

$nodeName = 's2dnode01.hci.lab'

# Get Host 's2dnode01.hci.lab'
$vmHost = Get-SCVMHost | Where-Object { $_.Name -eq $nodeName }

# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter'
$networkAdapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName -eq 'SLOT 1 PORT 1' }
$uplinkPortProfileSet = Get-SCUplinkPortProfileSet -Name 'S2D_UPP'
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $networkAdapter -UplinkPortProfileSet $uplinkPortProfileSet -JobGroup $jobGroup

# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter #2'
$networkAdapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName -eq 'SLOT 1 PORT 1' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $networkAdapter -UplinkPortProfileSet $uplinkPortProfileSet -JobGroup $jobGroup

$networkAdapter = @()
$networkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName -eq 'SLOT 1 PORT 1' }
$networkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object { $_.ConnectionName -eq 'SLOT 1 PORT 1' }

$logicalSwitch = Get-SCLogicalSwitch -Name S2dSwitch

#Management
$vmNetwork = Get-SCVMNetwork -Name 'Management'
$vmSubnet = Get-SCVMSubnet -Name 'Management_0'
New-SCVirtualNetwork -VMHost $vmHost -VMHostNetworkAdapters $networkAdapter -LogicalSwitch
```

```
$logicalSwitch -JobGroup $jobGroup -CreateManagementAdapter -ManagementAdapterName "Management" -ManagementAdapterVMNetwork $vmNetwork -ManagementAdapterVMSubnet $vmSubnet

#Storage1
$vmNetwork = Get-SCVMNetwork -Name 'Storage1'
$vmSubnet = Get-SCVMSubnet -Name 'Storage1_0'
$ipV4Pool = Get-SCStaticIPAddressPool -Name 'Storage-IPPool'
New-SCVirtualNetworkAdapter -VMHost $vmHost -Name "Storage1" -VMNetwork $vmNetwork -LogicalSwitch $logicalSwitch -JobGroup $jobGroup -VMSubnet $vmSubnet -IPv4AddressType "Static" -IPv4AddressPool $ipV4Pool -MACAddressType "Static" -MACAddress "00:00:00:00:00:00"

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

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

The above commands need to be executed for each Hyper-V host that will be a part of the cluster and Storage Spaces Direct configuration.

Configure fully configured logical switch on Hyper-V hosts with dedicated management

The following PowerShell commands enable creation of two logical switches on the hyper-V hosts for host management and storage respectively. The following commands must be executed on the VMM server for each node in the Storage Spaces Direct Cluster.

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

#Change the name of the host before each run
$nodeName = 's2d2n01.hci.lab'

# Get Host
$vmHost = Get-SCVMHost | Where-Object { $_.Name -eq $nodeName }

# Get Host Network Adapter 'Intel(R) Ethernet 10G 4P X710/I350 rNDC'
$mgmtUplinkPortProfileSet = Get-SCUplinkPortProfileSet -Name 'S2D_Management_UPP'
$mgmtNetworkAdapter1 = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'NIC1' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $mgmtNetworkAdapter1 -UplinkPortProfileSet
$mgmtUplinkPortProfileSet -JobGroup $jobGroup

# Get Host Network Adapter 'Intel(R) Ethernet 10G X710 rNDC'
$mgmtNetworkAdapter2 = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'NIC2' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $mgmtNetworkAdapter2 -UplinkPortProfileSet
$mgmtUplinkPortProfileSet -JobGroup $jobGroup

$mgmtNetworkAdapter = @()
$mgmtNetworkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'NIC1' }
$mgmtNetworkAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'NIC2' }
$MgmtLogicalSwitch = Get-SCLogicalSwitch -Name ManagementSwitch
New-SCVirtualNetwork -VMHost $vmHost -VMHostNetworkAdapters $mgmtNetworkAdapter -LogicalSwitch
$MgmtLogicalSwitch -DeployVirtualNetworkAdapters -JobGroup $jobGroup

# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter'
$storageUplinkPortProfileSet = Get-SCUplinkPortProfileSet -Name 'S2D_Storage_UPP'
$storage1Adapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 1' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $storage1Adapter -UplinkPortProfileSet
$storageUplinkPortProfileSet -JobGroup $jobGroup
```

```
# Get Host Network Adapter 'Mellanox ConnectX-4 Lx Ethernet Adapter #2'
$storage2Adapter = Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 2' }
Set-SCVMHostNetworkAdapter -VMHostNetworkAdapter $storage2Adapter -UplinkPortProfileSet
$storageUplinkPortProfileSet -JobGroup $jobGroup

$storageAdapter = @()
$storageAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 1' }
$storageAdapter += Get-SCVMHostNetworkAdapter -VMHost $vmHost | Where-Object
{ $_.ConnectionName -eq 'SLOT 1 Port 2' }
$storageSwitch = Get-SCLogicalSwitch -Name StorageSwitch
New-SCVirtualNetwork -VMHost $vmHost -VMHostNetworkAdapters $storageAdapter -LogicalSwitch
$storageSwitch -DeployVirtualNetworkAdapters -JobGroup $jobGroup

Set-SCVMHost -VMHost $vmHost -JobGroup $jobGroup -RunAsynchronously
```

The above commands need to be executed for each Hyper-V host that will be a part of the cluster and Storage Spaces Direct configuration.

Remove VLAN ID from physical adapter

NOTE: This step is required only if you have configured the VLAN ID on the physical port in an earlier step.

The following PowerShell command needs to be run on each node that will be a part of the cluster:

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 1' -DisplayName 'VLAN ID' -DisplayValue 0
```

Create Hyper-V cluster

As a final step in creating the cluster using Hyper-V hosts that are in the VMM management scope, a cluster can be created along with the Storage Spaces Direct configuration.

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.

The following PowerShell commands help achieve that:

```
$HostGroup = Get-SCVMHostGroup -Name 'SpacesDirectHosts'
$AdminRunAsAccount = Get-SCRunAsAccount -Name 'Administrator'
$jobGroup = (New-Guid).Guid

# Get Host
's2dnode01.s2dlab.local,s2dnode02.s2dlab.local,s2dnode04.s2dlab.local,s2dnode03.s2dlab.local'
$VMHosts = @()
$VMHosts += Get-SCVMHost | Where-Object { $_.Name -eq 's2dnode01.hci.lab' }
$VMHosts += Get-SCVMHost | Where-Object { $_.Name -eq 's2dnode02.hci.lab' }
$VMHosts += Get-SCVMHost | Where-Object { $_.Name -eq 's2dnode03.hci.lab' }
$VMHosts += Get-SCVMHost | Where-Object { $_.Name -eq 's2dnode04.hci.lab' }

$StaticIPAddress = @"172.16.102.55"

Install-SCVMHostCluster -ClusterName "S2DNodeCluster" -JobGroup $jobGroup -RunAsynchronously -
EnableS2D -Credential $AdminRunAsAccount -VMHost $VMHosts -ClusterIPAddress $StaticIPAddress
```

The section [Manual Cluster Deployment and Configuration](#) should be skipped in the VMM-based deployment of cluster and Storage Spaces Direct configuration. Proceed from section [RDMA configuration](#) for additional configuration on the cluster nodes.

Manual cluster deployment and configuration

In the absence of VMM 2016 in the existing data center environment, the cluster deployment and Storage Spaces Direct configuration can be done completely in a manual method. The following subsections provide an overview of these configuration tasks and provide PowerShell commands to complete this 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

For the Windows Server 2016 Storage Spaces Direct HCI cluster deployments, it is necessary that the host OS networking is configured by using a Switch Embedded Team (SET). SET configuration in Windows Server 2016 allows Remote Direct Memory Access (RDMA). It enables virtual adapters with RDMA capability to connect to the same virtual machine switch as the host management and other classes of network traffic.

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 dedicated host management).

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

Fully converged network

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




Figure 24. 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.

Perform the following steps to configure the OS network:

- 1 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
```

Dedicated host management

The following figure provides a high-level overview of VM switch and management adapter configuration in the Storage Spaces Direct Ready Node in a dedicated host management network configuration. This approach includes two different VM switches - one for Storage traffic and the other for host management and VM traffic.




Figure 25. Dedicated host management network configuration

Storage traffic is on the first SET switch with two storage virtual adapters in the management OS attached to it. The physical NIC ports from the Mellanox or QLogic network adapters are used to create this SET. The SFP+ ports or the 1Gb BASE-T from the rNDC are used to create an SET and the host management and VM adapters get attached to this VM switch.

- 1 Run the following command to create VM switches in the SET configuration by using the physical network ports from the Mellanox or QLogic Network adapters for the storage traffic:

```
New-VMSwitch -Name StorageSwitch -AllowManagementOS 0 -NetAdapterName 'SLOT 1 PORT 1','SLOT 1 PORT 2' -Verbose
```

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

- 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 MgmtSwitch -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 StorageSwitch -Passthru | Set-VMNetworkAdapterVlan -Access -VlanId 103 -Verbose
Add-VMNetworkAdapter -ManagementOS -Name 'Storage2' -SwitchName StorageSwitch -Passthru | Set-VMNetworkAdapterVlan -Access -VlanId 104 -Verbose
```

After the network adapters are added to the host OS, static IP addresses can be configured. Retrieve the argument for the *InterfaceAlias* parameter by using the *Get-NetAdapter* cmdlet.

Name	InterfaceAlias
vEthernet (Internal)	vEthernet (Internal)
vEthernet (Storage2)	vEthernet (Storage2)
vEthernet (Storage1)	vEthernet (Storage1)
vEthernet (Management)	vEthernet (Management)
NIC2	NIC2
NIC4	NIC4
NIC3	NIC3
NIC1	NIC1
SLOT 1 Port 2	SLOT 1 Port 2
SLOT 1 Port 1	SLOT 1 Port 1


```
#Host Management Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Management)' -IPAddress 172.16.102.51 -DefaultGateway 172.16.102.1 -PrefixLength 24 -AddressFamily IPv4 -Verbose

#DNS server address
Set-DnsClientServerAddress -InterfaceAlias 'vEthernet (Management)' -ServerAddresses 172.16.102.202

#Storage 1 Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Storage1)' -IPAddress 172.16.103.51 -PrefixLength 24 -AddressFamily IPv4 -Verbose

#Storage 2 Adapter
New-NetIPAddress -InterfaceAlias 'vEthernet (Storage2)' -IPAddress 172.16.104.51 -PrefixLength 24 -AddressFamily IPv4 -Verbose
```

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

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

The assigned IP address configuration can be verified using the following command:

```
Get-NetIPAddress -InterfaceAlias *vEthernet* -AddressFamily IPv4 | Select InterfaceAlias, IPAddress
```

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.

- 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
```

- ❗ **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.
- ❗ **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.

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.

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

```
Test-Cluster -Node S2DNode01, S2DNode02, S2dNode03, S2dNode04 -Include 'Storage Spaces Direct',
'Inventory', 'Network', 'System Configuration'
```

- ❗ **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.

```
New-Cluster -Name S2DSYSTEM -Node S2DNode01, S2DNode02, S2dNode03, S2dNode04 -StaticAddress
172.16.102.55 -NoStorage -IgnoreNetwork 172.16.103.0/24, 172.16.104.0/24 -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.

- ❗ **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 'RDMA Mode' -DisplayValue
'iWarp'
```

```
Set-NetAdapterAdvancedProperty -Name 'SLOT 1 PORT 2' -DisplayName 'RDMA Mode' -DisplayValue 'iWarp'
```

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 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)'
```

- 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:

```
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.

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

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

- 4 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 that use Mellanox CX4 adapters 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.

Table 13. QoS priorities

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

3	Enabled
---	---------

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

- 1 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
```

The argument '3' to the *PriorityValue8021Action* parameter indicates the IEEE 802.1p value and it 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
```

The above command specifies that the transmission algorithm used is Enhanced Transmission Selection (ETS) and the traffic class gets 50% of the bandwidth.

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 (fully converged or dedicated management). In case of a dedicated management network, this value can be higher than 50.

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

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

- 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.

Enabling jumbo frames

Enabling jumbo frames (MTU 9614) could help the workload running on the Storage Spaces Direct cluster by improving the overall read/write performance. You must configure this setting on storage network connections in the host OS and at the switch port.

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

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. For a cluster that has 4 or more nodes, witness configuration is not mandatory.

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: Configuration of cluster witness is mandatory for a 2-node cluster. 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://blogs.msdn.microsoft.com/clustering/2014/03/31/configuring-a-file-share-witness-on-a-scale-out-file-server/>.

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 S2D 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.
- Day 0 OS updates: This is required because there may be cumulative critical updates to the OS after the general availability.
- OS license activation: By default, the OS is installed in the evaluation mode. This needs to be activated immediately after the 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://dell.com/wssdmanuals>.

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](#)
- [Dell EMC Hyper-converged infrastructure](#)

Firewall port requirements

Table 14. 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)
Local Subnet	All Hosts and VMs	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 15. 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
VMM Server FQDN	vmm.s2dlab.local
VMM Administrator credentials	Username: hci\vmadmin Password: <DO NOT WRITE IT DOWN>
SCOM Server FQDN	scom.s2dlab.local
SCOM Administrator credentials	Username: hci\scomadmin Password: <DO NOT WRITE IT DOWN>
WSUS Server FQDN (if needed)	wsus.s2dlab.local

Table 16. Sample checklist

Traffic Class	Purpose	Minimum IP addresses needed	VLAN ID	Tagged / Untagged	IP address space	VLAN IP Addresses
Out of band	Required for OOB management of server nodes and TOR switches	18	100	Untagged	/27	OOB: 172.16.100.1
Host Management	Management of cluster and cluster nodes	17	102	Tagged	/27	TOR1: NA TOR2: NA
Storage 1	SMB traffic	16	103	Tagged	/27	TOR1: NA TOR2: NA
Storage 2	SMB Traffic	16	104	Tagged	/27	TOR1: NA TOR2: NA

Table 17. 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>
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 18. Sample checklist

Fields	Values
Virtual Switch/Logical Switch Name	S2DSwitch
Management Adapter/Logical and VM Network Name	Management
Storage 1 Adapter/Logical and VM Network Name	Storage1
Storage 2 Adapter/ Logical and VM Network Name	Storage2
Uplink Port Profile Name (VMM Only)	S2D_UP
Management IP Pool range (VMM only)	172.16.102.0/27
Storage1 IP pool range (VMM only)	172.16.103.0/27
Storage2 IP pool range (VMM only)	172.16.104.0/27

Table 19. 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

	Host Name	Management IP	Storage1 IP	Storage2 IP	OOB IP	OOB Host name
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
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

VMM preparation

If System Center Virtual Machine Manager (VMM) is used for creating and configuring Microsoft Storage Spaces Direct cluster, VMM needs to be prepared to ensure that the uplink port profile, logical networks, and logical switch required for cluster configuration are already present.

This section covers the steps involved in preparing VMM for cluster creation. For configuration specifics, see the [Deployment checklist](#) section.

Before running any of the commands listed in the below sections, ensure that you update the variables defined below. Values of these variables are used across these different sections.

```
#Variables to populate
$hostGroupName = 'SpacesDirectHosts'
$runAsAccountName = 'Administrator'
$runAsCredential = Get-Credential
$uplinkPortProfileName = 'S2D_UPP'
$logicalSwitchName = 'S2DSwitch'

#LogicalNetwork hash
#Do not update IsManagement and InheritsAddressFromPhysicalNetworkAdapter key values
$logicalNetworks = @(
    #Management Network
    @{
        Name      = 'Management'
        Subnet    = '172.16.102.0/24'
        VLANID    = 102
        IsManagement = $true
        InheritsAddressFromPhysicalNetworkAdapter = $true

        #Used for IP Pool
        StartIP = '172.16.102.51'
        EndIP   = '172.16.102.100'
        DNSServer = '172.16.102.2'
        DNSSuffix = 'HCI.lab'
        GatewayAddress = '172.16.102.1'
        IPPoolName = 'Management-IPPool'
    },
    #Storage 1 Network
    @{
        Name      = 'Storage1'
        Subnet    = '172.16.103.0/24'
        VLANID    = 103
        IsManagement = $false
        InheritsAddressFromPhysicalNetworkAdapter = $false

        #Used for IP Pool
        StartIP = '172.16.103.51'
        EndIP   = '172.16.103.100'
        IPPoolName = 'Storage1-IPPool'
    },
    #Storage 2 Network
    @{
        Name      = 'Storage2'
        Subnet    = "172.16.104.0/24"
        VLANID    = 104
        IsManagement = $false
        InheritsAddressFromPhysicalNetworkAdapter = $false
    }
)
```

```

        #Used for IP Pool
        StartIP = '172.16.104.51'
        EndIP = '172.16.104.100'
        IPPoolName = 'Storage2-IPPool'
    }
)

```

Topics:

- [Create VMM host group](#)
- [Add Run As Account for cluster creation](#)
- [VMM logical network details](#)
- [Create and configure logical and VM networks—fully converged](#)
- [Create and configure logical and VM networks—dedicated management](#)
- [Create IP address pools](#)
- [Create an uplink port profile](#)
- [Create a logical switch](#)

Create VMM host group

A new VM host group must be created to add the newly deployed Hyper-V cluster nodes. This can be done by using the *New-SCHostGroup* cmdlet.

```

#New SC Host group
New-SCVMHostGroup -Name $hostGroupName

```

You can verify the host group creation by using the *Get-SCVMHostGroup* cmdlet.

```

PS C:\> Get-SCVMHostGroup -Name $hostGroupName
AllowUnencryptedTransfers : False
CreationDate              : 3/30/2017 1:41:54 PM
Creator                   : hci\Administrator
Description                :
ID                         : ed4e638e-77e3-48f5-859f-7caef0c94915
InheritNetworkSettings    : True
IsFullyCached             : True
IsRoot                    : False
MarkedForDeletion         : False
ModificationDate          : 3/30/2017 1:41:54 PM
ModifiedBy                : hci\Administrator
Name                      : SpacesDirectHosts
ParentHostGroup           : All Hosts
Path                      : All Hosts\SpacesDirectHosts
ServerConnection          :
Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection

```

Add Run As Account for cluster creation

When using VMM for cluster creation and configuration, a run as account is needed for joining the computers to the domain.

NOTE: PowerShell Commands in this section and the subsequent sections for VMM preparation require System Center VMM PowerShell module. Ensure that the system where these commands are being run has access to the VMM server and has the VMM PowerShell module installed.

```

#Create RunAs Account
New-SCRunAsAccount -Name $runAsAccountName -Credential $runAsCredential

```

The account creation can be verified using the *Get-SCRunAsAccount* cmdlet.

```
PS C:\> Get-SCRunAsAccount -Name $runAsAccountName
Name                : Administrator
UserName            : administrator
Domain              : s2dlab
Enabled              : True
IsBuiltIn            : False
GrantedToList       : {}
UserRoleID           : 75700cd5-893e-4f68-ada7-50ef4668acc6
UserRole             : Administrator
Owner               : hci\Administrator
ObjectType           : RunAsAccount
Accessibility        : Public
IsViewOnly           : False
Description          :
AddedTime            : 3/30/2017 1:31:02 PM
ModifiedTime         : 3/30/2017 1:31:02 PM
MostRecentTask       : Create new RunAs Account
ServerConnection     : Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection
ID                   : 82f86a32-3708-404e-bdd1-6ad5c82a74aa
MarkedForDeletion    : False
IsFullyCached        : True
MostRecentTaskIfLocal : Create new RunAs Account
```

VMM logical network details

For creating VMM logical networks and VM networks, and assigning IP Pools, there are certain settings that the PowerShell cmdlets would require. Commands in the subsequent sections use the details defined in the PowerShell hashtable for configuration. See [VMM Preparation](#) section.

Create and configure logical and VM networks—fully converged

This section provides the steps to implement a fully-converged logical switch in VMM.

For Storage Spaces Direct HCI cluster deployment and configuration, three logical network definitions, which are host management, Storage 1, and Storage 2 networks, must be created in VMM. These logical networks should be created as VLAN-based independent networks. The Deployment Checklist section captures the necessary settings that are needed for creating these logical networks.

For VMM to create and configure the network adapters in the host OS, VM networks must be created and associated with the right logical network and the network site mapped as VM subnet in the VMM fabric.

```
#Create logical network
$hostGroup = Get-SCVMHostGroup -Name $hostGroupName
foreach ($logicalNet in $logicalNetworks)
{
    #VLAN-based independent logical network
    $logicalNetwork = New-SCLogicalNetwork -Name $logicalNet.Name -
LogicalNetworkDefinitionIsolation $true -EnableNetworkVirtualization $false -UseGRE $false -
IsPVLAN $false
    $subnetVlan = New-SCSubnetVlan -Subnet $logicalNet.Subnet -VlanID $logicalNet.VLANID
    $logicalNetDefinition = New-SCLogicalNetworkDefinition -Name "$($logicalNet.Name)_0" -
LogicalNetwork $logicalNetwork -VMHostGroup $hostGroup -SubnetVlan $subnetVlan -
RunAsynchronously

    #Create VM Network
    $vmNetwork = New-SCVMNetwork -Name $logicalNet.Name -LogicalNetwork $logicalNetwork -
IsolationType "VLANNetwork"
    $vmSubnet = New-SCVMSubnet -Name "$($logicalNet.Name)_0" -LogicalNetworkDefinition
$logicalNetDefinition -SubnetVlan $subnetVlan -VMNetwork $vmNetwork
```

The creation of logical networks can be verified by using the Get-ScLogicalNetwork cmdlet.

```
PS C:\> foreach ($logicalNet in $logicalNetworks) { Get-SCLogicalNetwork -Name $logicalNet.Name
| Select Name, IsLogicalNetworkDefinitionIsolated }
Name          IsLogicalNetworkDefinitionIsolated
-----
Management    True
Storage1       True
Storage2       True
```

The creation of VM Networks can be verified by using the Get-SCVMNetwork cmdlet.

```
Name          LogicalNetwork
-----
Management    Management
Storage1       Storage1
Storage2       Storage2
```

Create and configure logical and VM networks—dedicated management

This section provides the necessary steps to implement two logical switches that enable dedicated management scenario using VMM. In this scenario, the integrated 10 GbE ports are chosen to implement a SET switch for management, and the Mellanox or QLogic adapter ports are used to implement the storage virtual switch.

```
#Create management logical network
$hostGroup = Get-SCVMHostGroup -Name $hostGroupName
foreach ($logicalNet in $logicalNetworks)
{
    #VLAN-based independent logical network
    $logicalNetwork = New-SCLogicalNetwork -Name $logicalNet.Name -
LogicalNetworkDefinitionIsolation $true -EnableNetworkVirtualization $false -UseGRE $false -
IsPVLAN $false
    $subnetVlan = New-SCSubnetVlan -Subnet $logicalNet.Subnet -VlanID $logicalNet.VLANID
    $logicalNetDefinition = New-SCLogicalNetworkDefinition -Name "$($logicalNet.Name)_0" -
LogicalNetwork $logicalNetwork -VMHostGroup $hostGroup -SubnetVlan $subnetVlan -
RunAsynchronously

    #Create VM Network
    $vmNetwork = New-SCVMNetwork -Name $logicalNet.Name -LogicalNetwork $logicalNetwork -
IsolationType "VLANNetwork"
    $vmSubnet = New-SCVMSubnet -Name "$($logicalNet.Name)_0" -LogicalNetworkDefinition
$logicalNetDefinition -SubnetVlan $subnetVlan -VMNetwork $vmNetwork
}

#Create IP Pool
foreach ($logicalNet in $logicalNetworks)
{
    # Get Logical Network
    $logicalNetwork = Get-SCLogicalNetwork -Name $logicalNet.Name

    # Get Logical Network Definition
    $logicalNetworkDefinition = Get-SCLogicalNetworkDefinition -Name "$($logicalNet.Name)_0"

    $gateway = @()
    if ($logicalNet.GatewayAddress)
    {
        $gateway += New-SCDefaultGateway -IPAddress $logicalNet.GatewayAddress -Automatic
    }

    New-SCStaticIPAddressPool -Name $logicalNet.IPPoolName -LogicalNetworkDefinition
$logicalNetworkDefinition `
                                -Subnet $logicalNet.Subnet -IPAddressRangeStart
$logicalNet.StartIP `
                                -IPAddressRangeEnd $logicalNet.EndIP -DefaultGateway $gateway `
```

```

        -DNSServer $logicalNet.DNSServer -DNSSuffix "" -DNSSearchSuffix
$logicalNet.DNSSuffix -RunAsynchronously
}

#Create Storage UPP
$storageNetworkDefinition = @()
foreach ($logicalNet in $logicalNetworks)
{
    if (!$logicalNet.IsManagement)
    {
        $storageNetworkDefinition += Get-SCLogicalNetworkDefinition -Name "$
($logicalNet.Name)_0"
    }
}
New-SCNativeUplinkPortProfile -Name $storageUplinkPortProfileName -Description "Storage Uplink
Port profile" -LogicalNetworkDefinition $storageNetworkDefinition -EnableNetworkVirtualization
$false -LBFOLoadBalancingAlgorithm "HostDefault" -LBFOTeamMode "SwitchIndependent" -
RunAsynchronously

#Create Management UPP
$managementNetworkDefinition = @()
foreach ($logicalNet in $logicalNetworks)
{
    if ($logicalNet.IsManagement)
    {
        $managementNetworkDefinition += Get-SCLogicalNetworkDefinition -Name "$
($logicalNet.Name)_0"
    }
}
New-SCNativeUplinkPortProfile -Name $managementUplinkPortProfileName -Description "Management
Uplink Port profile" -LogicalNetworkDefinition $managementNetworkDefinition -
EnableNetworkVirtualization $false -LBFOLoadBalancingAlgorithm "HostDefault" -LBFOTeamMode
"SwitchIndependent" -RunAsynchronously

#Logical Switches - both management and storage
$managementLogicalSwitch = New-SCLogicalSwitch -Name $managementLogicalSwitchName -Description
'Management Logical Switch' -EnableSriov $false -SwitchUplinkMode 'EmbeddedTeam' -
MinimumBandwidthMode 'Weight'
$storageLogicalSwitch = New-SCLogicalSwitch -Name $storageLogicalSwitchName -Description
'Storage Logical Switch' -EnableSriov $false -SwitchUplinkMode 'EmbeddedTeam'

# Get Native Uplink Port Profile - Management
$managementNativeUppVar = Get-SCNativeUplinkPortProfile -Name $managementUplinkPortProfileName
$managementUppSetVar = New-SCUplinkPortProfileSet -Name $managementUplinkPortProfileName -
LogicalSwitch $managementLogicalSwitch -NativeUplinkPortProfile $managementNativeUppVar -
RunAsynchronously

# Get Native Uplink Port Profile - Storage
$storageNativeUppVar = Get-SCNativeUplinkPortProfile -Name $storageUplinkPortProfileName
$storageUppSetVar = New-SCUplinkPortProfileSet -Name $storageUplinkPortProfileName -
LogicalSwitch $storageLogicalSwitch -NativeUplinkPortProfile $storageNativeUppVar -
RunAsynchronously

foreach ($logicalNet in $logicalNetworks)
{
    # Get VM Network
    $vmNetwork = Get-SCVMNetwork -Name $logicalNet.Name

    # Get VMSubnet
    $vmSubnet = Get-SCVMSubnet -Name "$($logicalNet.Name)_0"

    if (-not ($logicalNet.IsManagement))
    {
        # Get Static IP Address Pool
        $ipV4Pool = Get-SCStaticIPAddressPool -Name $logicalNet.IPPoolName

        New-SCLogicalSwitchVirtualNetworkAdapter -Name $logicalNet.Name -UplinkPortProfileSet
$storageUppSetVar -RunAsynchronously -VMNetwork $vmNetwork -VMSubnet $vmSubnet -
IsUsedForHostManagement $logicalNet.IsManagement -InheritsAddressFromPhysicalNetworkAdapter

```

```
$logicalNet.InheritsAddressFromPhysicalNetworkAdapter -IPv4AddressType "Static" -
IPv6AddressType "Dynamic" -IPv4AddressPool $ipV4Pool
}
else
{
    $ipV4Pool = Get-SCStaticIPAddressPool -Name $logicalNet.IPPoolName
    New-SCLogicalSwitchVirtualNetworkAdapter -Name $logicalNet.Name -UplinkPortProfileSet
$managementUppSetVar -RunAsynchronously -VMNetwork $vmNetwork -VMSubnet $vmSubnet -
IsUsedForHostManagement $logicalNet.IsManagement -InheritsAddressFromPhysicalNetworkAdapter
$logicalNet.InheritsAddressFromPhysicalNetworkAdapter -IPv4AddressType "Static" -
IPv6AddressType "Dynamic" -IPv4AddressPool $ipV4Pool -Verbose
}
}
```

Create IP address pools

As a part of the Storage Spaces Direct cluster creation and configuration, VMM creates the network adapters in the host OS for management, Storage 1, and Storage 2, VMM assigns IP addresses for each VM adapter in the management OS. You must supply the necessary IP address for each subnet as an IP address pool in VMM. VMM retrieves the IP address from the pool and configures the addresses on each VM network adapter in the management OS.

```
#Create IP Pool
foreach ($logicalNet in $logicalNetworks)
{
    # Get Logical Network
    $logicalNetwork = Get-SCLogicalNetwork -Name $logicalNet.Name

    # Get Logical Network Definition
    $logicalNetworkDefinition = Get-SCLogicalNetworkDefinition -Name "$($logicalNet.Name)_0"

    $gateway = @()
    if ($logicalNet.GatewayAddress)
    {
        $gateway += New-SCDefaultGateway -IPAddress $logicalNet.GatewayAddress -Automatic
    }

    New-SCStaticIPAddressPool -Name $logicalNet.IPPoolName -LogicalNetworkDefinition
$logicalNetworkDefinition `
        -Subnet $logicalNet.Subnet -IPAddressRangeStart
$logicalNet.StartIP `
        -IPAddressRangeEnd $logicalNet.EndIP -DefaultGateway $gateway `
        -DNSServer $logicalNet.DNSServer -DNSSuffix "" -DNSSearchSuffix
$logicalNet.DNSSuffix -RunAsynchronously
}
```

Verify the IP pool creation using the Get-SCStaticIPAddressPool cmdlet.

```
PS C:\> foreach ($logicalNet in $logicalNetworks)
>> {
>>     Get-SCStaticIPAddressPool -Name "$($logicalNet.Name)-IPPool" | Select Name,
IPAddressRange*
>> }
>>
```

Name	IPAddressRangeStart	IPAddressRangeEnd
Management-IPPool	172.16.102.51	172.16.102.100
Storage1-IPPool	172.16.103.51	172.16.103.100
Storage2-IPPool	172.16.104.51	172.16.104.100

Create an uplink port profile

The Storage Spaces Direct HCI cluster requires a Switch Embedded Team (SET) configuration for which VMM should have an uplink port profile created with relevant configurations such as teaming mode set to Switch Independent and Load Balancing algorithm set to Host Default. The network configuration in the uplink port profile should have the relevant network sites selected as support networks. For example, this network should include all logical networks created in the section [Create and configure logical and VM networks](#).

```
#Create UPP
$definition = @()
foreach ($logicalNet in $logicalNetworks)
{
    $definition += Get-SCLogicalNetworkDefinition -Name "$($logicalNet.Name)_0"
}
New-SCNativeUplinkPortProfile -Name $uplinkPortProfileName -Description "Uplink Port profile" -
LogicalNetworkDefinition $definition -EnableNetworkVirtualization $false -
LBFOLoadBalancingAlgorithm "HostDefault" -LBFOTeamMode "SwitchIndependent" -RunAsynchronously
```

The creation of uplink port profile can be verified using the *Get-SCNativeUplinkPortProfile* cmdlet.

```
PS C:\> Get-SCNativeUplinkPortProfile -Name S2D_UPP
Name                               : S2D_UPP
Description                         : Uplink Port profilea
EnableNetworkVirtualization        : False
LogicalNetworkDefinitions          : {Storage2_0, Management_0, Storage1_0}
LBFOLoadBalancingAlgorithm         : HostDefault
LBFOTeamMode                       : SwitchIndependent
ServerConnection                   :
Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection
ID                                 : b4508131-16a7-4ad3-90c7-e82a78943f57
IsViewOnly                         : False
ObjectType                         : NativeUplinkPortProfile
MarkedForDeletion                  : False
IsFullyCached                      : True
```

Create a logical switch

After the VMM networking artifacts such as logical networks, VM networks, IP address pools, and uplink port profile are created, a logical switch must be configured. A logical switch combines all the other networking artifacts into a deployable entity. When performing a bare metal deployment, VMM uses the logical switch definition to identify the configuration required for the target host networking.

```
#Logical Switch
$logicalSwitch = New-SCLogicalSwitch -Name $logicalSwitchName -Description "" -EnableSriov
$false -SwitchUplinkMode "EmbeddedTeam" -MinimumBandwidthMode "Weight"

# Get Native Uplink Port Profile
$nativeUppVar = Get-SCNativeUplinkPortProfile -Name $uplinkPortProfileName
$uppSetVar = New-SCUplinkPortProfileSet -Name $uplinkPortProfileName -LogicalSwitch
$logicalSwitch -NativeUplinkPortProfile $nativeUppVar -RunAsynchronously

foreach ($logicalNet in $logicalNetworks)
{
    # Get VM Network
    $vmNetwork = Get-SCVMNetwork -Name $logicalNet.Name

    # Get VMSubnet
    $vmSubnet = Get-SCVMSubnet -Name "$($logicalNet.Name)_0"

    if (-not ($logicalNet.IsManagement))
    {
        # Get Static IP Address Pool
        $ipV4Pool = Get-SCStaticIPAddressPool -Name $logicalNet.IPPoolName

        New-SCLogicalSwitchVirtualNetworkAdapter -Name $logicalNet.Name -UplinkPortProfileSet
$uppSetVar -RunAsynchronously -VMNetwork $vmNetwork -VMSubnet $vmSubnet -
```

```

IsUsedForHostManagement $logicalNet.IsManagement -InheritsAddressFromPhysicalNetworkAdapter
$logicalNet.InheritsAddressFromPhysicalNetworkAdapter -IPv4AddressType "Static" -
IPv6AddressType "Dynamic" -IPv4AddressPool $ipV4Pool
}
else
{
    New-SCLogicalSwitchVirtualNetworkAdapter -Name $logicalNet.Name -UplinkPortProfileSet
$suppSetVar -RunAsynchronously -VMNetwork $vmNetwork -VMSubnet $vmSubnet -
IsUsedForHostManagement $logicalNet.IsManagement -InheritsAddressFromPhysicalNetworkAdapter
$logicalNet.InheritsAddressFromPhysicalNetworkAdapter -IPv4AddressType "Dynamic" -
IPv6AddressType "Dynamic" -Verbose
}
}

```

To get the details of the new logical switch, use the Get-SCLogicalSwitch cmdlet.

```

PS C:\> Get-SCLogicalSwitch -Name S2dSwitch
Name                : S2DSwitch
Description         :
EnableSriov         : False
MinimumBandwidthMode : Weight
UplinkMode          : EmbeddedTeam
VirtualSwitchExtensions : {}
ManagedByMSNetworkController : False
EnablePacketDirect  : False
ServerConnection    :
Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection
ID                  : 9efe8634-5eed-4ddb-a141-753ddca21790
IsViewOnly          : False
ObjectType          : LogicalSwitch
MarkedForDeletion    : False
IsFullyCached       : True

```