



IMPLEMENTING WINDOWS SERVER 2019 AZURE STACK HCI USING HPE PROLIANT SERVERS



CONTENTS

1.0	Introduction.....	3
2.0	HPE Azure Stack HCI solutions.....	3
2.1	HPE hybrid.....	3
2.2	HPE all-flash.....	3
3.0	Configuration best practices.....	4
3.1	Environment preparation.....	4
3.2	Network design.....	5
3.3	Deploying hyperconverged solutions.....	12
3.4	Deploying converged solutions.....	14
4.0	Troubleshooting.....	15
4.1	Cluster verification.....	15
4.2	Disk management.....	15
4.3	Disk rebuild.....	16
4.4	Disk clean.....	16
4.5	Cluster management.....	16
4.6	Network management.....	16
4.7	Remove Storage Spaces Direct.....	17
5.0	Additional resources.....	17
Appendix 1: Performance validation method.....		17
A1.1	VMFleet for Storage Spaces Direct.....	17
A1.2	How to monitor the performance of VMFleet.....	18
Appendix 2: Windows Admin Center (WAC) extension for Azure Stack HCI.....		21
A2.1	WAC Extension preview.....	21



1.0 INTRODUCTION

HPE's Azure Stack HCI solution uses Microsoft-validated HPE ProLiant Gen9 and Gen10 servers to create highly reliable Windows Server 2019 cluster to run virtualized workloads on-premises with optimal performance. Azure Stack HCI solution utilizes Windows Server 2019 Storage Spaces Direct (S2D), a software-defined storage feature, to combine individual server's local-attached storage devices such as SATA and SAS HDDs, SSDs, NVMe devices, and Persistent Memory to create cluster shared disk volumes. Distributed among several cluster nodes, these volumes can withstand disk, node, and network failures, and continue to operate normally without downtime. Different levels of resiliency can be achieved using mirror, parity, or a combination of the two. S2D is mainly targeted at cloud applications, Hyper-V, storage for backup and replication, and hyperconverged/converged infrastructures.

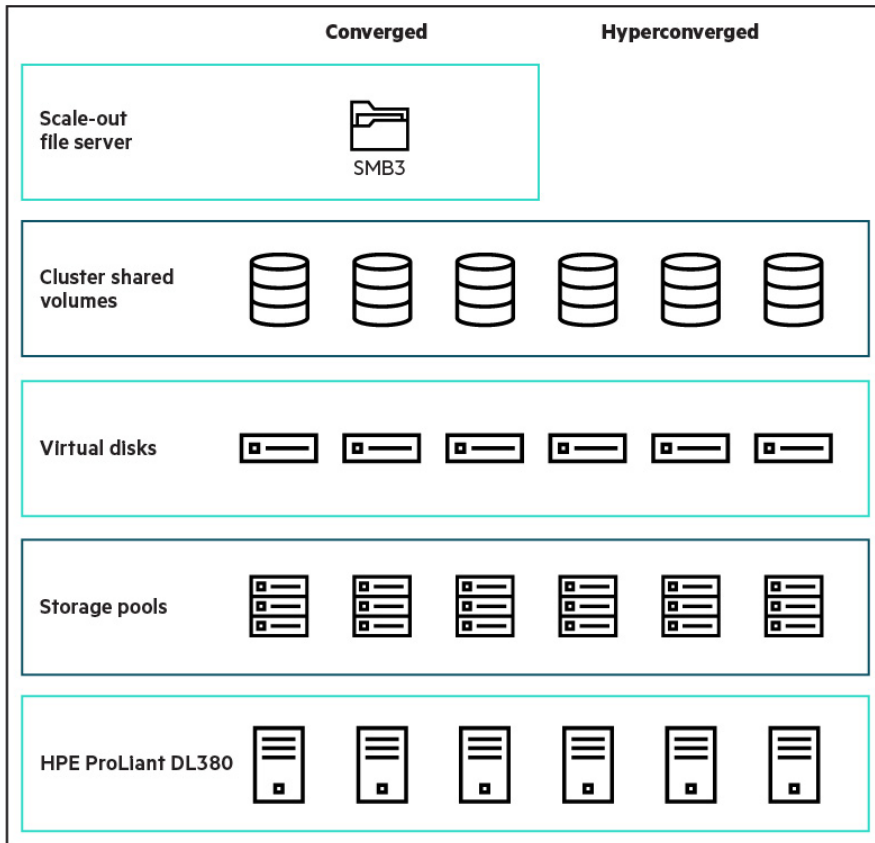


FIGURE 1. Converged and hyperconverged environments

2.0 HPE AZURE STACK HCI SOLUTIONS

HPE Azure Stack HCI solutions are configured to deliver solid performance, improved reliability, and high availability. They can be implemented as 2 to 16 node configurations, providing you with the ability to deploy an engineered, tested, and validated private cloud environment to meet your most demanding workloads. All HPE Azure Stack HCI solutions are available as hyperconverged or converged solutions.

2.1 HPE hybrid

HPE hybrid solutions start with two nodes of HPE ProLiant servers using RDMA-capable network adapters to provide high performance with minimum processor utilization. In order to make use of HPE solutions using RDMA networking, an RDMA compatible networking infrastructure is required. HPE hybrid solutions combine different types of drives (HDD/SSD/NVMe) to allow read/write caching and are available in both LFF and SFF form factors.

2.2 HPE all-flash

HPE all-flash solutions are configured entirely with either SSD or NVMe flash drives and 25Gb or faster RDMA network adapters to support performance-intensive workloads with HPE ProLiant and HPE Synergy Gen9/Gen10 servers.



3.0 CONFIGURATION BEST PRACTICES

This section provides some best practices for configuring Storage Spaces Direct. The procedures outlined in the examples are based on a 4-node HPE DL380 Gen10 SFF configuration. However, the same procedures, with some minor adjustments, can be applied to other configurations.

3.1 Environment preparation

3.1.1 System configuration

1. For each HPE DL380 server, install the SSD/HDD drives, network adapters, and storage controller as shown in the Figure 2.




FIGURE 2. HPE DL380 Gen10 front/rear view

2. The SSD and HDD drives intended for use in your S2D storage pool should be placed in zones that are physically approximate for the ease of management and maintenance. In the example above showing DL380 Gen10 SFF solution, all the S2D drives are placed in the front drive slots. The S2D storage pool drives must be connected to one or more Smart Array Controller running in HBA mode. Connecting S2D drives to Smart Array Controller running in non-HBA mode is not supported.
3. For OS boot drives, you can choose hot-pluggable SSD or HDD located in server storage cage, or non-hot-pluggable M.2 SATA SSD located inside server chassis attached to the embedded Smart Array software RAID controller, depending on your design choices and system capabilities. You can choose to create either a RAID boot volume or a simple AHCI boot volume, though RAID 1 is preferred for data redundancy. Boot drives must be connected to a Smart Array Controller or an onboard SATA controller that is not connected to any S2D storage pool drives. In the example above for DL380 Gen10 SFF solution, a recommended configuration is to place boot drives in rear cage that is connected to the onboard S100i controller which can be configured as a RAID volume by configuring UEFI options.

3.1.2 Firmware update

Download the latest [Service Pack for ProLiant \(SPP\)](#) and boot from this ISO image file to update the system firmware.

3.1.3 Storage controller configuration

1. Make sure each node has the default RBSU settings.
2. For Gen10 storage controllers, HBA mode is enabled by default.
 - a. S2D drives may only reside on a storage controller configured exclusively in HBA mode.
 - b. S2D drives cannot reside on a storage controller configured in either RAID or mixed mode (RAID logical drives and HBA physical drives simultaneously).
3. For Gen9 storage controllers, enable HBA mode on the storage controllers as follows:
 - a. Power on the server and press F10 during POST to enter the intelligent provisioning utility.
 - b. On the next screen, select HPE Smart Storage Administrator.
 - c. From the main HPE Smart Storage Administrator screen, select HPE Smart Array controller that is connected to S2D drives.
 - d. Under Actions, click the Configure button.
 - e. If the controller is already in HBA mode, you will see the Enable RAID Mode button. In this case, skip to Step g.
 - f. If the controller is in RAID mode, select Enable HBA Mode, then click OK and Finish.
 - g. Click the X button in the upper right-hand corner followed by OK to exit the application.
 - h. Click the green power icon in the upper right-hand corner followed by Reboot.
4. Use HPE SSA to create a RAID 1 volume with the boot drives.



3.1.4 Operating system installation and configuration

1. Install Windows Server 2019 Datacenter edition on the RAID volume.
2. Apply the Service Pack for ProLiant (SPP). Make sure to install the chipset, HPE iLO, network, and storage drivers. All others are optional. Reboot the system after completion.
3. Use PowerShell to install the required Windows features:

```
PS C:\> Install-WindowsFeature -Name File-Services, Failover-Clustering, Hyper-V, Data-Center-Bridging -IncludeManagementTools -Restart
```
4. Run Windows Update.

3.2 Network design

3.2.1 Network connections

Connect each server node to switches as shown in Figure 3. For added resiliency, we recommend dual switches per rack and dual network adapters per server. Each network adapter has its ports connected to both switches. If switch A or B fails, the S2D network will continue to operate using the remaining healthy switch. If one network adapter fails in a node, the S2D network will continue normal operation using the remaining healthy network adapter.




FIGURE 3. Recommended network configuration

For configuration consisting of only two nodes without an external switch, connect the two servers using direct cable connections as shown in Figure 4.




FIGURE 4. Two-node network configuration



3.2.2 Active Directory

Configure each node to join a Windows Active Directory domain. For the sample deployments in this document, we assume the following network configuration.

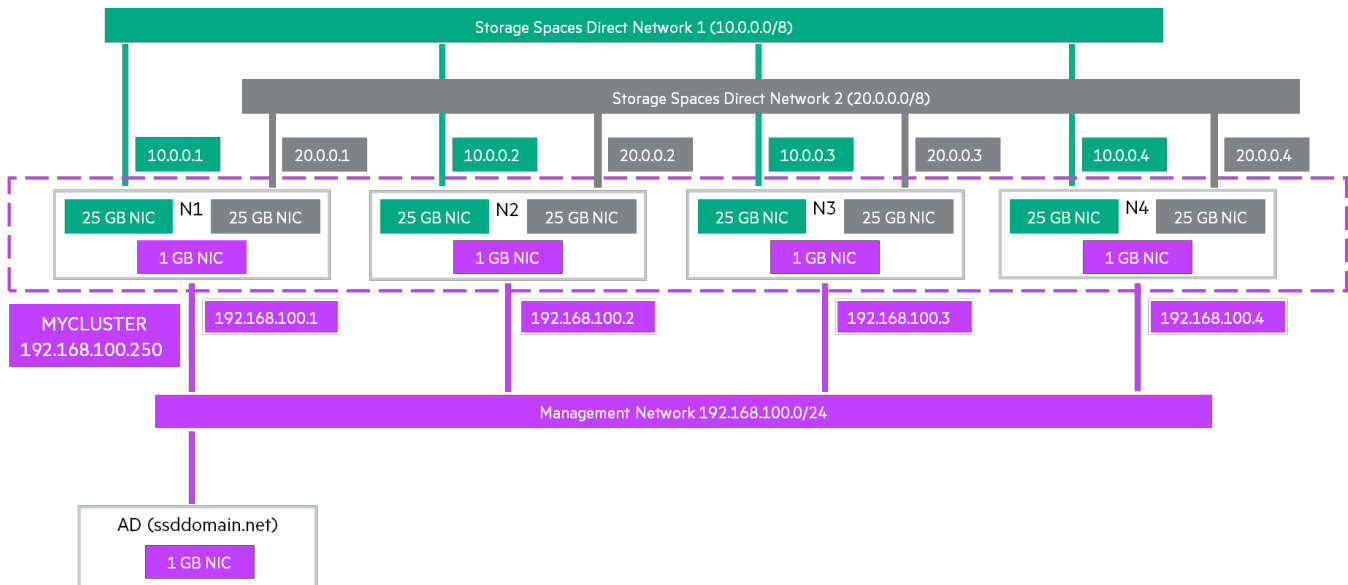


FIGURE 5. Storage Spaces Direct network configuration diagram

3.2.3 RDMA networking options

The remote direct memory access (RDMA) feature in HPE network adapters allows lower latency, reduces processor read/write cycles, and improves performance. RDMA can be implemented using either RDMA over Converged Ethernet (RoCE) or iWARP. RoCEv2 implementations require a lossless or lossy network configuration on the switch while iWARP implementations require minimal configuration and are easier to deploy.

3.2.3.1 iWARP networking

For iWARP deployments, set the RDMA Operational Mode in RBSU and Device Manager -> Advanced Properties as shown in Figures 6 and 7.

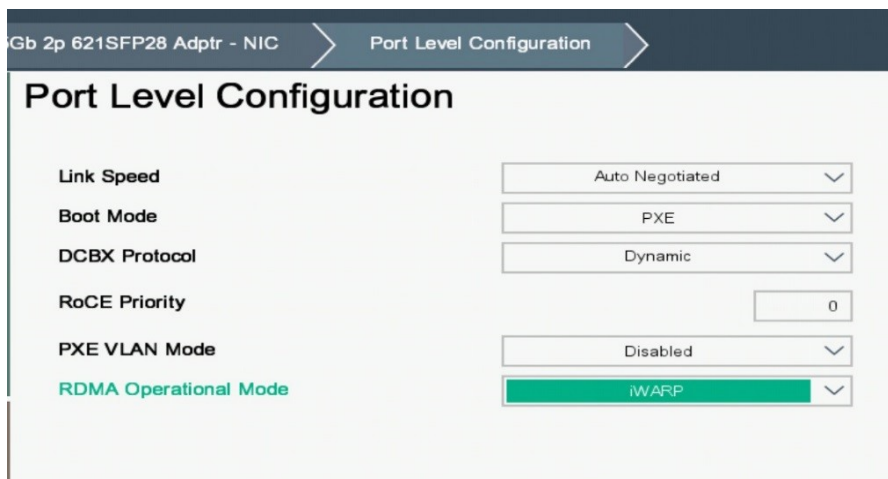


FIGURE 6. iWARP setting in RBSU



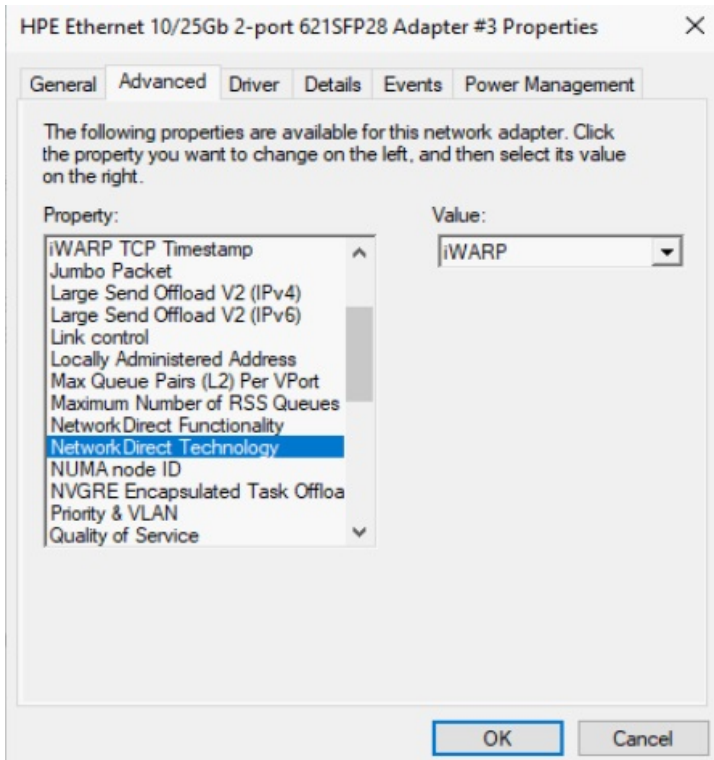


FIGURE 7. iWARP setting in Control Panel

3.2.3.2 RoCEv2 networking

For RoCEv2 deployments, configure Virtual LAN (VLAN) and Priority Flow Control (PFC) on the switches and network adapters.

Switch configuration

Use the following procedure for the HPE FlexFabric 5900AF-48XG-4QSFP+ 48-port switch:

1. Use an RJ-45 to RS-232 cable to connect the console port of the switch to the serial port of your PC.




FIGURE 8. Switch connection

2. Open a terminal emulation program (for example, PuTTY) on your PC.
3. Set your serial connection to 9600bps/8-N-1.
4. The following screen will appear. Press Enter to continue.






FIGURE 9. Set your serial connection on PuTTY

- For each port, create and assign VLAN and Priority Flow Control values. In the following example, we use VLAN 4 and a Priority Flow Control value of 3, but the actual values will depend on your particular network configuration.

```
<HPE>system-view
[HPE]vlan 4
[HPE-vlan4]quit
[HPE]interface Ten-GigabitEthernet 1/0/1
[HPE-Ten-GigabitEthernet1/0/1]port link-type trunk
[HPE-Ten-GigabitEthernet1/0/1]port trunk permit vlan 4
[HPE-Ten-GigabitEthernet1/0/1]priority-flow-control auto
[HPE-Ten-GigabitEthernet1/0/1]priority-flow-control no-drop dot1p 3
[HPE-Ten-GigabitEthernet1/0/1]qos trust dot1p
[HPE-Ten-GigabitEthernet1/0/1]quit
```

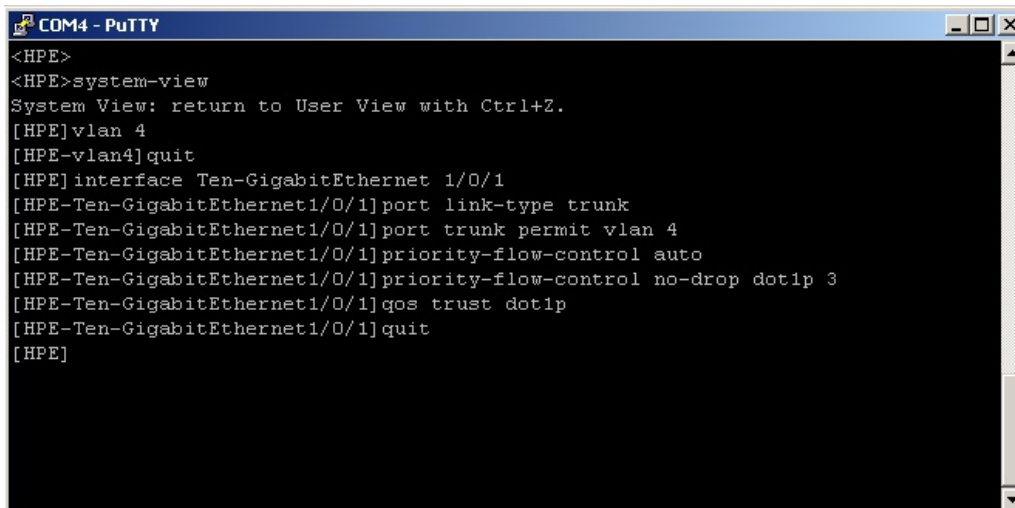


FIGURE 10. Sample parameters for VLAN 4 setting

- Repeat Step 5 for each S2D node port interface on the switch (that is, Ten-GigabitEthernet 1/0/1 to Ten-GigabitEthernet 1/0/16).
- After all ports are configured, use “display vlan” command (in our example, “display vlan 4”) to verify that the required ports are tagged properly.



Use the following commands for the HPE StoreFabric SN2700M 100GbE switch:

```
switch (config) # vlan 4
switch (config) # interface ethernet 1/1 switchport mode trunk
switch (config) # interface ethernet 1/1 switchport trunk allowed-vlan 4
#Repeat this command for ports 1/1 to 1/32
switch (config) # interface ethernet 1/32 switchport mode trunk
switch (config) # interface ethernet 1/32 switchport trunk allowed-vlan 4
switch (config) # interface ethernet 1/1-1/32 traffic-class 3 congestion-control ecn minimum-absolute 150 maximum-absolute 1500
switch (config) # interface ethernet 1/1-1/32 ingress-buffer iPort.pg6 bind switch-priority 6
switch (config) # interface ethernet 1/1-1/32 ingress-buffer iPort.pg3 bind switch-priority 3
switch (config) # interface ethernet 1/1-1/32 ingress-buffer iPort.pg3 map pool iPool1 type lossless reserved 67538 xoff 18432 xon 18432 shared alpha 2
switch (config) # interface ethernet 1/1-1/32 ingress-buffer iPort.pg6 map pool iPool0 type lossy reserved 10240 shared alpha 8
switch (config) # interface ethernet 1/1-1/32 egress-buffer ePort.tc3 map pool ePool1 reserved 1500 shared alpha inf
switch (config) # interface ethernet 1/1-1/32 traffic-class 6 dcb ets strict
switch (config) # interface ethernet 1/1-1/32 qos trust L3
switch (config) # dcb priority-flow-control enable force
switch (config) # dcb priority-flow-control priority 3 enable
switch (config) # interface ethernet 1/1-1/32 dcb priority-flow-control mode on force
```

Network adapter configuration

Configure PFC and create an SMB Direct policy to work with RDMA using PowerShell as follows:

```
#Set QoS policy for SMB-Direct
PS> Set-NetQosDcbxSetting -Willing $false
PS> New-NetQosPolicy "SMB" -SMB -PriorityValue8021Action 3
PS> New-NetQosPolicy "DEFAULT" -Default -PriorityValue8021Action 0
#Enable PFC 3
PS> Enable-NetQosFlowControl -Priority 3
PS> Disable-NetQosFlowControl -Priority 0,1,2,4,5,6,7
#Apply the policy to all adapters
PS> Enable-NetAdapterQos -InterfaceAlias "Embedded FlexibleLOM 1 Port 1", "Embedded FlexibleLOM 1 Port 2"
PS> Enable-NetAdapterQos -InterfaceAlias "PCIe Slot 2 Port 1", "PCIe Slot 2 Port 2"
#Configure VLAN ID 4
PS> Set-NetAdapter -Name "Embedded FlexibleLOM 1 Port 1" -VlanID 4
PS> Set-NetAdapter -Name "Embedded FlexibleLOM 1 Port 2" -VlanID 4
PS> Set-NetAdapter -Name "PCIe Slot 2 Port 1" -VlanID 4
PS> Set-NetAdapter -Name "PCIe Slot 2 Port 2" -VlanID 4
#Assign 95% bandwidth for SMB-Direct
PS> New-NetQosTrafficClass "SMB" -Priority 3 -BandwidthPercentage 95 -Algorithm ETS
```



NOTE

“Embedded FlexibleLOM 1 Port x” and “PCIe Slot x Port x” are the friendly names of the RDMA network adapters.

3.2.3.3 RDMA verification

Use the Performance Monitor to verify RDMA is set up correctly. The performance counters RDMA Activity and SMB Client Shares can list relevant objects such as active RDMA connections, inbound bytes per second, and SMB data bytes per second.

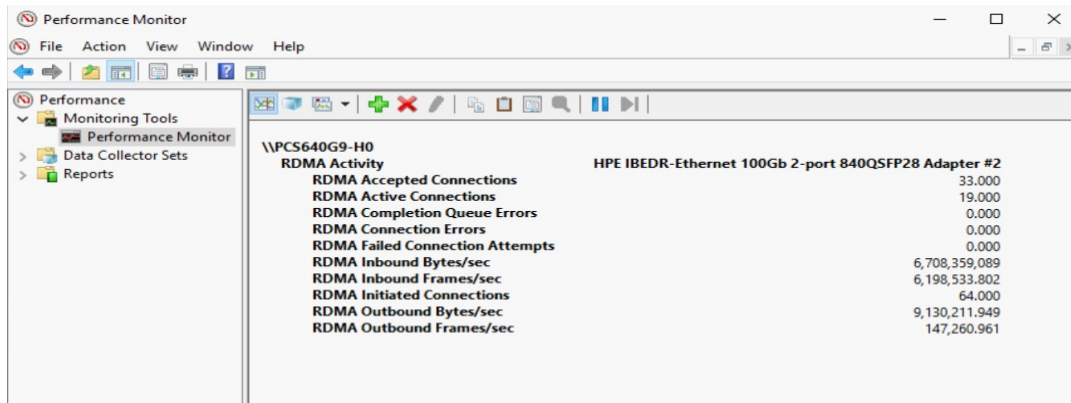


FIGURE 11. Performance Monitor

Microsoft provides a PowerShell script Test-RDMA and a cmdlet Validate-DCB to diagnose RDMA configurations.

Download Test-RDMA.ps1 from github.com/microsoft/SDN/blob/master/Diagnostics/Test-Rdma.ps1 to test port traffic.

Use “Get-NetAdapter” to retrieve the interface ifIndex of the Network port, and run the following command:

```
PS> Test-Rdma.ps1 -IfIndex 7 -IsRoCE $True -RemoteIp 172.16.101.2 (as shown below in Figure 12)
```

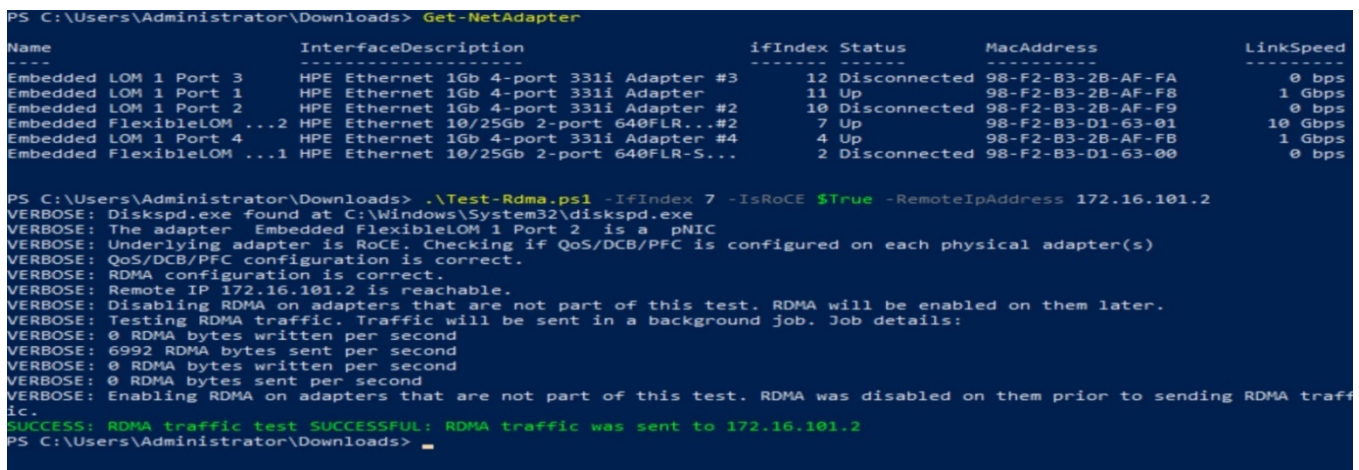


FIGURE 12. Test-RDMA script

SUCCESSFUL transmissions will display green messages. Otherwise, the script will provide steps to troubleshoot your connection.

The Validate-DCB cmdlet can be installed from the PowerShell gallery using Install-Module

```
PS> Install-Module -Name Validate-DCB
```

```
PS> Validate-DCB -LaunchUI
```

In the following wizard, enter information about your S2D cluster and save the settings in Config.ps1



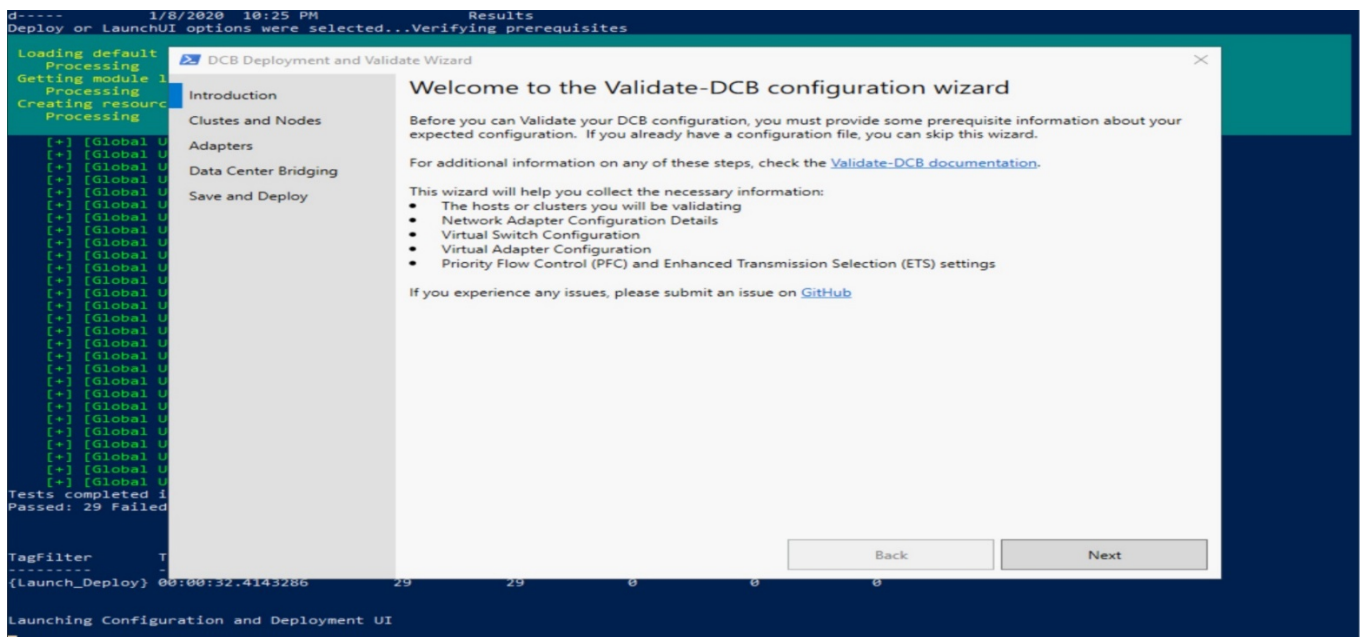


FIGURE 13. Validate-DCB

Execute Validate-DCB with parameter `-ConfigFilePath Config.ps1`

3.2.4 HPE Persistent Memory

HPE Persistent Memory modules featuring Intel® Optane™ DC Persistent Memory technology provides faster data access at a reasonable price point, which makes it an ideal S2D storage option for data-intensive workload. Servers equipped with HPE persistent memory, also known as PMEM, need to follow “Population Guidelines for HPE Persistent Memory” as documented in [Server Memory and Persistent Memory Population Rules for HPE Gen10 Servers with Intel® Xeon® Scalable Processors](#). Server firmware and drivers must be updated to Service Pack for ProLiant (SPP) version 2019.09.0 or newer.

Before using HPE PMEM for S2D storage pool, please follow Microsoft’s instruction to deploy persistent memory: docs.microsoft.com/en-us/windows-server/storage/storage-spaces/deploy-pmem. Below is an example of using PowerShell to configure persistent memory:

1. Use “Get-PmemPhysicalDevice” to check PMEM status. For each PMEM device, HealthStatus should be Healthy, OperationalStatus should be OK.

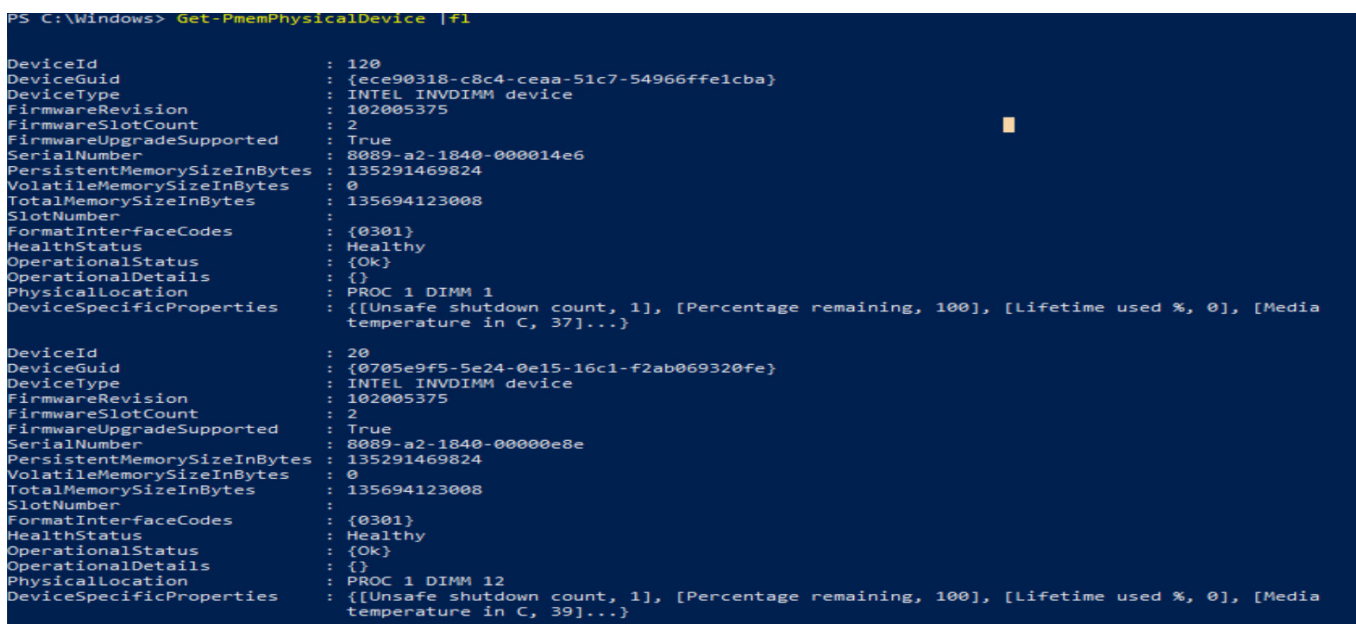


FIGURE 14. Checking PMEM



In order to use PMEM as storage devices for use in Storage Spaces Direct, the PMEM Modules must be configured in App Direct Mode. See HPE Persistent Memory modes section in the in the [HPE Persistent Memory User Guide](#) for additional information about the available PMEM modes. If you find that your PMEM status is not OK or Healthy, this is likely caused by the PMEM modules not being configured in the correct mode for use by Storage Spaces Direct. To resolve this, you must set a “goal configuration” as described in the Setting the goal configuration of the [HPE Persistent Memory User Guide](#). We configure the PMEM devices for App Direct Mode using the embedded UEFI Shell. This can be accomplished by following the instructions described in the Setting the goal configuration with ipmctl section of the [HPE Persistent Memory User Guide](#). To accomplish this requires rebooting the system, pressing F9 and then going to “System Utilities” → “Embedded Applications” → “Embedded UEFI Shell”; then run following command at UEFI shell:

```
Shell> ipmctl create -goal PersistentMemoryType=AppDirectNotInterleaved
```

2. Once PMEM status is OK and Healthy, run “New-PmemDisk” command to create new persistent memory disk:

```
PS C:\> Get-PmemUnusedRegion

RegionId TotalSizeInBytes DeviceId
-----
1        135291469824 {20}
2        135291469824 {120}

PS C:\> Get-PmemUnusedRegion | New-PmemDisk
Creating new persistent memory disk. This may take a few moments.
Creating new persistent memory disk. This may take a few moments.
```

FIGURE 15. Create PMEM disk

3. At this point, new PMEM disk has been created. Use “Get-PmemDisk” to check PMEM disk status:

```
PS C:\> Get-PmemDisk

DiskNumber Size HealthStatus AtomicityType CanBeRemoved PhysicalDeviceIds UnsafeShutdownCount
-----
6          126 GB Healthy None True {20} 1
7          126 GB Healthy None True {120} 1
```

FIGURE 16. PMEM status

4. In “Get-PhysicalDisk”, the newly created PMEM disks should show up as other regular SSD or HDD devices and ready to participate S2D storage pool:

```
PS C:\> Get-PhysicalDisk

DeviceId FriendlyName SerialNumber MediaType CanPool OperationalStatus HealthStatus
-----
6 Persistent memory disk 03018089246e7d08dfbbfb4383eb8c0511d2990d SCM True OK Healthy
0 HPE LOGICAL VOLUME Unspecified False OK Healthy
1 HP M0000800JWFWP 28E0A1JRTLRE1807 SSD True OK Healthy
3 HP M0000800JWFWP 28E0A1JMTLRE1807 SSD True OK Healthy
2 HP M0000800JWFWP 28E0A1JFTLRE1807 SSD True OK Healthy
5 Persistent memory disk 03018089107cea9017aa564381dab57862302ac8 SCM True OK Healthy
4 HP M0000800JWFWP 28E0A1JSTLRE1807 SSD True OK Healthy
```

FIGURE 17. Physical Disks with PMEM devices

3.3 Deploying hyperconverged solutions

In a hyperconverged environment, the compute and storage components are located in the same cluster. For this reason, the recommended hardware requirements are higher in order to handle the load. Storage Spaces Direct volumes are accessible from all the nodes under a local directory named \ClusterStorage.






FIGURE 18. Hyperconverged solution

3.3.1 Sample procedure

The following procedure creates a 3-way mirror volume on a sample 4-node cluster. For more information about different volume types, see the following document: docs.microsoft.com/en-us/windows-server/storage/storage-spaces/plan-volumes. These commands should be executed from one of the cluster nodes using the domain administrator account.

1. Create cluster

Create a cluster named MYCLUSTER using the four nodes N1, N2, N3, N4, and assign an unused IP address to the cluster.

```
PS C:\> New-Cluster -Name MYCLUSTER -Node N1, N2, N3, N4 -NoStorage -StaticAddress 192.168.100.250
```

Assign the file share witness as follows:

```
PS C:\> Set-ClusterQuorum -FileShareWitness \\{Witness Server Name}\{Share Name}
```

2. Enable Storage Spaces Direct

Set the storage system to Storage Spaces Direct mode with the following PowerShell command.

```
PS C:\> Enable-ClusterStorageSpacesDirect
```

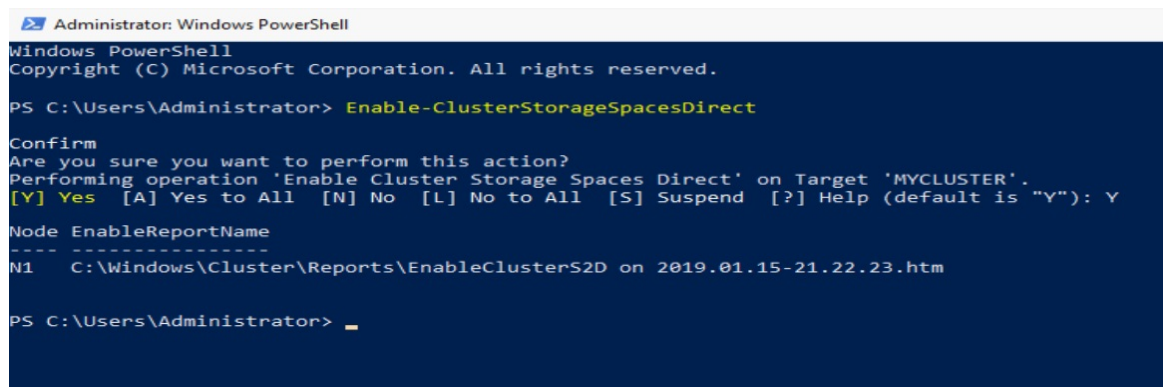


FIGURE 19. Enable Storage Spaces Direct

3. Create disk volume

Volumes in Storage Spaces Direct provide resiliency, using the following mechanisms: **mirroring** and **parity**. We strongly recommend use of **mirroring** for optimal performance. The following example shows how to create a 1 TB mirror volume on the Storage Spaces Direct pool named “S2D*”:

```
PS C:\> New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName MyDisk -FileSystem CSVFS_ReFS -Size 1TB or
PS C:\> New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName MyDisk -FileSystem CSVFS_ReFS -Size 1TB -
ResiliencySettingName Mirror -PhysicalDiskRedundancy 2
```



```

Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName MyDisk -FileSystem CSVFS_ReFS -Size 1TB

DriveLetter FriendlyName FileSystemType DriveType HealthStatus OperationalStatus SizeRemaining Size
-----
MyDisk      CSVFS_ReFS Fixed      Healthy      OK              1015.35 GB ... GB

PS C:\Users\Administrator>
    
```

FIGURE 20. Create disk volume

This example shows how to create a 1 TB parity volume on Storage Spaces Direct:

```

PS C:\> New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName MyDisk -FileSystem CSVFS_ReFS -Size 1TB -
ResiliencySettingName Parity
    
```

3.4 Deploying converged solutions

A converged environment separates compute and storage, placing them in different clusters. This allows the compute and storage functions to use independent hardware resources and provide independent scalability. The S2D volumes are accessed through the network using the SMB3 protocol. To understand more about this deployment option, see the following documentation: docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-direct-overview.

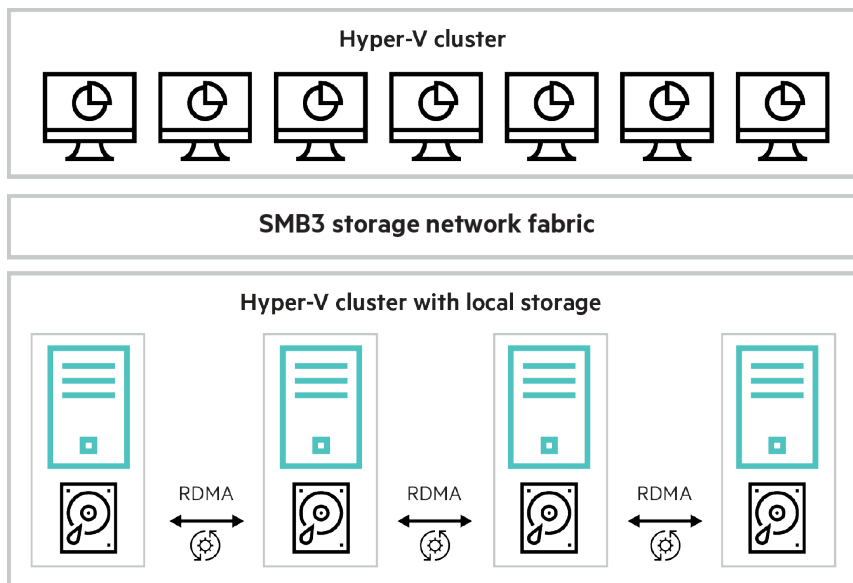


FIGURE 21. Converged solution



4.0 TROUBLESHOOTING

4.1 Cluster verification

Use the Test-Cluster cmdlet to verify your system configuration. A full detailed report will identify the areas where the requirements are not met.

PS C:\> Test-Cluster -Node N1, N2, N3, N4 -Include "Storage Spaces Direct", Inventory, Network, "System Configuration"

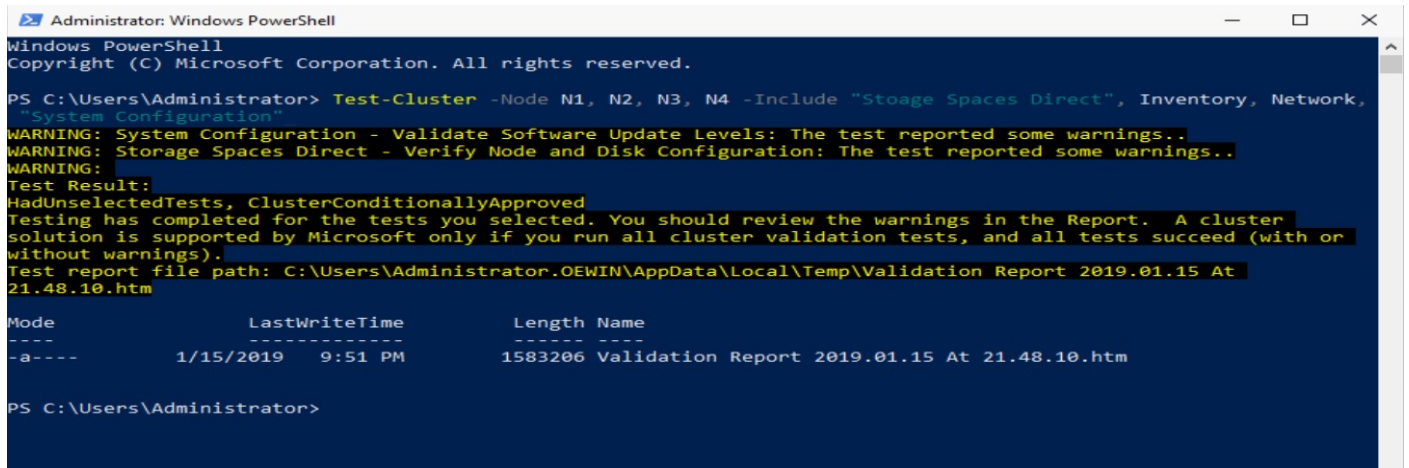


FIGURE 22. A full detailed report after using the Test-Cluster cmdlet

Use a browser to view the report that is generated in the specified folder location.

4.2 Disk management

In the following report, there were no eligible disks found on Node 3 for S2D. Replace the storage with the appropriate media to fix the problem.

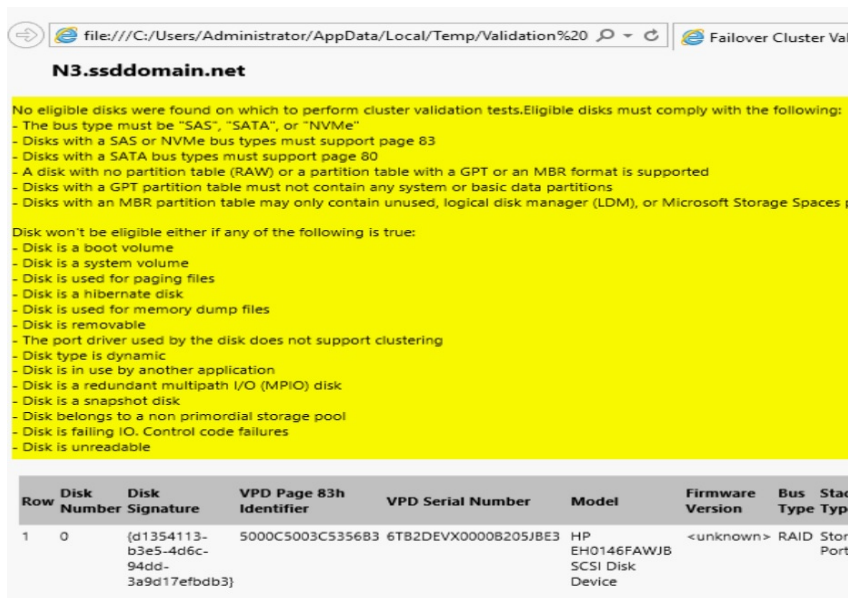


FIGURE 23. No eligible disks found on Node 3 for S2D

Ensure the proper HPE Smart Array controller driver is installed using the SPP. The Windows inbox storage controller driver will make your SSDs/HDDs not eligible for S2D. The following are some useful tips for disk management:

Disks cannot contain any partitions and should be cleaned using the "Clean" command in DiskPart.



Use the following cmdlet to view the list of disks used by the storage pool:

```
PS C:\> Get-StoragePool -FriendlyName "S2D*" | Get-PhysicalDisk
```

Use the following cmdlet to delete a virtual disk:

```
PS C:\> Remove-VirtualDisk -FriendlyName "MyDisk"
```

4.3 Disk rebuild

When replacing failed disks with new ones, verify data reconstruction status using Get-StorageJob as follows.

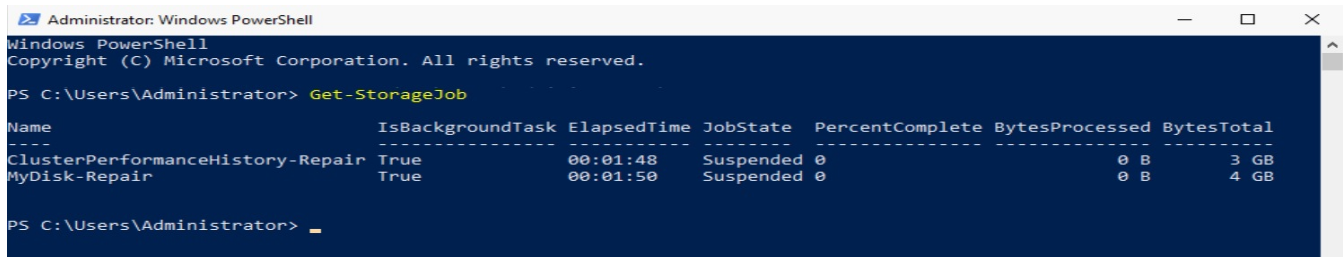


FIGURE 24. Storage jobs

4.4 Disk clean

All disks need to be empty before enabling S2D. You can clean all disks using PowerShell scripts. For more information, see the **Clean drives** section in following documentation: docs.microsoft.com/en-us/windows-server/storage/storage-spaces/hyper-converged-solution-using-storage-spaces-direct#step-3-configure-storage-spaces-direct.

4.5 Cluster management

Adding a new node to your existing cluster can extend your storage pool size. However, only volumes created after adding the new node will be able to make use of the additional storage space available in the new node. Add a new node using the Get-Cluster cmdlet as follows:

```
PS C:\> Get-Cluster MYCLUSTER | Add-ClusterNode NEWNODE
```




FIGURE 25. Cluster management

NOTE

Each cluster may contain up to a maximum of 16 nodes.

4.6 Network management

Use the following cmdlets to disable RDMA:

```
PS C:\> Disable-NetAdapterRdma "Embedded FlexibleLOM 1 Port 1"
```

```
PS C:\> Disable-NetAdapterRdma "Embedded FlexibleLOM 1 Port 2"
```

```
PS C:\> Set-NetOffloadGlobalSetting -NetworkDirect Disabled
```

Alternatively, enable RDMA again using:

```
PS C:\> Enable-NetAdapterRdma "Embedded FlexibleLOM 1 Port 1"
```




```
PS C:\> Enable-NetAdapterRdma "Embedded FlexibleLOM 1 Port 1 2"
```

```
PS C:\> Set-NetOffloadGlobalSetting -NetworkDirect Enabled
```

4.7 Remove Storage Spaces Direct

Before you remove a Storage Spaces Direct Cluster, you should remove all virtual disks created in the Storage Spaces Direct storage pool (see [Disk Management](#)). Once all virtual disks are removed, you can run the following commands to remove the Storage Spaces Direct Cluster.

```
PS C:\> Disable-ClusterStorageSpacesDirect
```

```
PS C:\> Get-Cluster MYCLUSTER | Remove-Cluster
```

Reboot all nodes and clean all used disks.

5.0 ADDITIONAL RESOURCES

[Microsoft Storage Spaces Direct](#)

[Microsoft Storage Replica](#)

[Microsoft Azure Stack HCI home page](#)

[HPE Azure Stack HCI home page](#)

APPENDIX 1: PERFORMANCE VALIDATION METHOD

A1.1 VMFleet for Storage Spaces Direct

In this section, we describe the standardized environment and the procedure used to test our performance validation. Unless explained otherwise, all the tests are based on following criteria:

1. Four-node isolated S2D cluster
2. Three-way mirror volume resilience setup for use by the VM guests
3. VMFleet running 20 VM guests per node. Each VM guest runs one instance of diskspd

VMFleet is part of diskspd and provides performance stress against an S2D deployment. Our validations are carried out in the following manner:

1. Download diskspd from GitHub. The VMFleet files are located in diskspd-master\Frameworks\VMFleet folder
2. Copy the VMFleet folder onto one of the S2D nodes
3. Run following commands:

```
Get-ClusterNode |% { New-Volume -StoragePoolFriendlyName S2D* -FriendlyName $_ -FileSystem CSVFS_ReFS -Size 12TB }  
(Volume size depends on VM size and number of VMs)
```

```
New-Volume -StoragePoolFriendlyName S2D* -FriendlyName Collect -FileSystem CSVFS_ReFS -Size 1TB
```

4. Go to the VMFleet folder and run ".\install-vmfleet.ps1 -source."
5. Copy DISKSPD.exe into C:\ClusterStorage\Collect\Control\Tools
6. Open Hyper-V Manager and install Windows Server 2019 Core VM for VMFleet. Copy VM file (YourVM.vhdx) to C:\ClusterStorage\Collect
For testing purposes, we recommend using a fixed size VHDX (20 GB).
7. To create 20 running VMs, run the following command:
. \create-vmfleet.ps1 -basevhd "C:\ClusterStorage\collect\YourVM.vhdx" -adminpass xxx -connectpass xxx -connectuser "Administrator" -vms 20
-adminpass and -connectpass are the same password as the VM administrator.
Allow some time for 20 VM deployment.



8. The following command is optional for VM configuration:

```
.\set-vmfleet.ps1 -ProcessorCount 1 -MemoryStartupBytes 8GB -DynamicMemory $false
```

9. All VMFleet files are ready; run `.\start-vmfleet.ps1` to start all VMs.

```
Run .\test-clusterhealth.ps1 to check for any problems in the cluster
```

10. To monitor the storage performance, run `.\Watch-Cluster.ps1` (Ctrl-C to stop)

11. Run test cases on VMs; the following commands run the test for 600 seconds using seven threads to drive 0%/10%/30% random 64 KiB writes at a depth of 40 overlapped I/Os to a regular file.

```
.\start-sweep.ps1 -b 64 -t 7 -o 40 -w 0 -p r -d 600
```

```
.\start-sweep.ps1 -b 64 -t 7 -o 40 -w 10 -p r -d 600
```

```
.\start-sweep.ps1 -b 64 -t 7 -o 40 -w 30 -p r -d 600
```

12. Log files will be generated when all tests are completed in the VMs.

13. To collect the report files, copy *.xml files from `C:\ClusterStorage\Collect\Control\report` to a new folder.

- Copy `process-diskspd.ps1` from `diskspd-master` folder to the folder you created
- Run `process-diskspd.ps1`, report files will be collected into single file named `result.tsv`

14. To stop all VMs, run `.\stop-vmfleet.ps1`

15. To destroy all VMs, run `.\destroy-vmfleet.ps1`

A1.2 How to monitor the performance of VMFleet

Performance Monitor provides counters and recorders for monitoring processor utilization, disk read/write bytes, and network activity.

The following steps outline how to monitor specific counters.

1. Open Performance Monitor, browse to **Data Collector Sets** and create a **new Data Collector Set**.




FIGURE 26. Create a new Data Collector Set in Performance Monitor



2. Name the collector set and choose **Create manually (Advanced)**. Click **Next**.

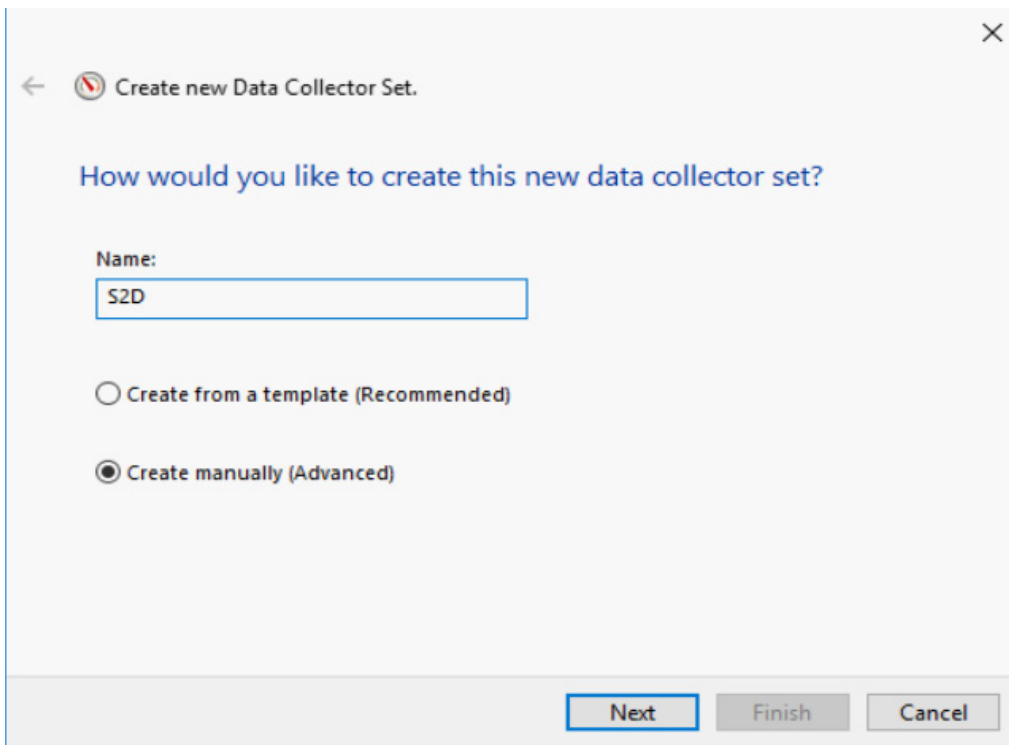


FIGURE 27. Name the new Data Collector Set

3. Select **Performance counter** and click **Next**.

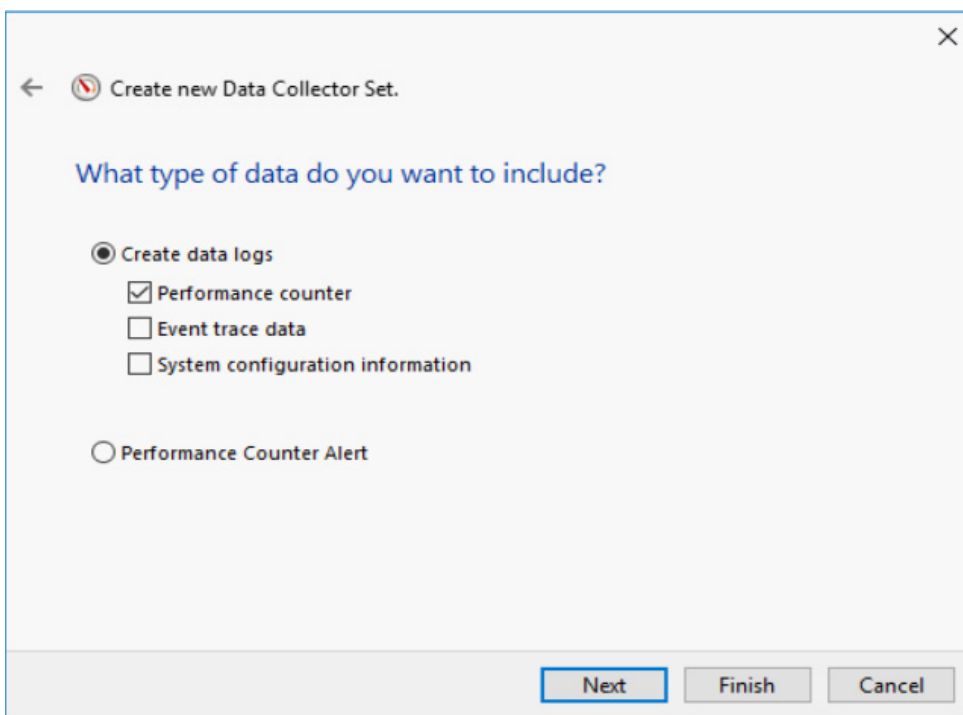


FIGURE 28. Choose the data type

4. Click **add** and select the desired counters. We recommend the following counters:

Processor



Cluster CSVFS

Sample interval can be 5 seconds or longer. Click **Next** and select the location of report.

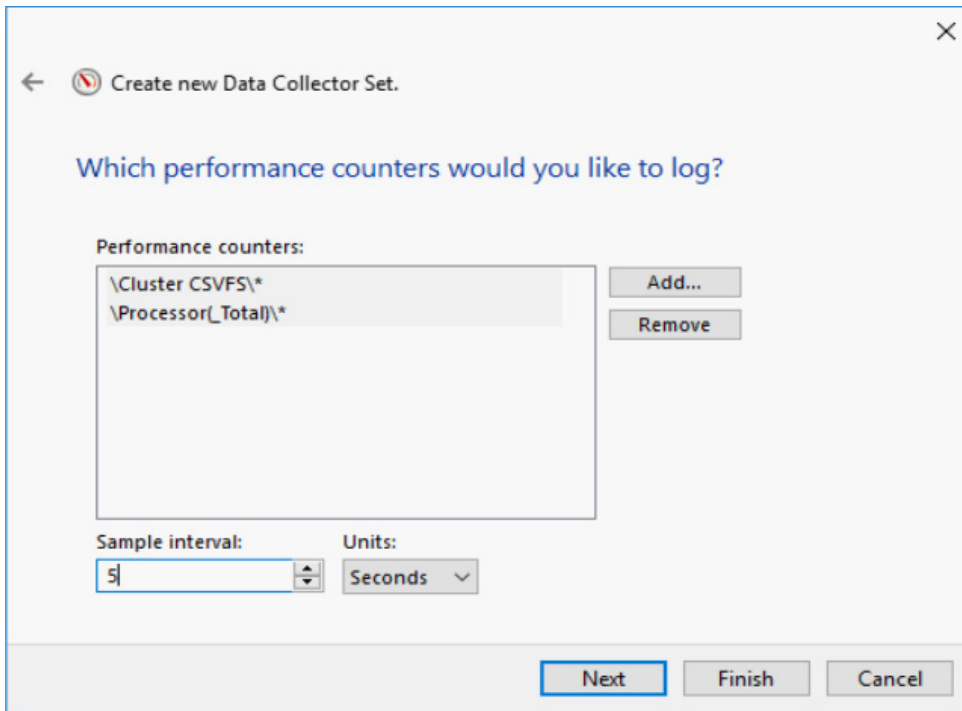


FIGURE 29. Add performance counter for monitoring

5. The new collector set will be added. Select the data collector and edit its properties. You can set **Maximum samples**.

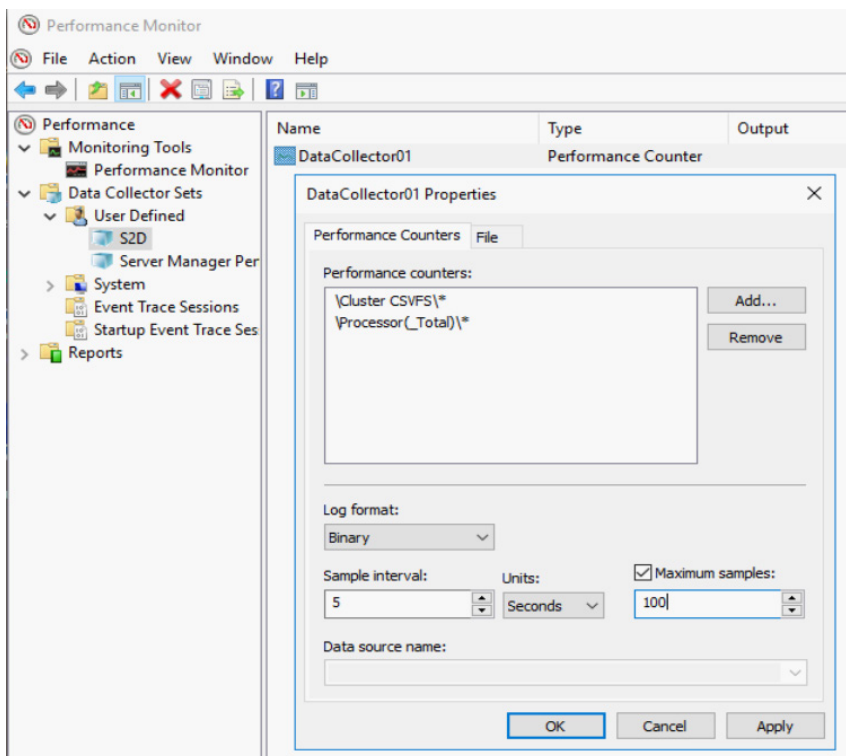


FIGURE 30. Set maximum samples in DataCollector properties



- When VMFleet starts running, allow five seconds for it to **warm up**. Then, click **Start** on your Data Collector Sets. If you set **Maximum Samples** already, the collector will stop automatically when the maximum is reached.

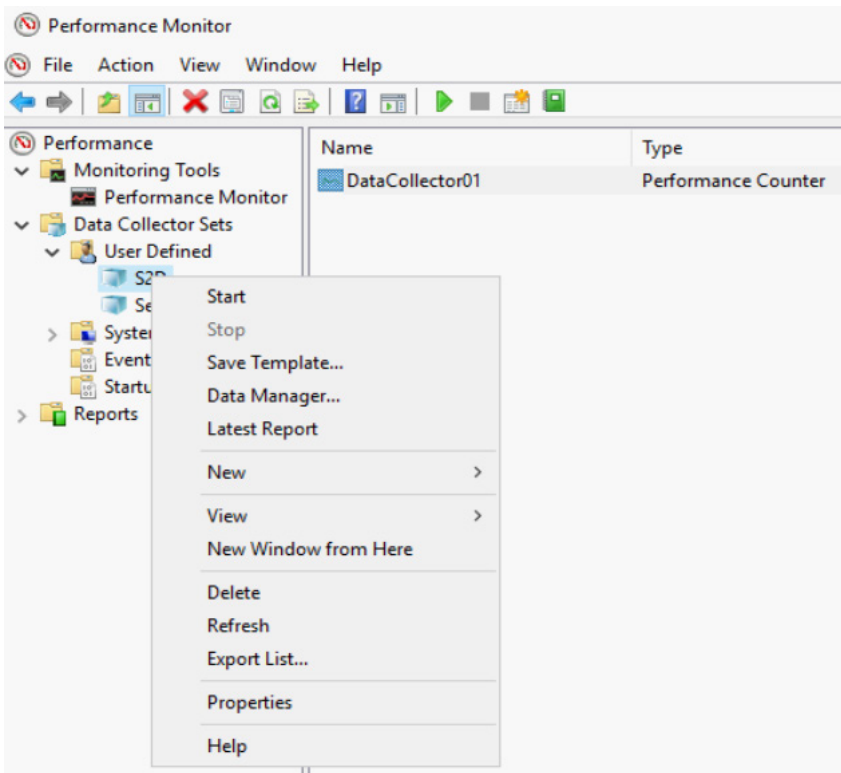


FIGURE 31. Content window of the Data Collector Sets

APPENDIX 2: WINDOWS ADMIN CENTER (WAC) EXTENSION FOR AZURE STACK HCI

A2.1 WAC Extension preview

The HPE WAC extension for Azure Stack HCI preview is ready for download from the Microsoft Windows Admin Center Feed.

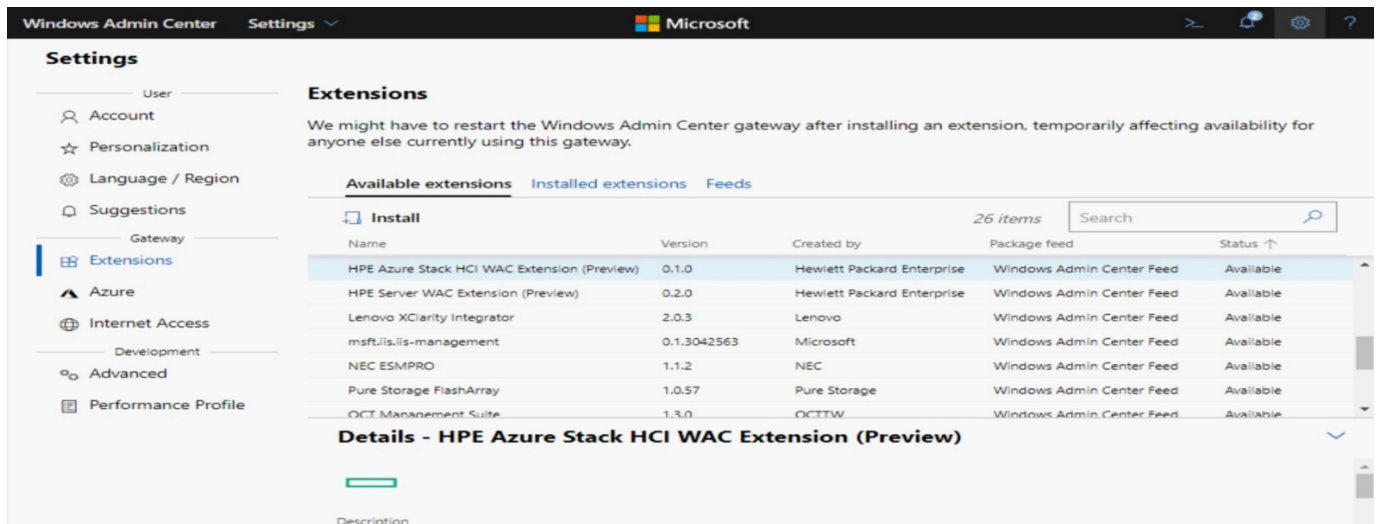


FIGURE 32. HPE WAC Extension

From the Settings icon the Top Right, select Gateway “Extensions”. The HPE Azure Stack HCI Extension Preview version 0.1.0 is listed. Click on “Install” to download.



LEARN MORE AT

hpe.com/solutions/azurestackhci

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

© Copyright 2019–2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel Optane DC and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Active Directory, Azure, Hyper-V, Microsoft, PowerShell, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.

a00067845ENW, March 2020, Rev. 3