

# Microsoft SQL Server 2019 on Windows Hyper - V using PowerFlex

## Abstract

This paper demonstrates the deployment of Microsoft SQL Server 2019 on Microsoft Hyper-V Failover clustering using PowerFlex two-layer deployment architecture.

July 2020

## Revisions

Date	Description
July 2020	Initial release

## Acknowledgements

Author: Swathi Nagaram

Contributors: Robert Sonders, Kevin Jones, Simon Stevens and S Ranjan

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [8/14/2020] [White Paper] [000069]

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents .....	3
Executive summary.....	5
<b>1 Introduction.....</b>	<b>6</b>
1.1 Objective.....	6
1.2 Audience.....	6
1.3 Terminology .....	6
<b>2 Product overview .....</b>	<b>7</b>
2.1 PowerFlex family .....	7
2.2 PowerFlex software components .....	7
2.2.1 PowerFlex .....	7
2.2.2 PowerFlex Manager .....	8
2.3 PowerFlex consumption options.....	8
2.3.1 PowerFlex rack.....	8
2.3.2 PowerFlex appliance .....	8
2.3.3 VxFlex Ready Nodes.....	8
2.4 PowerFlex deployment architectures .....	8
2.4.1 Two-layer architecture.....	8
2.4.2 Single-layer (HCI) architecture.....	9
2.4.3 Mixed architecture.....	9
2.5 Microsoft Hyper-V Failover Clustering.....	9
2.6 Microsoft SQL Server 2019 .....	9
<b>3 Solution architecture.....</b>	<b>11</b>
3.1 Logical architecture .....	11
3.2 Two-layer network topology.....	13
3.3 Storage layout.....	14
3.4 SQL database layout .....	14
<b>4 Deploying Windows Hyper-V Failover Cluster using Windows CO nodes.....</b>	<b>15</b>
4.1 Prerequisites.....	15
4.2 Provisioning PowerFlex volume to Windows CO nodes .....	16
4.3 Configuring Hyper-V Failover cluster using PowerShell commands.....	17
<b>5 Conclusion.....</b>	<b>24</b>

## Table of contents

<b>A</b>	<b>Appendix</b> .....	<b>25</b>
A.1	Hardware and Software components .....	25
A.2	Best Practices.....	26
<b>B</b>	<b>Technical support and resources</b> .....	<b>28</b>
B.1	Related resources .....	28

## Executive summary

With Microsoft SQL Server 2019, it is possible to build intelligent, mission-critical applications using a scalable, hybrid database platform that has everything built-in, from in-memory performance, and advanced security to database analytics.

Many organizations demand uninterrupted operations with minimum downtime and SLAs in services, that can be achieved using Microsoft Hyper-V Failover Clustering with PowerFlex.

PowerFlex is a software-defined storage platform designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high-performance storage resources in a managed unified fabric. PowerFlex comes in flexible deployment options that enables disaggregated (two-layer), HCI (single-layer), or mixed architectures.

This white paper focuses on the deployment details of Microsoft SQL Server 2019 on Windows Server Hyper-V Failover Clustering using Dell EMC PowerFlex two-layer deployment architecture. In a two-layer architecture, nodes that provide storage capacity and host datasets are separated from nodes that host applications and workloads. This segregation of compute and storage resources can be helpful to minimize software licensing costs in certain situations. This architecture is suitable for hosting high-performance high-value databases and application workloads.

# 1 Introduction

## 1.1 Objective

The primary focus of this White Paper is to provide the reader with details of Windows Hyper-V Failover Clustering and its prerequisites. This paper also provides best practices to be followed while deploying the Microsoft SQL Server 2019 on Dell EMC PowerFlex.

## 1.2 Audience

This White Paper is intended for Microsoft SQL Server database administrators, system engineers, partners, and members of Dell EMC and partner professional service community who are looking for deploying SQL Server database on Dell EMC PowerFlex rack using two-layer deployment architecture with Microsoft Hyper-V hypervisor.

Readers are expected to have an understanding and working knowledge of Dell EMC PowerFlex rack, PowerFlex Manager, PowerFlex (**previously VxFlex OS**), Microsoft PowerShell commands, Microsoft Failover Cluster Manager and Microsoft SQL Server database administration.

## 1.3 Terminology

The following table defines acronyms and terms that are used throughout this document:

Table 1 Terms and definitions

Term	Definition
MDM	Meta Data Manager
SDC	Storage Data Client
SDS	Storage Data Server
SVM	Storage Virtual Machine
OLTP	Online Transaction Processing
RCM	Release Certification Matrix
VLAN	Virtual Local Area Network
TOR	Top of Rack Switch
CO	Compute Only Node
VM	Virtual Machine
GUI	Graphical User Interface

## 2 Product overview

### 2.1 PowerFlex family

PowerFlex is a software-defined storage platform designed to significantly reduce operational and infrastructure complexity, empowering organizations to move faster by delivering flexibility, elasticity, and simplicity with predictable performance and resiliency at scale. The PowerFlex family provides a foundation that combines compute as well as high performance storage resources in a managed unified fabric. PowerFlex comes in flexible deployment options - rack, appliance or ready nodes - that enables disaggregated (two-layer), HCI (single-layer), or mixed architectures. PowerFlex is ideal for high performance applications and databases, building an agile private cloud, or consolidating resources in heterogeneous environments.

## PowerFlex family: Flexibility Unleashed

**PowerFlex**  
Software-defined block-storage service that enables a scale-out storage infrastructure using X86 hardware nodes

<p><b>PowerFlex rack</b></p> <p>Compute &amp; high performance storage with integrated networking at rack scale</p> 	<p><b>PowerFlex appliance</b></p> <p>Compute &amp; high performance storage</p> 	<p><b>VxFlex Ready Node</b></p> <p>Compute &amp; high performance storage with customer configurability</p> 
<p><b>PowerFlex Manager</b> Comprehensive IT Operations Management (ITOM) software</p>		<p><b>Element Managers</b> Manage and monitor node clusters</p>

Figure 1 PowerFlex family

### 2.2 PowerFlex software components

Software is the key differentiation and the “secret sauce” in the PowerFlex offering. The PowerFlex software components not only provide software-defined storage services, they also help simplify infrastructure management and orchestration with comprehensive ITOM and LCM capabilities that span compute as well as storage infrastructure, from BIOS and firmware to nodes, software and networking.

The core foundational component in the PowerFlex family that enables Software Defined Storage (SDS) services is called PowerFlex, to represent the core value it enables for the platform. Additionally, PowerFlex Manager is a comprehensive IT Operational Management (ITOM) and Life Cycle Management (LCM) tool that drastically simplifies management and ongoing operation.

#### 2.2.1 PowerFlex

PowerFlex (previously VxFlex OS) is the software foundation of PowerFlex software-defined storage. It is a scale-out block storage service designed to deliver flexibility, elasticity, and simplicity with predictable high performance and resiliency at scale. PowerFlex management is available using a GUI, CLI, and REST clients.

## 2.2.2 PowerFlex Manager

PowerFlex Manager is the software component in PowerFlex family that enables ITOM automation and lifecycle management capabilities for PowerFlex systems.

## 2.3 PowerFlex consumption options

PowerFlex SDS platform is available in multiple consumption options to help customers meet their project and datacenter requirements. PowerFlex appliance and PowerFlex rack provide customers comprehensive IT Operations Management (ITOM) and lifecycle management (LCM) of the entire infrastructure stack in addition to sophisticated high-performance, scalable, resilient storage services. PowerFlex appliance and PowerFlex rack are the two preferred and proactively marketed consumption options. PowerFlex is also available on VxFlex Ready Nodes without the ITOM and LCM capabilities. Please note that the brand for Ready Nodes continues to be PowerFlex.

### 2.3.1 PowerFlex rack

PowerFlex rack is a software-defined storage platform designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high performance storage resources in a managed unified network. This rack-scale engineered system, with integrated networking, enables customers to achieve the scalability and management requirements of a modern data center.

### 2.3.2 PowerFlex appliance

PowerFlex appliance is a software-defined storage platform designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high performance storage resources in a managed unified network. This turnkey offer allows customers the flexibility and savings to bring their own compatible networking. With PowerFlex, customers deploy to match their initial needs and easily expand with massive scale potential, without having to compromise on performance and resiliency.

### 2.3.3 VxFlex Ready Nodes

VxFlex Ready Nodes are validated server building blocks configured for use with PowerFlex. They are available with thousands of configuration options and are available for customers who prefer to build their own environments.

---

**Note: The brand for Ready Nodes continues to be VxFlex.**

---

## 2.4 PowerFlex deployment architectures

PowerFlex software-define storage offers flexibility of deployment architecture to help best meet the specific deployment and architectural requirements. PowerFlex can be deployed in a two-layer (Server SAN), single-layer (HCI), or in storage-only architectures.

### 2.4.1 Two-layer architecture

In this White Paper solution, PowerFlex is deployed using two-layer architecture.

In a two-layer architecture, nodes that provide storage capacity and host datasets are separated from nodes that host applications and workloads. PowerFlex manager provides LCM and IOTM for the entire

infrastructure, including nodes that provide storage and nodes that host the applications. Compute and storage resources can be scaled by adding respective node to the cluster. This segregation of compute and storage resources can be helpful to minimize software licensing costs in certain situations. This architecture could be most suitable for hosting high-performance high-value databases and application workloads.

### 2.4.2 Single-layer (HCI) architecture

In this architecture, each node in the cluster contributes storage resources as well as hosts applications and workloads. This architecture allows you to scale your infrastructure uniformly and with a pre-defined building block that adds both storage and compute resources. PowerFlex Manager provides ITOM and LCM capabilities for the entire infrastructure. This architecture is most suitable for datacenter and workload consolidation.

### 2.4.3 Mixed architecture

In this architecture, PowerFlex provides only storage resources. Using PowerFlex storage-only nodes, a software-defined block storage environment is created that can be accessed and consumed by a number of applications and workloads that are hosted outside of PowerFlex cluster. PowerFlex Manager provides LCM and ITOM for the storage infrastructure. This is a suitable architecture where customer has existing compute infrastructure but needs high-performance SDS. This can be a starting point with the customer, and may expand to a two-layer Server SAN deployment in the future as the external compute is migrated to PowerFlex.

## 2.5 Microsoft Hyper-V Failover Clustering

Microsoft Hyper-V Failover clustering offers high availability and scalability. Whenever a host fails, VMs running on that can be restarted onto another healthy host, which allows business critical application to be brought back up very quickly with minimal disruptions, it also offers flexibility for the host that needs downtime or a reboot for any planned maintenance activity.

Additionally, Failover clusters provide Cluster Shared Volume (CSV) allowing multiple nodes in a failover cluster to access the same LUN provisioned with NTFS for read /write operation simultaneously.

For completed details pertaining to new functionalities implemented for Windows Server 2019 Failover clustering, see [Failover Clustering](#).

## 2.6 Microsoft SQL Server 2019

Microsoft® SQL Server® 2019 delivers breakthrough mission critical capabilities with in-memory performance and operational analytics that are built in while comprehensive security features help protect your data at-rest, data in-motion, and now support for built-in big data.

The SQL Server 2019 release adds new security features, querying capabilities, R analytics, Hadoop and cloud integration, and so on. Besides, numerous improvements and enhancements.

The following are some of the new features provided by Microsoft SQL Server 2019 release:

- Mission-critical security using Always Encrypted with secured enclave.
- Improved High Availability.
- Data virtualization with Big Data Clusters.
- Intelligent database with In-Memory database, intelligent query processing, and performance.
- Improved developer experience and monitoring capabilities.
- Provides flexibility to choose platform of your choice with Linux and Containers support.

For more details about new features and enhancements, see [Microsoft SQL Server 2019 Documentation](#).

### 3 Solution architecture

For this paper, the Microsoft SQL Server 2019 solution is deployed on a two-layer (Server SAN) deployment architecture using compute-only nodes with Windows Hyper-V hypervisor and dedicated storage only nodes with PowerFlex to provide the required storage capacity.

#### 3.1 Logical architecture

The following diagram shows logical architecture of Microsoft SQL Server 2019 deployment on Dell EMC PowerFlex rack two-layer setup that had four SDS and three SDC.




Figure 2 Solution architecture

A PowerFlex storage cluster was created from four RHEL storage only nodes configured in a single protection domain that played the role of SDS. Each host is populated with ten 1.92 TB SSD drives. These 40 disks created a storage pool from which the usable storage volumes were carved out.

PowerFlex Manager supports the deployment of a compute-only service with Microsoft Windows Server 2019. This applies to compute-only (CO) nodes on PowerFlex appliance or PowerFlex rack. PowerFlex Manager also installs the SDC to the nodes in order to access the volumes created in the storage pool and enables the Microsoft Hyper-V service. However, the setup is incomplete and so the user must create Windows Hyper-V Failover Cluster using CLI or GUI.

These PowerFlex volumes are provisioned as Clustered Shared Volumes to Windows CO hosts.

One SQL Server VM on each host with 16 vCPU and 64 GB RAM was created, and the solution was validated by simulating real time OLTP workloads using TPC-C benchmark tool.

For detail configuration on nodes, see [Appendix A.1](#).

### 3.2 Two-layer network topology

The following figure demonstrates high-level network architecture of the two-layer deployment option on PowerFlex rack system:




Figure 3 Network architecture

Table 2 PowerFlex rack networking details

Term	Definition
Cisco Nexus 93180YC-EX	10 Gbps & 25 Gbps TOR switches
Cisco Nexus 9332PQ	25 Gbps Aggregation switches
Cisco Nexus 3172TQ	1 Gbps & 10 Gbps Management switches
PowerFlex storage traffic	2 x 25 Gbps links

### 3.3 Storage layout

Each storage only node includes two Intel Xeon gold 12-core processors, 192 GB RAM, and ten 1.92 TB SSDs. A PowerFlex storage cluster was created from four Red Hat Enterprise Linux 7.6 storage only nodes configured in a single protection domain that played the role of SDS.

A single storage pool was configured, and multiple logical volumes were created to meet the SQL Server database requirements including volumes for data, log etc. These volumes were mapped to the Hyper-V cluster as disks drive to SQL Server virtual machine using Paravirtual SCSI (PVSCSI) adapters.

For detailed configuration of PowerFlex nodes, see [Appendix A.1](#).

### 3.4 SQL database layout

For this solution, Microsoft SQL Server 2019 was running on Windows Server 2019 guest VMs consisting of 16 virtual CPUs, 64 GB RAM, and thin-provisioned disks in the layout as shown in the following table:

Table 3 SQL layout

Disk size (GB)	Drive	Disk Purpose
90	C	Windows operating system disk
100	D	Database disk1
102	E	Database disk2
104	F	Database disk3
106	G	Database disk4
150	H	DB Log1
150	L	DB Log2
120	T	TempDB

---

**Note: The vCPUs and memory can be modified based on the customer use cases.**

---

After the drives are provisioned, the SQL Server data and log drives were formatted with a 64 KB NTFS cluster size. This size optimizes I/O performance with no overhead and offers a good balance between flexibility, performance, and ease of use. The operating system and the SQL binary drives use the standard 4 KB NTFS cluster size.

## 4 Deploying Windows Hyper-V Failover Cluster using Windows CO nodes


For this solution, the deployment process is automated with PowerFlex Manager.

1. Deployed each R740 compute node by installing or enabling the following options:
  - Windows Server 2019 Data Center version.
  - Hyper-V service in the Windows Server configuration.
  - PowerFlex SDC component.
2. Created Microsoft Hyper-V Failover clustering using PowerShell commands.
3. Created and managed SQL VMs using Failover Cluster Manager.


### 4.1 Prerequisites

PowerFlex Manager creates a NIC team between two NIC ports and assigns the management VLAN to it during a Windows compute-only standalone deployment process. It is recommended to untag the VLAN and make it to default before creating a Windows Failover cluster.

1. Go to **Server Manager > Local Server > NIC Teaming > Team Interfaces**.



2. Right-click **Team Interfaces** and select **Properties > Change to Default**.



3. Click **OK** to apply the new changes.

## 4.2 Provisioning PowerFlex volume to Windows CO nodes

1. Open the VxFlex OS 3.1 GUI click front-end and select **SDC**.
2. Windows-based compute-only nodes are listed as **SDCs** if configured correctly.
3. Click **front-end** again, select **Volumes**.
4. Right-click a storage pool of your choice and select **Add Volume**.
5. Provide the following details for the new volume and click **OK** to proceed.
  - Name: Windows\_vol1
  - Size in GB:3000
  - Provisioning: Thick
  - **(Optional)** Use Read RAM Cache
6. Right-click the volume that has been created (Windows\_vol1 as shown in example here) and select **Map**.
7. In the Select Nodes panel, select the Windows-based compute-only nodes and then select **Map Volumes**.




Figure 4 Mapping volumes

### 4.3 Configuring Hyper-V Failover cluster using PowerShell commands

Use respective PowerShell commands to configure the Hyper-V Failover cluster.

1. Install Failover Clustering Windows feature.

```
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools

Success Restart Needed Exit Code Feature Result
-----
True Yes SuccessRest... {Failover Clustering, Failover Clustering ...
WARNING: You must restart this server to finish the installation process.
```

2. Install Hyper-V management tools.

```
Install-WindowsFeature -Name Hyper-V -IncludeManagementTools -Restart

Success Restart Needed Exit Code Feature Result
-----
True Yes SuccessRest... {Hyper-V, Hyper-V Module for Windows Power...
WARNING: You must restart this server to finish the installation process.
```

3. Verify the roles and features installed.

```
Get-WindowsFeature
```

4. Rename the NIC teaming adapter created by PowerFlex Manager to **PhysicalTeam**.

```
Rename-NetAdapter -Name "ESX_Mgmt_105 - VLAN 105" -NewName "PhysicalTeam - VLAN 105"
```

5. Create a Hyper-V virtual switch using the team and enable QoS in Weight.

```
New-VMSwitch -Name TeamSwitch -NetAdapterName "PhysicalTeam - VLAN 105" - MinimumBandwidthMode Weight -AllowManagementOS 0
```

6. Configure a default bandwidth for the switch.

```
Set-VMSwitch -Name "TeamSwitch" -DefaultFlowMinimumBandwidthWeight 0
```

7. Configure the virtual network adapters on the virtual switch for management operating system.

```
Add-VMNetworkAdapter -ManagementOS -Name Management -SwitchName TeamSwitch
```

8. Set the VLAN associated with the adapter and configure the VMQ weight and minimum bandwidth weight.

```
>Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Management - Access -VlanId 105  
>Set-VMNetworkAdapter -ManagementOS -Name Management -VmqWeight 80 - MinimumBandwidthWeight 10
```

9. Configure the static IP, gateway and DNS for the management network.

```
>Set-NetIPInterface -InterfaceAlias "vEthernet (Management)" -dhcp Disabled  
>New-NetIPAddress -PrefixLength 24 -InterfaceAlias "vEthernet (Management)" -IPAddress 192.168.105.55 -DefaultGateway 192.168.105.254  
>Set-DnsClientServerAddress -InterfaceAlias "vEthernet (Management)" - ServerAddresses 192.168.105.250
```

10. Configure the virtual network adapters on the virtual switch for virtual machine migration between nodes and set the VLAN associated with the adapter.

```
>Add-VMNetworkAdapter -ManagementOS -Name Migration -SwitchName TeamSwitch
>Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Migration -
Access -VlanId 106
```

11. Configure the VMQ weight and minimum bandwidth weight and assign the static IP for the migration network.

```
>Set-VMNetworkAdapter -ManagementOS -Name Migration -VmqWeight 90 -
MinimumBandwidthWeight 40
>Set-NetIPInterface -InterfaceAlias "vEthernet (Migration)" -dhcp Disabled
>New-NetIPAddress -PrefixLength 24 -InterfaceAlias "vEthernet (Migration)" -
IPAddress 192.168.106.61
```

12. Configure the virtual network adapters on the virtual switch for private cluster traffic and set the VLAN associated with the adapter.

```
>Add-VMNetworkAdapter -ManagementOS -Name Cluster -SwitchName TeamSwitch
>Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Cluster -
Access -VlanId 107
```

13. Configure the VMQ weight and minimum bandwidth weight and assign the static IP for the cluster network.

```
>Set-VMNetworkAdapter -ManagementOS -Name Cluster -VmqWeight 80 -
MinimumBandwidthWeight 10
>Set-NetIPInterface -InterfaceAlias "vEthernet (Cluster)" -dhcp Disabled
>New-NetIPAddress -PrefixLength 24 -InterfaceAlias "vEthernet (Cluster)" -
IPAddress 192.168.107.61
```

14. Configure the ports dedicated to PowerFlex data and enable jumbo frames only if the switches are also configured to use them.

```
>New-NetIPAddress -PrefixLength 24 -InterfaceAlias VxFlexData1 -IPAddress
192.168.152.61
>New-NetIPAddress -PrefixLength 24 -InterfaceAlias VxFlexData2 -IPAddress
192.168.160.61
>Get-NetAdapterAdvancedProperty -Name VxFlexData* -DisplayName "Jumbo
Packet" | Set-NetAdapterAdvancedProperty -RegistryValue 9014
```


15. Configure Windows update, download and install all updates.

```
>Install-Module PSWindowsUpdate  
>Get-WindowsUpdate  
>Install-WindowsUpdate
```

16. Add the host to Windows domain and rename the host.

```
>Add-Computer -DomainName FLEXVDI-AD.flexvdi.net -DomainCredential  
flexvdi\administrator  
>Rename-Computer -NewName WinHost1 -LocalCredential administrator -  
DomainCredential flexvdi\administrator
```

You will be prompted to enter your admin credentials.



17. Reboot the node.

```
shutdown.exe /r /t 0
```

18. Log in with the new domain credentials and enable Remote Desktop.

```
>Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Control\Terminal  
Server' -Name "fDenyTSConnections" -Value 0  
>Enable-NetFirewallRule -DisplayGroup "Remote Desktop"
```

19. Open firewall ports to enable ping.


```
>New-NetFirewallRule -DisplayName "Allow inbound ICMPv4" -Direction Inbound
-Protocol ICMPv4 -IcmpType 8 -Action Allow
>New-NetFirewallRule -DisplayName "Allow inbound ICMPv6" -Direction Inbound
-Protocol ICMPv6 -IcmpType 8 -Action Allow
```

**Note: Repeat all the above steps on other two Windows CO nodes.**

20. Once all the nodes are ready, perform cluster validation. Ensure all are in success state, if any issues are encountered fix it and perform the validation step.

```
Test-Cluster WinHost1, WinHost2, WinHost3

Mode lastWriteTime LengthName
-----
-a--- 4/7/2020 6:33AM 1425044 Validation Report 2020.04.07 At 06.29.07.htm
```



21. Create cluster.

```
New-Cluster -name Win2019Cluster -Node WinHost1, WinHost2, WinHost3 -
StaticAddress 192.168.105.200

Name
-----
Win2019cluster
```

## 22. Rename the cluster networks, if needed.

```
>Get-ClusterNetwork | Where-Object {$_.Address -eq
"192.168.105.0"}).Name = "Management_Network"
>Get-ClusterNetwork | Where-Object {$_.Address -eq
"192.168.106.0"}).Name = "Migration_Network"
>Get-ClusterNetwork | Where-Object {$_.Address -eq
"192.168.107.0"}).Name = "Cluster_Network"
>Get-ClusterNetwork | Where-Object {$_.Address -eq
"192.168.152.0"}).Name = "VxFlexData1"
>Get-ClusterNetwork | Where-Object {$_.Address -eq
"192.168.160.0"}).Name = "VxFlexData2"
```

## 23. Configure the cluster network roles.

```
>Get-ClusterNetwork -Name "Management_Network").Role = 3
>Get-ClusterNetwork -Name "Cluster_Network").Role = 1
>Get-ClusterNetwork -Name "Migration_Network").Role = 1
>Get-ClusterNetwork -Name "VxFlexData1").Role = 0
>Get-ClusterNetwork -Name "VxFlexData2").Role = 0
```

## 24. Remove all networks except Migration network from Live Migration Settings / Configure the live migration network.

```
Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -
Name MigrationExcludeNetworks -Value ([String]::Join(";", (Get-
ClusterNetwork | Where-Object {$_.Name -ne "Migration_Network"}).ID))
```

## 25. List the disks within Windows Server that was mapped using VxFlex OS GUI.

```
Get-Disk -FriendlyName "*EMC ScaleIO*"

Number Friendly Name Serial Number HealthStatus OperationalStatus Total
Size Partition Style
-----
-----
1      EMC ScaleIO EMC-447c94612b4dfc0f-d67687b0... Healthy      Offline
2.93 TB      RAW
```

26. Initialize the disks and format with Allocation unit 64k as recommended for Windows Cluster Shared Volumes.

```
Get-Disk -FriendlyName "*EMC ScaleIO*" | Initialize-Disk -PassThru |
New-Partition -UseMaximumSize | Format-Volume -AllocationUnitSize 65536
```

DriveLetter	FriendlyName	FileSystemType	DriveType	HealthStatus
OperationalStatus	SizeRemaining	Size		
-----	-----	-----	-----	-----
		NTFS	Fixed	Healthy
2.93 TB	2.93 TB			OK

27. List the disks available for use by the cluster.

```
Get-ClusterAvailableDisk
```

```
Cluster      : Win2019Cluster
Id           : 3485dc14-5b7b-452e-a263-93b121e24968
Name        : Cluster Disk 1
Number      : 2
Size        : 3221225472000
Partitions  : {\?\GLOBALROOT\Device\Harddisk2\Partition2\}
```

28. Add the disks to cluster.

```
Get-ClusterAvailableDisk | Add-ClusterDisk
```

Name	State	OwnerGroup	ResourceType
-----	-----	-----	-----
Cluster Disk 1	Online	Available Storage	Physical Disk

29. Add Cluster Disk 1 as Clustered Shared Volume.

```
Add-ClusterSharedVolume -Name "Cluster Disk 1"
```

30. Create a directory on the Cluster Shared Volume to host the Virtual Machines.

```
mkdir C:\ClusterStorage\Volume1\Hyper-V
```

## 5 Conclusion

This White Paper demonstrates the deployment of SQL Server 2019 on Dell EMC PowerFlex two-layer architecture along with the details of Windows Hyper-V Failover Clustering and relevant prerequisites. In addition, it states the best practices to be followed while deploying the Microsoft SQL Server 2019 on Dell EMC PowerFlex to meet performance, resiliency, and scale.

The segregation of compute and storage resources can be helpful to minimize software licensing costs in certain situations. This architecture is particularly suitable for hosting high-performance databases and application workloads to achieve optimal and fast performance.

## A Appendix

### A.1 Hardware and Software components

Table 4 Storage-only node details

Hardware	Configuration
Node	PowerEdge R640
CPU	2 x Intel® Xeon® Gold 6126 CPU @ 2.60 GHz, 12 cores
Memory	12 x 16 GB DDR-4
NIC	2x Mellanox ConnectX-4 LX 25GbE SFP Adapter
Storage	Dell HBA330 controller 10 x 1.92 TB SAS SSD
Operating system	RHEL 7.6
VxFlex OS (rebranded as PowerFlex from version 3.5 onwards)	3.0.1

Table 5 Compute-only node details

Hardware	Configuration
Node	PowerEdge R740
CPU	2 x Intel® Xeon® Gold 6240 CPU @ 2.60 GHz, 18 cores
Memory	24 x 32 GB DDR-4
NIC	2x Mellanox ConnectX-4 LX 25GbE SFP Adapter
Operating system	Microsoft Windows Server 2019

Table 6 SQL VM details

Hardware	Configuration
CPU cores	16
Memory	64
Operating system	Microsoft Windows Server 2019
SQL Server Release	Microsoft SQL Server 2019

## A.2 Best Practices

The following sections outline the best practices followed in this solution:

### **PowerFlex rack**

- Minimum number of nodes for a production workload is 8.
- Homogenous node types are recommended for predictable performance.
- Maximum number of devices in a storage pool is 300.
- Recommended maximum number of nodes in a protection domain is 32.
- Change the passwords for all default accounts.
- Use secure communication - HTTPS (TCP port 443) to remotely access PowerFlex nodes.
- Ensure PowerFlex rack is compliant to an RCM.

### **PowerFlex**

- Configure high-performance profile for MDM, SDS, and SDC.
- Disable Read Flash cache and Read RAM cache for all flash clusters.
- Check with PowerFlex platform team to increase the per device queue length value to 256 per host for improving the I/O concurrency.
- Ensure that the customize power plan is set to High Performance.

### **PowerFlex rack network**

- Use the Para virtual SCSI (PVSCSI) controller on guest VMs to achieve high performance.
- Enable secure network protocol options only (for example, HTTPS and Secure Shell (SSH)).
- Avoid autonomous certificate deployments to ones that are fully integrated with site trust infrastructures and train people to not accept self-signed certificates.

### **Windows Hyper-V Failover Clustering**

- Enable Virtual machine Queue (VMQ) on Physical NIC for efficient network connectivity.
- Create Separate networks for traffic types such as management, cluster, migration.
- Enable jumbo frames where possible.
- Use Static IP addresses for network configuration.
- Use redundant NICs to avoid single points of failure.
- Do not enable Management traffic on new virtual switches created.
- Provide dedicated networks for PowerFlex data.
- Follow recommended networking best practices for Microsoft Hyper-V failover clustering.

## **SQL Server 2019**

The following best practices were used for the standard version of SQL Server 2019 VM configurations:

### **General**

- Perform a current state analysis to identify workloads and sizing.
- Start with a proof of concept, and then test, optimize, iterate, and scale.

### **RAM**

- More RAM can increase SQL database read performance.
- From the total Windows operating system memory, reserve 4 GB to the operating system itself. Configure SQL Server memory as per suggestion provided by Microsoft. See, [Server Memory Configuration](#).
- Size each VM to fit within a NUMA node's memory footprint.

### **vCPUs**

- Do not over allocate vCPUs to VMs.
- At virtual level, 1 socket has 8 CPU cores.

### **Drive configuration**

- Distribute databases and logfiles across multiple VMDKs.
- Distribute vDisks across four SCSI controllers.
- Use 64 KB NTFS allocation for database and log drives.
- Size for at least 20 percent free disk space on all drives.
- Create drives of slightly different sizes.

### **SQL Server datafiles**

- Do not shrink databases as it causes severe Index fragmentation.

### **SQL Server logfiles**

- Under most circumstances, one log per database (including TempDB) should be enough.
- Log files fill sequentially, so extra files do not increase performance.

### **TempDB**

- Use multiple TempDB files of the same size.
- If cores < 8, the number of TempDB files = cores.
- If cores > 8, start with eight TempDB files and monitor for performance.
- One TempDB drive should be enough for most environments.
- If NVMe storage is available, put TempDB into it.

## B Technical support and resources

[Hammer DB](#)

[Microsoft SQL Server 2019](#)

[Failover Clustering.](#)

### B.1 Related resources

[Dell.com/support](#)

[Storage technical documents and videos](#)