

Enabling Mission Critical Workloads on APEX Cloud Platform for Azure with Dell PowerFlex White Paper

Abstract

This white paper outlines the Dell APEX Cloud Platform for Microsoft Azure's integration with Dell PowerFlex software-defined storage. This integration, a collaborative effort between Dell Technologies and Microsoft, supports linear scaling of storage resources, independent of Azure Stack HCI compute and Storage Spaces Direct resources. It is designed to handle mission-critical workloads with diverse scalability and capacity needs.

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.

Revision history.....	4
Chapter 1: Introduction.....	5
Business challenge.....	5
Document purpose.....	6
Terminology.....	6
Audience.....	6
Chapter 2: Architecture concepts and requirements.....	7
Overview.....	7
Storage Spaces Direct as primary storage.....	8
PowerFlex integration.....	8
Data mobility.....	8
Data encryption.....	9
Chapter 3: Proof of concept configuration.....	10
Overview.....	10
Azure Stack HCI cluster node networking.....	11
PowerFlex custom node networking.....	13
SDC driver.....	13
PowerFlex CSV special consideration.....	14
Node configurations.....	14
Physical networking.....	15
Management software.....	15
Chapter 4: Presenting PowerFlex volumes to the Azure Stack HCI cluster.....	17
Overview.....	17
Prerequisites.....	17
Creating volumes in PowerFlex Manager.....	17
Creating the CSV on the Azure Stack HCI cluster.....	19
Validating volumes.....	20
Deploying virtual machines onto PowerFlex CSVs.....	20
Using PowerFlex CSVs for AKS persistent volumes.....	22
Chapter 5: Summary and conclusion.....	23
Summary and conclusion.....	23
We value your feedback.....	23
Chapter 6: References.....	24
References.....	24

Revision history

Table 1. Revision history

Date	Document revision	Description of changes
April 2024	1.0	Initial release

Introduction

Topics:

- [Business challenge](#)
- [Document purpose](#)
- [Terminology](#)
- [Audience](#)

Business challenge

Dell APEX Cloud Platform for Microsoft Azure is an on-premises, turnkey infrastructure platform that Dell Technologies and Microsoft collaboratively engineered to optimize the Azure hybrid cloud experience. It is the first offer in the market for Premier Solutions for Microsoft Azure Stack HCI. This is a new category in the Azure Stack HCI catalog that is reserved for Microsoft's top-tier partners that deliver the highest levels of integration and continuous validation. Dell APEX Cloud Platform Foundation Software uses a Dell extension in Microsoft Windows Admin Center to provide automated management and orchestration for the platform.

Azure Stack HCI is a hyperconverged infrastructure that consists of proven enterprise-class, software-defined technologies that have been running at-scale in Microsoft Azure and in enterprise data centers for years. These familiar technologies include Hyper-V, Storage Spaces Direct (S2D), and Software Defined Networking (SDN). Since the inception of the APEX Cloud Platform for Azure in September 2023, S2D has functioned as the primary storage solution. It combines the internal drive performance and capacity across cluster nodes into a single resilient and highly performant software-defined pool of virtual shared storage. You can scale out the storage capacity and performance of your cluster by adding more drives or adding more nodes to the cluster.

For many workloads, S2D adequately meets the current and projected performance and resiliency requirements. Many customers appreciate its tight integration into the Azure Stack HCI ecosystem and the ability to fully standardize on Microsoft technologies. In addition, APEX Cloud Platform Foundation Software fully automates the deployment and lifecycle management of Azure Stack HCI with S2D as the primary storage. However, certain workloads might benefit from alternative storage solutions.

APEX Cloud Platform for Azure integrated with Dell PowerFlex software-defined storage currently provides block storage. Using Dell's enterprise-class storage solution to extend the Azure Stack HCI storage fabric enables new use cases and provides many benefits:

- Run a highly diverse set of workloads – including mission-critical databases and applications that require significantly low latency and high transactional performance at large scale beyond S2D's default capability.
- Deliver consistent I/O performance for streaming, ingestion, and reporting transactions within AI and data analytics applications.
- Achieve 99.9999% availability with extremely fast rebuild and rebalance operations.
- Address exponential data growth with a modular scale-out architecture.

APEX Cloud Platforms with PowerFlex is a key building block in the creation of a universal storage layer. A universal storage layer is created when common software-defined storage services are deployed across all locations where your data lives and applications run. With Dell software, a universal storage layer is created with Dell enterprise storage on-premises, Dell storage for public cloud, and APEX Cloud Platforms with PowerFlex. Intelligent management with automation can streamline storage management and orchestrate data movement between locations where the common software-defined storage services are deployed.

To help onboard your data to this universal storage layer, Dell has developed a data mobility utility that seamlessly moves data from APEX Cloud Platforms for Azure running S2D for primary storage to platforms with PowerFlex. This utility enables you to quickly take advantage of the increased performance, scalability, and data management features of PowerFlex to unleash innovation and unlock new use cases, propelling your organization forward in this highly competitive market.

Document purpose

The Dell engineering team set up an APEX Cloud Platform for Azure with PowerFlex in the Dell labs as a proof of concept to ensure it supported each use case. This document focuses on the software driver installation and network configuration steps taken to successfully extend the Azure Stack HCI storage fabric to PowerFlex software-defined storage in the lab.

NOTE: You cannot order the configuration presented in this guide nor does Dell Technologies support any attempt to configure the presented solution. APEX Cloud Platform for Azure with PowerFlex will be available when our fully supported Dell Validated Design is published.

Terminology

The following table provides definitions for some of the terms and concepts that are used in this white paper.

Table 2. Terminology

Term	Definition
MC node	Multicloud (MC) nodes are deliberately configured common servers for APEX Cloud Platforms based on PowerEdge Servers. For APEX Cloud Platform for Azure, MC nodes run the Azure Stack HCI operating system with Hyper-V, S2D, and virtual networking. This guide uses MC node and Azure Stack HCI cluster node interchangeably.
CSVs	Cluster Shared Volumes (CSV) enable multiple nodes in an Azure Stack HCI cluster to simultaneously have read-write access to the same disk.
Dell PowerFlex	Dell PowerFlex is Dell Technologies' software-defined storage solution. With its exceptional performance, scale, and ability to consolidate diverse workloads, PowerFlex instills confidence in your infrastructure to securely deliver under extreme requirements and enables businesses to remain competitive and respond quickly to changing conditions.
PowerFlex custom node	PowerFlex custom nodes offer you the option to deploy PowerFlex software-defined storage on a broad range of Dell PowerEdge servers. The PowerFlex R760 custom node was leveraged in the testing performed for this white paper.
PowerFlex Manager	PowerFlex Manager is the unified management and orchestration platform for Dell PowerFlex with automation to enable high availability, serviceability, scalability, efficiency, and agility. In the test lab, this software ran co-resident on the PowerFlex custom nodes.
SDC	Storage Data Client (SDC) is a client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It acts like a virtual HBA and presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections with all storage nodes and enables parallel I/O with native multipathing.

Audience

This white paper is intended for IT experts who are interested in building an optimized Azure hybrid cloud environment. The document provides important context for organizations considering a migration to Dell APEX Cloud Platform for Microsoft Azure from competing solutions. This information is also helpful for users who are interested in unlocking new use cases by extending Azure Stack HCI to include Dell enterprise-class software-defined storage.

The audience for this guide can include IT managers, system engineers, field consultants, and others with expertise in virtualization and cloud computing.

Architecture concepts and requirements

Topics:

- Overview
- Storage Spaces Direct as primary storage
- PowerFlex integration
- Data mobility
- Data encryption

Overview

This chapter begins with a high-level overview of the solution architecture for the testing in the Dell Technologies engineering labs. First, it is important to understand the role of S2D and PowerFlex software-defined storage in the test environment. The following diagram depicts the integration between APEX Cloud Platform for Azure running S2D as primary storage and an external PowerFlex cluster used to extend the Azure Stack HCI storage fabric.

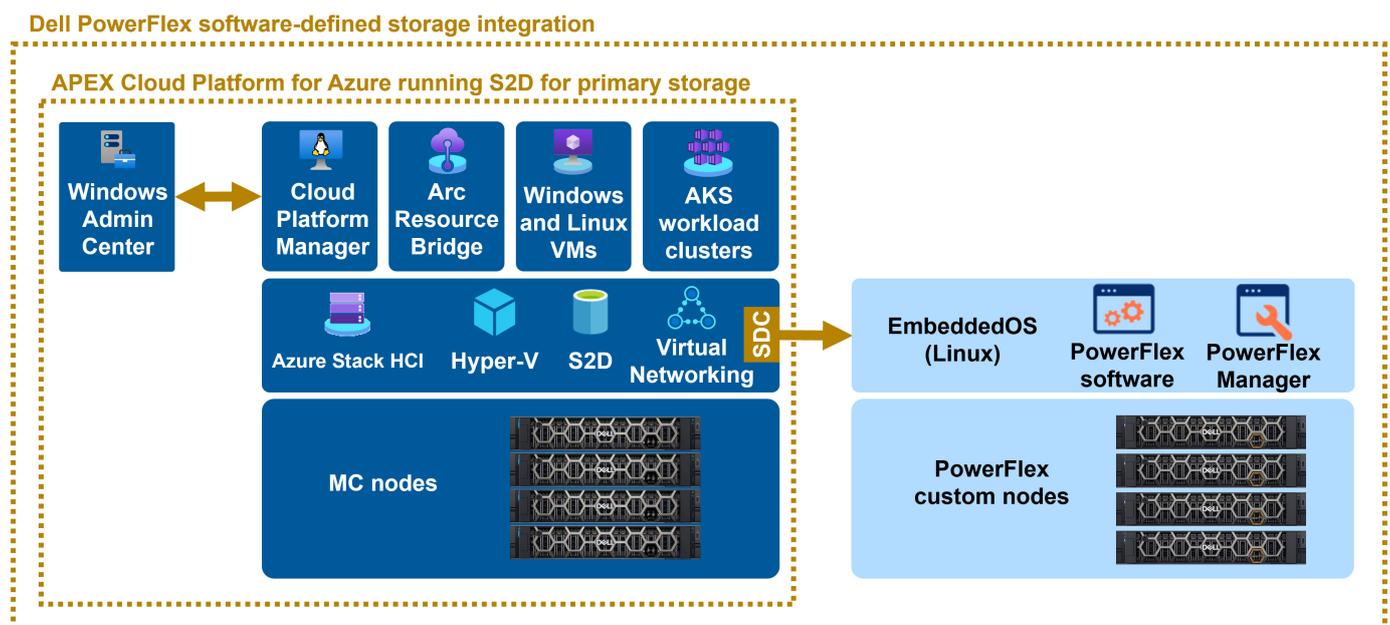


Figure 1. APEX Cloud Platform for Azure storage configurations

During the proof of concept testing, the team prioritized the following configuration quality attributes:

- High availability: Avoid single points of failure at the compute, storage, and networking layers.
- Scalability: Enable you to right-size the infrastructure to meet the demands of the current workload profile, while also letting you scale up and out for future growth projections.
- Manageability: Leverage familiar tools to manage the configuration.
- Simplicity: Provide a seamless experience for presenting this new external storage to Azure Stack HCI for IT administrators already familiar with S2D.

Storage Spaces Direct as primary storage

APEX Cloud Platform for Azure is deployed as a single-tier hyperconverged infrastructure running the Azure Stack HCI operating system and Microsoft software-defined technologies – Hyper-V, S2D, and virtual networking. During initial deployment, the Cloud Platform Manager VM and Azure Arc Resource Bridge VM are brought online automatically. The Cloud Platform Manager VM runs the microservices-based APEX Cloud Platform Foundation Software, which communicates with the MC nodes and integrates automation workflows into Windows Admin Center. The Azure Arc Resource Bridge VM is running a Kubernetes management cluster and is designed to host other Azure Arc services. These VMs can only be provisioned to Cluster Shared Volumes (CSVs) created from the S2D storage pool.

The Azure Stack HCI cluster is registered with your Azure subscription and onboarded to Azure Arc. Arc Resource Bridge enables self-service provisioning of VMs and Azure Kubernetes Service (AKS) workload clusters onto Azure Stack HCI clusters from Azure.

For more details about the S2D configuration, see [References](#).

PowerFlex integration

For this integration, the APEX Cloud Platform for Azure running S2D as primary storage is deployed first. Then, a separate, disaggregated PowerFlex storage cluster is deployed and configured. Any of the PowerFlex product flavors (rack, appliance, or custom node) can be used to satisfy this solution. The Azure Stack HCI cluster nodes communicate with the PowerFlex cluster by installing a software driver and by setting up the host-based virtual networking and the network infrastructure in a specific way. After the integration is complete, you can provision VMs on CSVs created from the S2D storage pool or PowerFlex block storage. You can also create and access persistent volumes mapped to PowerFlex CSVs from your AKS worker nodes.

S2D is a core technology of Azure Stack HCI and therefore cannot be disabled after the integration of PowerFlex is complete. Even if you intend to exclusively consume PowerFlex block storage capacity, you must still allocate the minimum S2D capacity that is required by Azure Stack HCI plus enough capacity to host critical VMs like the Cloud Platform Manager and Azure Arc Resource Bridge. These VMs are created at initial Azure Stack HCI cluster creation and must remain on S2D CSVs for the life of the cluster.

Enabling the Azure Stack HCI cluster to consume block storage from the PowerFlex cluster is a simple process. Dell and Microsoft have collaboratively engineered the ability to extend the Azure Stack HCI storage fabric to external storage. The team performed all the proof of concept testing in the Dell engineering labs to ensure optimal performance and resiliency. The following list describes the primary tasks that are required for the integration:

1. Create two discrete, non-routable storage VLANs in the network infrastructure between the Azure Stack HCI and PowerFlex clusters.
2. Create a new Network ATC intent on the Azure Stack HCI cluster with associated storage vNICs.
3. Install and configure the Storage Data Client (SDC) driver on each Azure Stack HCI cluster node. This SDC driver facilitates communication with the PowerFlex cluster over the IP network.
4. Create and map new volumes on the PowerFlex cluster.
5. Create new Azure Stack HCI CSV(s) on PowerFlex volumes using familiar tools. CSVs created on S2D are formatted with the Resilient File System (ReFS), and CSVs created on PowerFlex are formatted with the New Technology File System (NTFS).

Data mobility

To facilitate the replication and migration of data from an APEX Cloud Platform for Azure running S2D as primary storage to a platform with PowerFlex, Dell provides the Dell APEX Cloud Platform for Microsoft Azure data mobility utility. Dell has created this utility to address scenarios where the requirements for a virtualized application that is running in a VM on an S2D CSV have evolved and could benefit from the performance, scalability, and resiliency delivered by a PowerFlex CSV.

Another reason to use the data mobility utility is to create a copy of data for use by other applications, such as those performing artificial intelligence training. This utility delivers the following functionality:

- Select one or many VMs to replicate the Virtual Hard Disks (VHDs) for or migrate from the source to the destination platform.
- If a VM has multiple VHDs, you can choose which VHDs are in scope.
- You can request a differential copy of the VHD files after the initial full copy to keep the data synchronized.
- Leverages native Microsoft checkpoints and can benefit from SMB3 to transfer data across.
- Provides regular point-in-time checkpoints and the ability to restore to individual point-in-time checkpoints.

- Data does not require the VM metadata to be moved.
- Data can be restored without requiring the full VM to be rebuilt.

Data encryption

Data encryption is supported in both S2D CSVs and PowerFlex CSVs, and the approach for each configuration differs. In the APEX Cloud Platform for Azure, each S2D CSV is encrypted by default using BitLocker. BitLocker is a Windows security feature that provides encryption for entire volumes, addressing the threats of data theft or exposure from lost, stolen, or inappropriately decommissioned devices.

BitLocker is not supported on PowerFlex CSVs, and it must not be enabled on these volumes. If you have to encrypt PowerFlex volumes, Dell offers support for CloudLink. CloudLink provides software-based Data at Rest Encryption (D@RE) that is transparent to the features and operation of the PowerFlex software. It uses dm-crypt, a proven high-performance volume encryption solution that is widely implemented for Linux machines. CloudLink must be deployed outside of the APEX Cloud Platform for Azure on customer-provided infrastructure.

For more information about CloudLink, see [References](#).

Proof of concept configuration

Topics:

- Overview
- Azure Stack HCI cluster node networking
- PowerFlex custom node networking
- SDC driver
- PowerFlex CSV special consideration
- Node configurations
- Physical networking
- Management software

Overview

This chapter describes the configuration built in the Dell engineering labs. Only the unique aspects of the PowerFlex integration will be explored in-depth. The Dell engineering team used eight MC-760 Azure Stack HCI nodes and eight PowerFlex R760 custom nodes for these tests, as shown in Figure 2. During these tests, the team considered both the ToRs and aggregation layers of the network infrastructure.

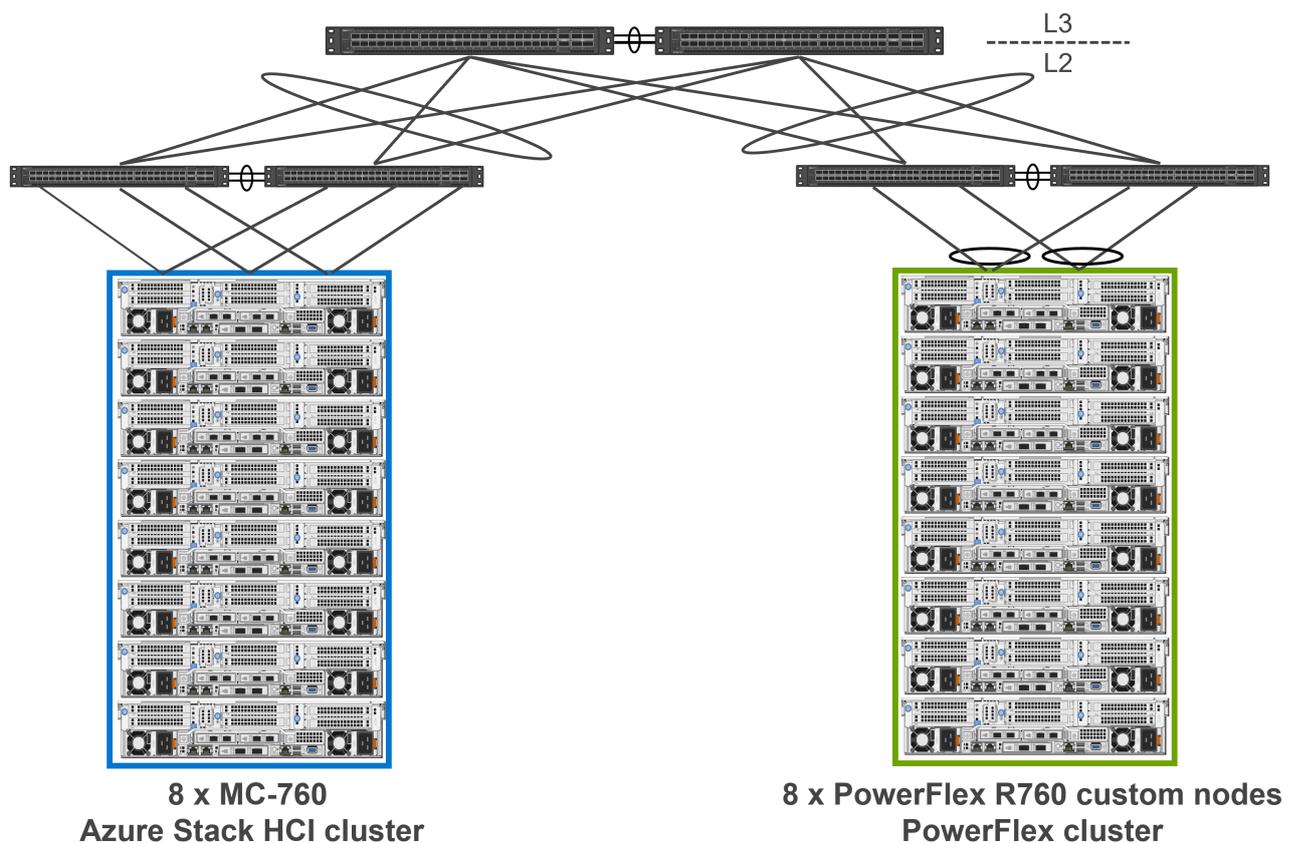


Figure 2. Lab network infrastructure

The following diagram depicts the host networking configuration and connectivity to the ToRs for both the MC nodes and PowerFlex custom nodes.

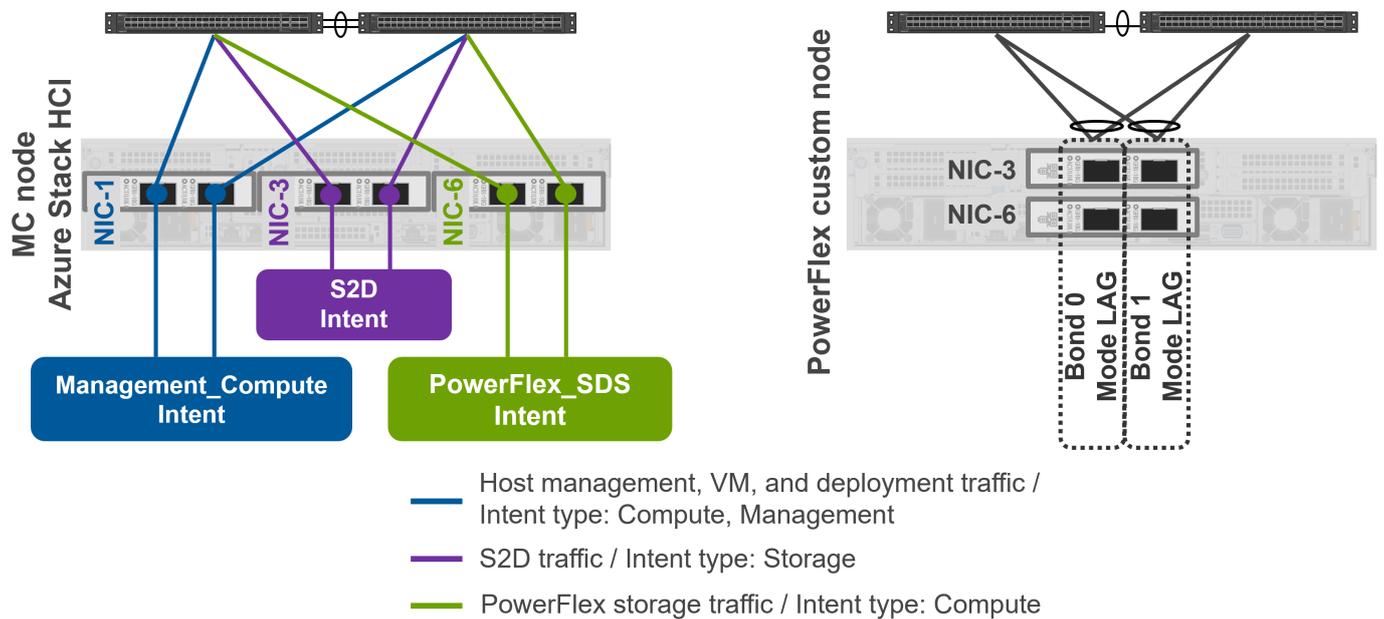


Figure 3. Host networking configuration

Azure Stack HCI cluster node networking

Dell engineering created the following Network ATC intents on each Azure Stack HCI cluster node, as shown in Figure 3.

Table 3. Network ATC intents

Intent Name	Intent Type	Network Adapter	Adapter Label in Figure 3
Management_Compute	Management and compute	NVIDIA ConnectX-6 Lx Dual-Port 10/25GbE SFP28 OCP NIC 3.0	NIC-1
S2D	Storage	NVIDIA ConnectX-6 Lx Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile Add-In	NIC-3
PowerFlex_SDS	Compute	NVIDIA ConnectX-6 Lx Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile Add-In	NIC-6

The network topology setup was a non-converged network configuration. In this configuration, management/compute traffic was separated from S2D and PowerFlex storage traffic. The PowerFlex_SDS intent was configured to use the compute intent type. The following list includes additional details about the PowerFlex_SDS intent:

- The compute intent type was chosen because the Azure Stack HCI operating system only allows one intent to be configured with the storage intent type. The storage intent type is reserved for use by the S2D intent.
- By default, all intents with the compute intent type had Single Root I/O Virtualization (SR-IOV) enabled. However, this feature was not used for the integration.
- Quality of Service (QoS) policies were not configured on the PowerFlex_SDS intent because there was no potential for contention with other traffic types.
- RDMA was enabled by default but only used on the S2D intent. It was not required on the PowerFlex_SDS intent.
- Jumbo Frames were enabled and set to 9014.

- The network adapter port connectivity speed on all physical ports was set to 25 GbE.

A summary of this information is available in the following table.

Table 4. Powerflex_SDS Intent settings

Feature	Default State	Configured State
RDMA	Enabled but not used	Enabled but not used
SR-IOV	Enabled but not used	Enabled but not used
QoS	Default (no configuration)	Default (no configuration)
Jumbo Frames	Disabled	Enabled (MTU:9014)

The PowerFlex_SDS intent was mapped to the two ports on physical NIC-6 on each MC node, as shown in Figure 4. Each port was configured on a separate, dedicated, non-routable VLAN. These VLANs were configured in the network infrastructure to facilitate communication between the MC nodes and PowerFlex custom nodes.

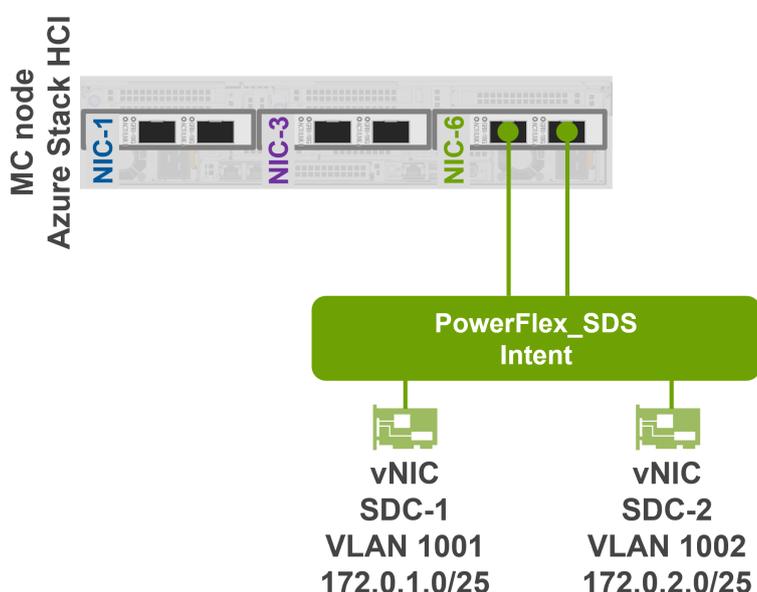


Figure 4. PowerFlex SDS intent vNIC configuration

The following table shows how the vNICs were configured with affinity to specific ports on NIC-6.

Table 5. PowerFlex_SDS vNIC affinity

Intent	vNIC	Physical NIC	Port	VLAN
PowerFlex_SDS	SDC-1	NIC-6	1	1001
PowerFlex_SDS	SDC-2	NIC-6	2	1002

In the event of a port failure, the virtual NIC bound to that port could fail over and use the other physical port. These vNICs were labeled SDC-1 and SDC-2, as the SDC that was installed on the Azure Stack HCI operating system facilitated the communication down these vNICs to the PowerFlex cluster. On each node, SDC-1 was configured with the first SDC VLAN ID, and SDC-2 was configured with the second SDC VLAN ID. Both nodes had Jumbo Frames enabled and the MTU was set to 9014.

PowerFlex custom node networking

NIC high availability was set up in the SLES operating system on the PowerFlex custom nodes by creating a bond across the respective ports on each network adapter. A routable VLAN interface was created on Bond 0 for PowerFlex management, and a non-routable interface was also created on Bond 0 for the PowerFlex storage traffic from the Azure Stack HCI cluster nodes on the SDC-1 VLAN. Bond 1 was configured to carry storage traffic exclusively over the SDC-2 non-routable VLAN. The following diagram illustrates this configuration.

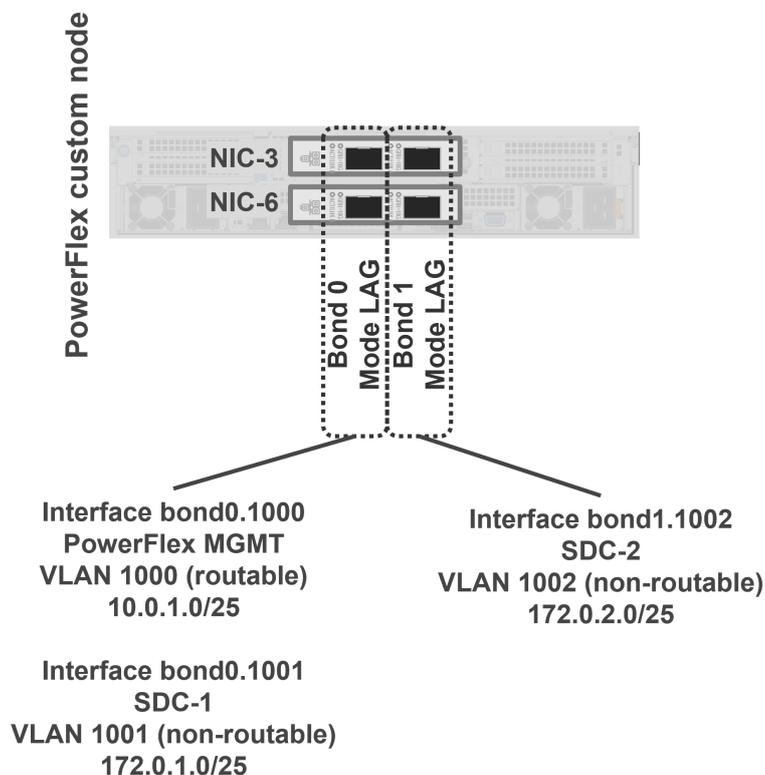


Figure 5. PowerFlex custom node host networking configuration

The following table shows how the NICs for each MC Storage node were cabled to the corresponding ToR Switches to ensure high availability.

Table 6. PowerFlex custom node NIC cabling

NIC/Port	ToR Switch
Bond 0 - NIC 3 Port 1	ToR A
Bond 0 - NIC-6 Port 1	ToR B
Bond 1 - NIC 3 Port 2	ToR B
Bond 1 - NIC-6 Port 2	ToR A

SDC driver

The Dell engineering team installed the SDC driver on each of the Azure Stack HCI cluster nodes. This driver enabled the communication to the PowerFlex cluster over the dedicated, non-routable storage VLANs. The driver also was responsible for presenting the block storage volumes that were created in PowerFlex Manager to the Azure Stack HCI operating system on each cluster node as raw disks. This allowed the creation of NTFS-formatted Clustered Shared Volumes that could be consumed by all nodes on the Azure Stack HCI Cluster.

For more information about the role of the SDC driver, see [References](#).

PowerFlex CSV special consideration

After the networking and SDC driver installation and configuration were complete, Dell engineering created the PowerFlex volumes and mapped them to the Azure Stack HCI cluster node SDCs. Once mapping was completed, Dell engineering created the CSVs. They fully populated the CSVs with data before using a synthetic load testing tool for performance and resiliency testing.

Instructions for creating the PowerFlex CSVs are available in the chapter [Presenting PowerFlex Volumes to the Azure Stack HCI Cluster](#) of this document.

Node configurations

The following table contains the software and hardware components and versions that were used for the MC nodes in the Dell labs for this proof of concept testing.

Table 7. MC node specifications

Resources per cluster node	Description	Firmware and driver versions
CPU	Dual-socket Intel Xeon Gold 6430 2.1G, 32C/64T	N/A
Memory	256 GB DDR5, 4.8 MT/s RDIMMs	N/A
Storage controller for the operating system	BOSS-N1 controller card + 2 x M.2 960GB (RAID 1)	2.1.13.2025
Physical drives for operating system	2	N/A
Storage controller for the Storage Spaces Direct storage pool	Dell HBA355i	FW: 24.15.14.00 Driver: 23.7.26138.0
Physical drives for Storage Spaces Direct	24 x 800GB SSD SAS MU	N/A
Network adapters	3 x NVIDIA ConnectX-6 Lx Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile	FW: 26.38.10.02 Driver: 2.61.48.00
Operating system	Azure Stack HCI baseline version 2311 (operating system version 23H2)	N/A
BIOS	N/A	1.6.6
iDRAC	N/A	7.00.60.201

The following table contains the software and hardware components and versions that were used for the PowerFlexR760 custom nodes in the Dell labs for the proof of concept testing.

Table 8. PowerFlex custom node specifications

Resources per cluster node	Description	Firmware and driver versions
CPU	Dual-socket Intel Xeon Gold 6430 2.1G, 32C/64T	N/A
Memory	512 GB DDR5, 4.8 MT/s RDIMMs	N/A
Storage controller for the operating system	BOSS-N1 controller card +2 x M.2 960GB (RAID 1)	2.1.13.2025
Physical drives for the operating system	2	N/A
Physical drives for Dell SDS	24 x1.6TB EnterpriseNVMemixed Use	FW: 1.1.1

Table 8. PowerFlex custom node specifications (continued)

Resources per cluster node	Description	Firmware and driver versions
Network adapters	2 x NVIDIA ConnectX-6 Lx Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile	FW: 26.38.10.02 Driver: (kernel based) 5.14.21-150400.24.100-default
Operating system	SUSE Enterprise Linux Server (SLES) 15 SP4	N/A
Kernel	5.14.21-150400.24.100-default	N/A
PowerFlex SDS software	PowerFlex 4.5.1	N/A
BIOS	N/A	1.6.6
iDRAC	N/A	7.00.60.00

Physical networking

For this testing, Dell engineers used the Dell PowerSwitch S5248F-ON running SmartFabric operating system 10 for the ToRs and Dell PowerSwitch S5232F-ON for the aggregation switches.

Table 9. Network switch specifications

Component	Component	Driver/software version
PowerFlex SDS cluster ToRs	2x Dell PowerSwitch S5248F-ON	10.5.6.0.47
Azure Stack HCI cluster ToRs	2x Dell PowerSwitch S5248F-ON	10.5.6.0.47
Aggregation switches	2x Dell PowerSwitch S5232F-ON	10.5.6.0.47

ToR switches

One pair of ToR switches was configured to provide network connectivity to the MC nodes in one rack. A separate pair of ToR switches provided connectivity to the PowerFlex custom nodes in a separate rack. The switches in each pair of ToRs were connected by a Virtual Link Trunking interconnect (VLTi), offering high availability across the switches. The VLTi required sufficient bandwidth to cater to storage, management, and workload traffic in case any of the server facing ports fail. The engineering team used two QSFP dual-density (2x200G) ports for the VLTi (these are shown as 4x100G ports logically in the switch configuration). These switchports participated in the VLTi Switch configurations for this solution.

Aggregation switches

A pair of PowerSwitch S5232F-ON switches were used for the aggregation switches during testing. To minimize oversubscription, each pair of ToR switches connected to the aggregation layer with an 800 Gbps VLT port channel. Furthermore, a VLTi was configured using two 100 Gbps physical interfaces between the aggregation switches to ensure N+1 availability and adequate resiliency through the networking design.

The storage traffic between the MC nodes and the PowerFlex custom nodes were non-routed, layer 2 networks which traversed the ToR and aggregation switches. All ports that were configured for this storage traffic required support for Jumbo Frames for optimal storage performance.

Management software

The following table describes the management tools that were used in the lab proof of concept testing.

Table 10. Management software versions

Management software	Version
Dell APEX Cloud Platform Foundation Software	1.01
Windows Admin Center	Version 2311 Build 1.5.2312.09001
APEX Cloud Platform extension in Windows Admin Center	1.1.0
PowerShell	7.4.0
PowerFlex Manager	4.5.1.0

PowerFlex Manager is the Web-based UI that was installed and configured in the lab as part of the PowerFlex cluster deployment. It runs co-resident on the PowerFlex cluster and is accessed using the PowerFlex management IP. PowerFlex Manager was used to create volumes and map them to the SDC drivers running on the Azure Stack HCI nodes. The team also used the PowerFlex Manager to monitor the health of the PowerFlex cluster volumes.

Windows Admin Center and PowerShell were used to create the PowerFlex Clustered Shared Volumes to host Windows and Linux VMs and AKS workload cluster persistent volumes. The instructions for presenting PowerFlex block storage to Azure Stack HCI and creation of the PowerFlex CSVs are available in the chapter [Presenting PowerFlex volumes to the Azure Stack HCI cluster](#).

Presenting PowerFlex volumes to the Azure Stack HCI cluster

Topics:

- [Overview](#)
- [Prerequisites](#)
- [Creating volumes in PowerFlex Manager](#)
- [Creating the CSV on the Azure Stack HCI cluster](#)
- [Validating volumes](#)
- [Deploying virtual machines onto PowerFlex CSVs](#)
- [Using PowerFlex CSVs for AKS persistent volumes](#)

Overview

This chapter describes the high-level process that Dell engineering followed to create PowerFlex volumes, map the volumes to Azure Stack HCI clusters, and create Cluster Shared Volumes for use by VMs and AKS workload clusters.

Prerequisites

Before creating any volumes, the Dell engineering team validated the following items in the test environment:

- Network connectivity was properly established between the MC nodes and the PowerFlex custom nodes on the two non-routable data networks. This IP communication optimally incurred a latency of under 1 millisecond. Also, Jumbo Frames were correctly set to 9014 on the host interfaces, and the switch ports were set to 9216 throughout the network infrastructure.
- The SDC driver was correctly installed and configured on all MC nodes.
- The PowerFlex integration was proven to be highly available from a compute, storage, and network perspective. This was an N+1 design, so no single hardware component failure would result in a failure of the platform. Dell engineering performed all failure testing while generating workload traffic with a synthetic load generation tool to confirm availability. This included but was not limited to:
 - Single Azure Stack HCI cluster node failure
 - Single PowerFlex custom node failure
 - Single ToR switch failure
 - Single aggregation switch failure
 - Single physical NIC failure
 - Individual switchport and cabling failures throughout the design

Creating volumes in PowerFlex Manager

The first step to present PowerFlex volumes to the Azure Stack HCI cluster involved creating the volume in PowerFlex Manager. Dell engineering created one volume at a time. This ensured that the volume would be easily identifiable from the Azure Stack HCI cluster.

Create Volume ✕

Number of volumes

Volume name

Provisioning Size

Thin Thick GB

Currently set to: Thin

Select Storage Pool

Cancel Create

Figure 6. Create volume

Next, Dell engineering used PowerFlex Manager to map the new volume to all the Azure Stack HCI cluster nodes. The Azure Stack HCI cluster nodes in Figure 7 are identified by their respective SDC GUIDs.

Map Volumes to Hosts



Select the host/s protocol type.

Select the hosts from the filtered list below to be mapped to the selected volumes.

SDC NVMe

8 Hosts, 8 Selected Show Selected

Filter Host: Add filter Clear Filter

<input type="checkbox"/>	Name	IP Address	Connected	Approved	Identifier	Protocol
<input checked="" type="checkbox"/>	SDC:5031085e00000000	172.16.22.92	Yes	Yes	8E75DA84-0871-6A41-85DC-A7DC2F48049A	SDC
<input checked="" type="checkbox"/>	SDC:5031085f00000001	172.16.21.93	Yes	Yes	C6494385-BA67-864E-BD18-7540B6D8C310	SDC
<input checked="" type="checkbox"/>	SDC:5031086000000002	172.16.22.94	Yes	Yes	FFACC073-43F3-B347-9261-882E821E419F	SDC
<input checked="" type="checkbox"/>	SDC:5031086100000003	172.16.22.95	Yes	Yes	8022720A-DD45-7147-A64B-952562E6EC0F	SDC
<input checked="" type="checkbox"/>	SDC:5031086200000004	172.16.21.96	Yes	Yes	8E7CEC25-0754-3E42-846D-506B04450C01	SDC
<input checked="" type="checkbox"/>	SDC:5031086300000005	172.16.22.97	Yes	Yes	CF8D2AD4-7F50-3743-8FB2-BA71F01651D4	SDC
<input checked="" type="checkbox"/>	SDC:5031086400000006	172.16.22.98	Yes	Yes	1F593BC7-3684-334D-99F9-95CBC779C1CB	SDC
<input checked="" type="checkbox"/>	SDC:5031086500000007	172.16.21.99	Yes	Yes	DF2A6EFC-A2FD-494D-B062-6DFE62F19F00	SDC

Cancel

Map

Figure 7. Azure Stack HCI cluster nodes mapped to new volume

Volumes + Create Volume

Mapping Modify More Actions Showing 1 of 1 Volumes, 1 Selected Show Selected

<input checked="" type="checkbox"/>	Name	Volume ID	Type	Size	Mapped	# Hosts	Hosts Protocol
<input checked="" type="checkbox"/>	Dell_SDS_Volume_01	45478d0600000000	Thin	800 GB	Yes	8	SDC

Figure 8. Volume mapped to eight Azure Stack HCI nodes

Creating the CSV on the Azure Stack HCI cluster

Once the PowerFlex volume was mapped to all the Azure Stack HCI nodes, Dell engineering used Windows Admin Center to create the CSV. Then, they selected Create non-S2D volume in the Windows Admin Center.

Create volume

- Create S2D volume
- Create non-S2D volume

Name *

Dell_SDS_Volume_01

Disk

Disk number: 34 (800 GB)

File system

CSVFS_NTFS

Figure 9. Create CSV in Windows Admin Center

This caused PowerFlex CSV to appear in the list of volumes along with existing S2D CSV.

ClusterPerformanceHistory	OK	ReFS	Three-way mirror	20 GB	A0760R29C02N07	SU1_Pool	9%
Dell_SDS_Volume_01	OK	CSVFS_NTFS	Unknown	800 GB	A0760R29C02N05		0%
Infrastructure_1	OK	CSVFS_ReFS	Three-way mirror	256 GB	A0760R29C02N04	SU1_Pool	32%
Volume-APEXCP	OK	CSVFS_ReFS	Three-way mirror	304 GB	A0760R29C02N03	SU1_Pool	59%

Figure 10. New volume in inventory

Validating volumes

Once volumes were presented to the Azure Stack HCI cluster and confirmed as CSVs, volume health was verified using the following methods:

- From one of the Azure Stack HCI nodes, the following PowerShell command was run:

```
Get-ClusterSharedVolume
```

- For more in-depth information regarding the CSVs, the following PowerShell command was run:

```
Get-ClusterSharedVolumeState | Sort-Object Name,Node | FT
```

The output of this command showed all CSVs and the nodes that they were mapped to. If a redirection was detected, the output would have indicated failure.

Deploying virtual machines onto PowerFlex CSVs

Provisioning VMs using Windows Admin Center onto PowerFlex CSVs was similar to provisioning them onto S2D CSVs. When choosing the location for the VM paths, Dell engineering used the path of the PowerFlex CSV. The C:\ClusterStorage folder on any Azure Stack HCI cluster node showed all the S2D CSVs and PowerFlex CSVs. These paths were then used to create VMs.

New virtual machine

Name *

DellSDSVM01

Generation

Generation 2 (Recommended)

Host

a0760r06c03n04.lab.azsdemo.com (Recommended)

Path ⓘ

C:\ClusterStorage\Dell_SDS_Volume_01 (792 GB availab

Browse

C:\ProgramData\Microsoft\Windows\Hyper-V (Hyper-V Host...

C:\ClusterStorage\Dell_SDS_Volume_01 (792 GB available)

C:\ClusterStorage\Infrastructure_1 (135 GB available)

C:\ClusterStorage\UserStorage_1 (3.71 TB available)

C:\ClusterStorage\UserStorage_2 (4.45 TB available)

C:\ClusterStorage\UserStorage_3 (4.42 TB available)

C:\ClusterStorage\UserStorage_4 (6.82 TB available)

C:\ClusterStorage\Volume-APEXCP (124 GB available)

Hard Disks

Figure 11. Selecting Dell_SDS volume as VM path

New virtual machine

Name *

DellSDSVM01

Generation

Generation 2 (Recommended)

Host

a0760r06c03n04.lab.azsdemo.com (Recommended)

Path ⓘ

C:\ClusterStorage\Dell_SDS_Volume_01 (792 GB availab

Browse

i The VM configuration and virtual hard disks are saved under C:\ClusterStorage\Dell_SDS_Volume_01\Hyper-V\DellSDSVM01

Figure 12. Full VM path visible in Windows Admin Center

For more information about deploying virtual machines on Azure Stack HCI, see [References](#).

Using PowerFlex CSVs for AKS persistent volumes

For instructions on how to create and access persistent volumes mapped to PowerFlex CSVs from an AKS workload cluster, see the Microsoft documentation in [References](#).

Summary and conclusion

Topics:

- [Summary and conclusion](#)
- [We value your feedback](#)

Summary and conclusion

This white paper demonstrates that use cases can be optimized by Dell APEX Cloud Platform for Microsoft Azure with Dell PowerFlex software-defined storage. This includes running a highly diverse set of workloads that are mission critical to your organization and require greater levels of performance and scalability. This solution supports linear scaling of storage resources independently from Azure Stack HCI compute and S2D.

This white paper also describes the processes deployed by the Dell engineering team to test this integration in the Dell labs. The specific compute, storage, and network requirements were shared. This document describes how to create volumes on the PowerFlex cluster and present them to the Azure Stack HCI cluster. Using the PowerFlex integration, you can provision your workloads on S2D or PowerFlex Cluster Shared Volumes based on your application requirements.

NOTE: Dell APEX Cloud Platform for Microsoft Azure with Dell PowerFlex will become available to order and configure when we release our fully supported Dell Validated Design. For more information about ordering availability, contact your Dell account team.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Authors: Graham Barker and Michael Lamia

For additional information and resources, see the [Dell Technologies Info Hub for APEX Cloud Platforms](#).

References

Topics:

- [References](#)

References

The following documentation provides additional information about this solution.

- [Dell APEX Cloud Platform for Microsoft Azure Support site with official user manuals and guides](#)
- [Dell PowerFlex Software technical documentation](#)
- [Using CloudLink with PowerFlex](#)
- [White papers, blogs, videos, and interactive demos about APEX Cloud Platforms](#)
- [How to create virtual machines on Azure Stack HCI](#)
- [Creating AKS workload cluster persistent volumes](#)