

Dell EMC VxFlex for Microsoft SQL Server 2019

Abstract

This paper highlights the benefits of hosting SQL Server 2019 on Dell EMC VxFlex systems.

January 2020

Revisions

Date	Description
January 2019	Initial release

Acknowledgements

This paper was produced by the following:

Author: Kailas Goliwadekar

Support: Robert Sonders, Shashikiran Chidambara

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, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [1/14/2020] [White Paper] [000061]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
1.1 Terminology.....	5
1.2 Objective.....	6
1.3 Audience.....	6
2 Product overview.....	7
2.1 VxFlex family	7
2.1.1 VxFlex integrated rack.....	7
2.1.2 VxFlex appliance	8
2.1.3 VxFlex Ready Nodes	8
2.1.4 VxFlex OS.....	8
2.1.5 VxFlex Manager.....	9
2.2 Microsoft SQL Server 2019	10
3 Microsoft SQL Server 2019 on VxFlex - Solution architecture.....	11
3.1 Logical architecture	11
3.2 Networking layout	12
3.3 Storage layout.....	13
3.4 SQL database layout	13
4 Best practices.....	14
4.1 VxFlex integrated rack.....	14
4.2 VxFlex OS.....	14
4.3 VxFlex integrated rack network	14
4.4 SQL Server 2019.....	14
5 Conclusion.....	16
6 Technical support and resources	17
6.1 Related resources.....	17
6.2 Additional resources	17
A Appendix.....	18
A.1 VxFlex integrated rack cluster details.....	18

Executive summary

In today's enterprise database environments such as Microsoft SQL Server 2019, organization demand uninterrupted operations with minimum downtime and SLAs. The hyperconverged infrastructure platform facilitates improved flexibility, productivity, performance, protection, and ease of infrastructure management.

The VxFlex portfolio offers "Transformation without Trade-offs". VxFlex makes it possible to transform from a traditional three-tier architecture to a modern data center without any trade-offs between performance, resilience or flexibility. VxFlex Ready Nodes powered by Intel® Xeon® Scalable Processors are for organizations wanting a full build experience. The VxFlex appliance offers a turnkey experience with components pre-validated and tested, while allowing organization's to 'bring your own' networking. VxFlex integrated rack offers a turnkey experience at rack-scale with integrated networking.

VxFlex systems handle a wide variety of production workloads, both traditional and cloud-native workloads, and customers are steadily increasing the number and, due to its unparalleled flexibility, the range of production workloads they run on their VxFlex systems. This paper highlights in detail the best practice of deploying Microsoft SQL Server 2019 on Dell EMC VxFlex integrated system.

1 Introduction

Modern data center workloads have varying business value and characteristics for the workload and data that governs the performance, throughput, capacity, availability, data protection, and data services requirements. Shrinking IT budgets, push for greater efficiency, and consolidation and workload requirements have made it necessary for the underlying infrastructure to deliver high performance, scalability, resiliency, and most importantly -- flexibility. VxFlex integrated rack is an engineered system for Dell EMC designed on five super power principals to meet the key infrastructure requirements. VxFlex integrated rack delivers:

- Unmatched **performance**.
- Unprecedented **scale** (1000 nodes and 100s-PB storage capacity).
- Built in redundant hardware components and VxFlex OS mesh mirror architecture delivers unparalleled **resiliency**.
- Infrastructure **flexibility**: VxFlex integrated rack is second to none. A VxFlex integrated rack system has 1000s of hardware and software configuration option combinations that can co-exist freeing customers from T-shirt size, dedicated and siloed environments, and accelerating the data center consolidation.
- **Engineered system with single call support and life cycle management.**

This white paper outlines how you can deploy Microsoft SQL server on VxFlex systems to meet performance, resiliency, scale, and availability requirements and take full advantage of the five integrated super powers. The automated deployment of VxFlex cluster was done using VxFlex Manager.

1.1 Terminology

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

Term	Definition
MDM	Meta Data Manager
SDC	Storage Data Client
SDS	Storage Data Server
SVM	Storage Virtual Machine
FG	Fine Granularity
MG	Medium Granularity
TPM	Transactions Per Minute
OLTP	Online Transaction Processing
DSS	Decision Support System
RHEL	Red Hat Enterprise Linux
RCM	Release Certification Matrix

1.2 Objective

This paper shows:

- Overview of VxFlex family
- An overview of SQL Server 2019 on VxFlex integrated rack use case
- Solution architecture for SQL server 2019 on VxFlex integrated rack
- VxFlex integrated rack and SQL Server 2019 deployment best practices
- Conclusion

1.3 Audience

This white paper is intended for 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 VxFlex integrated rack with ESXi hypervisor.

The reader of this document must have a working knowledge of the following technologies:

- Dell EMC VxFlex integrated rack
- VMware vSphere
- SQL Server database administration

2 Product overview

2.1 VxFlex family

VxFlex family helps transforming from a traditional three-tier architecture to a modern data center without any trade-offs, meeting business requirements without compromise. The central software layer for VxFlex family is VxFlex OS, scale-out block storage service that enables customers to create a scale-out Server SAN or hyperconverged infrastructure. The VxFlex family currently include VxFlex Ready Nodes, VxFlex appliance, and VxFlex integrated rack.

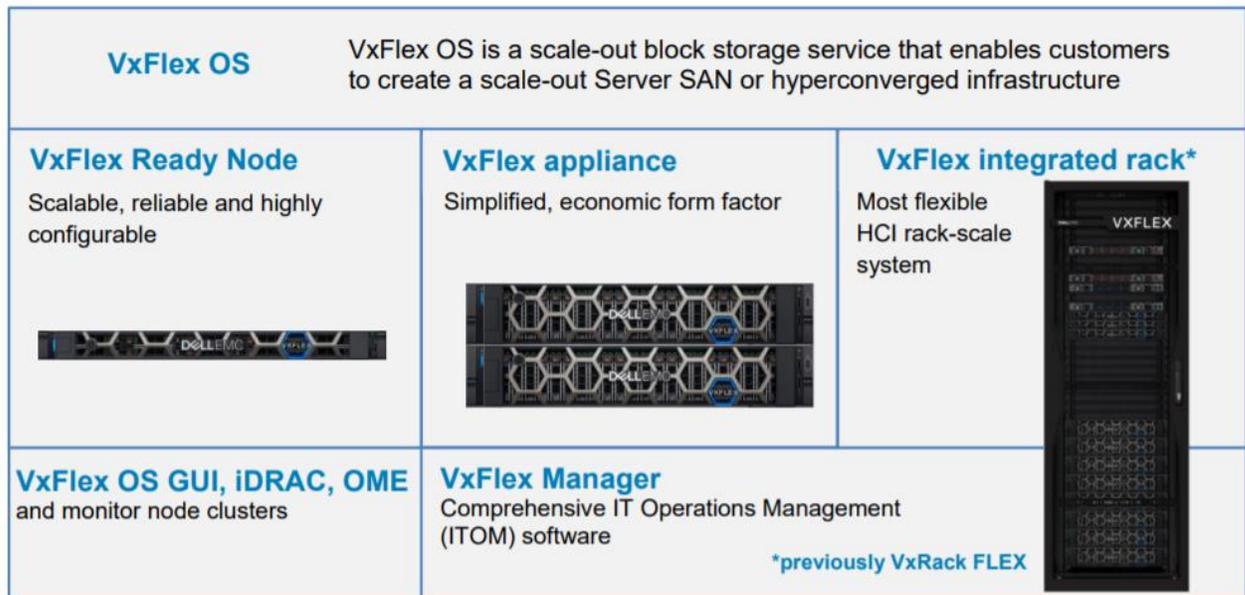


Figure 1 VxFlex family

2.1.1 VxFlex integrated rack

VxFlex integrated rack is a rack-scale engineered system, with integrated networking, that provides linear scalability and enterprise-grade availability. VxFlex integrated rack is engineered, manufactured, managed, supported, and sustained as one system for single end-to-end lifecycle support.

VxFlex integrated systems create a server-based SAN by combining virtualization software, which is known as VxFlex OS, with Dell EMC PowerEdge servers to deliver flexible, scalable performance and capacity on demand. Local storage resources are combined to create a virtual pool of block storage with varying performance tiers. The VxFlex integrated rack enables you to scale from a small environment to enterprise scale with over a thousand nodes. In addition, it provides enterprise grade data protection, multi-tenant capabilities, and add-on enterprise features such as QoS, thin provisioning, and snapshots. VxFlex systems deliver the performance and time-to-value required to meet the demands of the modern enterprise data center.

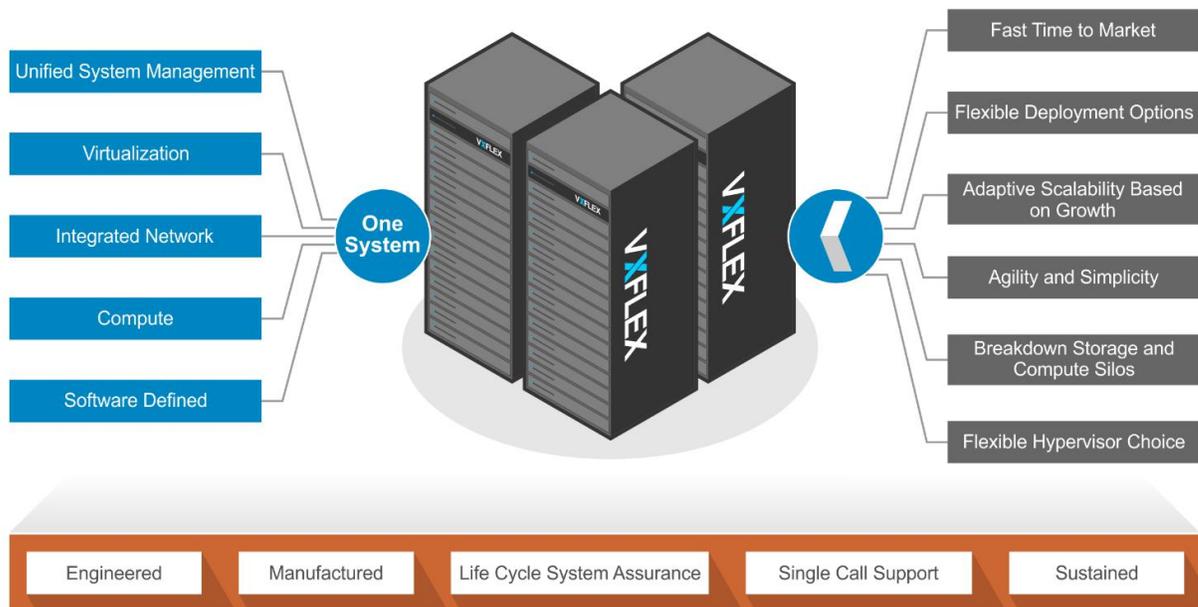


Figure 2 VxFlex integrated rack benefits

2.1.2 VxFlex appliance

VxFlex appliance is a preconfigured and validated for fast, easy deployment, VxFlex appliance offers a turnkey experience in an economic form factor. With VxFlex appliance, customers benefit from a smaller starting point, with massive scale potential, without having to compromise on performance and resiliency.

2.1.3 VxFlex Ready Nodes

VxFlex Ready Nodes combine Dell EMC PowerEdge servers that are powered by Intel® Xeon® Scalable Processors and VxFlex OS software to create scalable, reliable, and easy-to-deploy building blocks for hyperconverged or server SAN architecture, multi-hypervisor or bare metal environments, and high-performance databases.

2.1.4 VxFlex OS

VxFlex OS applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. VxFlex OS abstracts the local storage out of each contained within each server, including HDDs, SSDs, and all-flash. VxFlex OS uses three lightweight pieces of software to create, consume, and coordinate the storage layer in VxFlex systems.

VxFlex OS can be deployed in the following ways:

- **Two-layer:** In a two-layer (also known as storage-only) deployment, compute resources exist on one set of nodes, and storage resources exist on another set of nodes.
- **Hyperconverged (HCI):** The SDCs and SDSs run on the same set of nodes.
- **Hybrid:** VxFlex OS also enables deployments that mix the two-layer and HCI deployments.

VxFlex OS 3.0 comes with lot of new features around space efficiency, data management, and data integrity. These features make it ideal to host database workloads like SQL on VxFlex integrated rack. Few of the newly added features are listed here.

- **Fine Granularity Layout**

VxFlex adds a new storage layout option that is called as Fine Granularity (FG) in addition to the existing Medium Granularity (MG). MG has a 1 MB allocation units and FG has an allocation unit of 4 KB and a physical data placement scheme based on Log Structure Array (LSA) architecture. This new FG requires both Flash media (SSD or NVMe) and NVDIMM in order to create a storage pool. FG gives inline compression that can reduce the total amount of physical data that must be written to SSD. This greatly helps in efficient data storage management and snapshot capabilities. NVDIMM can be used for persistent memory technology which is cacheable and has lesser latencies and throughput than flash storage.

SQL Server with NVMe results in high performance and low latency which translates into more processing of data and analysis. SQL Server thus benefits in providing higher workload consolidation and achieving greater cost savings.

Intel Optane cards also can be used to store SQL Server TempDB files or transaction log files that give excellent random read and write performance at low queue depths. They are particularly used for heavy OLTP workloads.

- **Inline compression**

Fine Granularity layout enables an inline compression capability that can reduce the data footprint.

- **Persistent Checksum**

FG also has persistent checksum that maintains the data integrity for the data and metadata of FG storage pools.

For complete list of new features of VxFlex OS, see [VxFlex OS v3.0 Release Notes](#).

VxFlex OS management is available using a GUI, CLI, and REST clients. There is a VMware vSphere® plug-in that allows VMware admins to deploy, upgrade, configure, and manage VxFlex OS in an ESXi environment within VMware vSphere.

2.1.5 VxFlex Manager

VxFlex Manager is a unified management and automation solution for the VxFlex integrated rack and the VxFlex appliance that enables template-based provisioning of infrastructure and workloads. It provides monitoring, alerting, and health status of hardware and services. It simplifies and automates lifecycle management of diverse and heterogeneous IT resources. It enables users to respond rapidly and reliably to dynamic business needs.

VxFlex Manager is a comprehensive IT Operations Management (ITOM) software purpose that is built for VxFlex to automate and simplify implementation, expansion, and lifecycle management.

VxFlex Manager brings together multiple management consoles, workflow automation, and an intuitive interface that allows customers to monitor, manage, deploy, and maintain physical and virtual resources with the click of a button. Key tenets of the VxFlex Manager architecture include:

- System assurance: compliancy and non-disruptive remediation
- Insights: monitoring, alerting, and health checks
- Implementation simplification: simplified and automated system deployment and workflows
- Node serviceability: single-button take node out of and back into service
- Hypervisor update: single-button upgrade of hypervisor

2.2 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 table shows key features that are new or enhanced in SQL Server 2019:

Table 1 SQL Server 2019 features

Feature	Description
SQL Server Big Data Clusters	SQL Server Big Data Clusters allow you to deploy scalable clusters of SQL Server, Spark, and HDFS containers running on Kubernetes.
Intelligence over any data	SQL Server is a hub for data integration. Deliver transformational insights over structured and unstructured data with the power of SQL Server and Spark.
Choice of language and platform	Build modern applications with innovative features using your choice of language and platform on Windows, Linux, and containers.
Industry-leading performance	Take advantage of breakthrough scalability, performance, and availability for mission-critical, intelligent applications, data warehouses, and data lakes.
Advanced security features	Protect data at rest and in use. SQL Server has been the least vulnerable database over the last 8 years in the NIST vulnerabilities database.
Make faster, better decisions	Power BI Report Server gives your users access to rich, interactive Power BI reports, and the enterprise reporting capabilities of SQL Server Reporting Services.
Better High Availability (HA)	SQL Server 2019 has made significant improvements to availability groups in this version.

For complete information about new features and enhancements, see [SQL Server 2019 Documentation](#).

3 Microsoft SQL Server 2019 on VxFlex - Solution architecture

For this paper, the SQL Server 2019 solution is deployed on Dell EMC VxFlex integrated rack. The VxFlex integrated rack two-layer system is running ESXi hypervisor for compute and network, and VxFlex OS for software-defined storage.

3.1 Logical architecture

The below diagram shows logical architecture diagram of SQL Server 2019 deployment on Dell EMC VxFlex integrated rack two-layer setup that had four SDS and three SDC. The setup of VxFlex two-layer is done using VxFlex Manager which automates the entire deployment process.

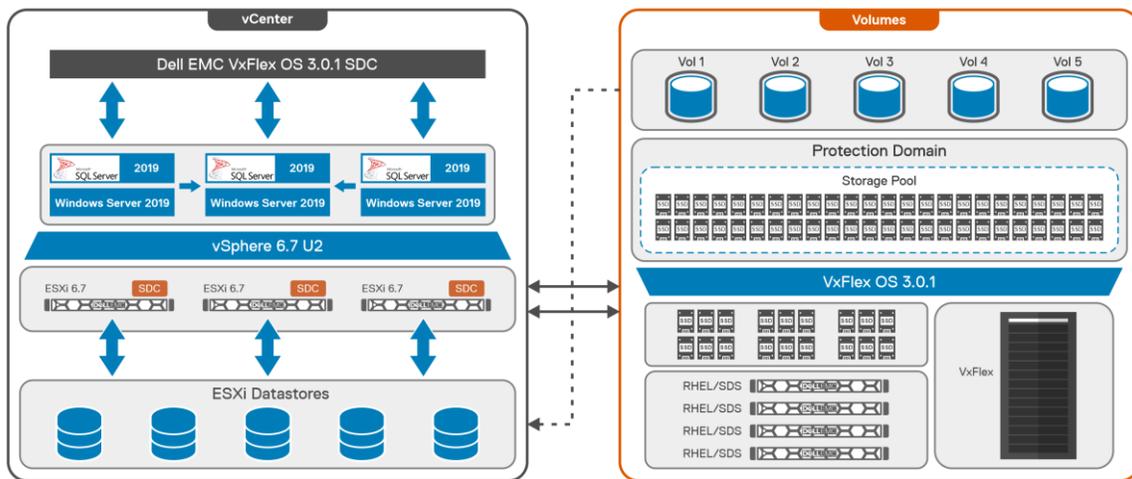


Figure 3 VxFlex integrated rack logical architecture

In the lab environment, for the compute side, one VM on each host (16 vCPU and 64 GB RAM) was created. For the storage side, a protection domain using the four SDS hosts was created. The solution architecture of the VxFlex integrated rack system for a protection domain that had four nodes used for testing along with key components as shown in the diagram.

3.2 Networking layout

The following figure demonstrates high-level network architecture of the two-layer set up on VxFlex integrated rack system:

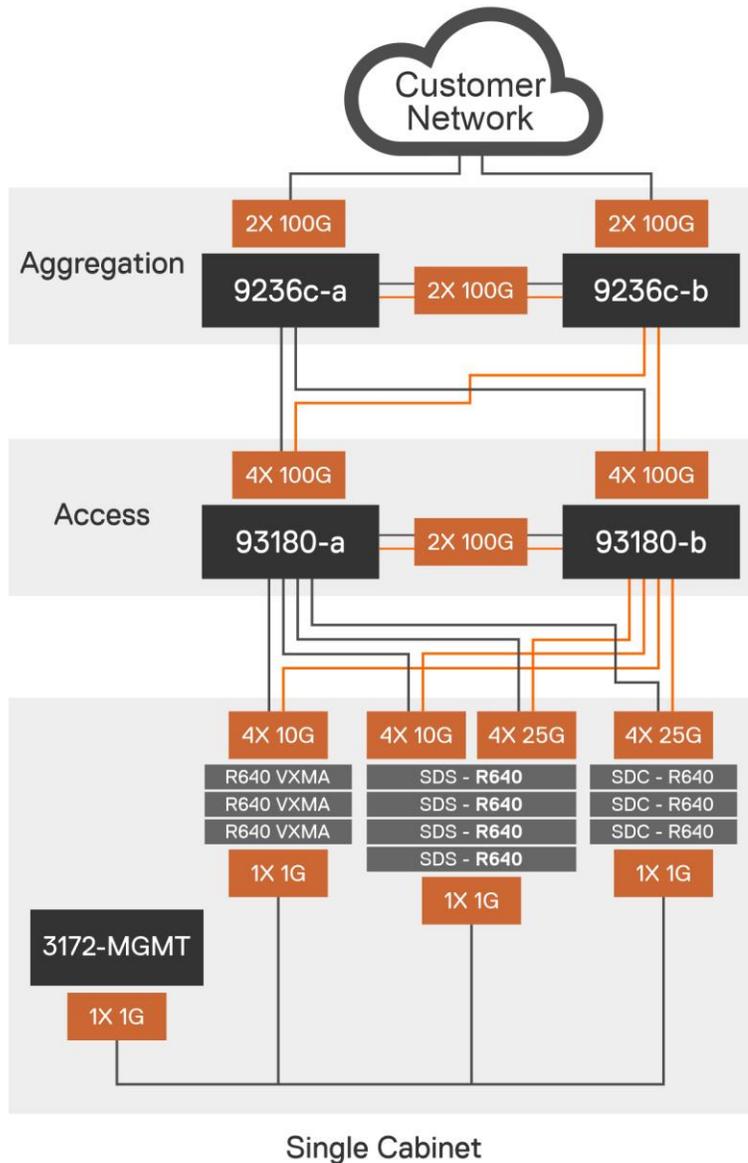


Figure 4 Network architecture

Table 2 VxFlex integrated rack networking details

Components	Description
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
VxFlex storage traffic	2 x 25 Gbps links

3.3 Storage layout

From VxFlex OS standpoint, there was a single VxFlex OS cluster with a protection domain, consisting of four RHEL storage only nodes that play the role of SDS. Each host had ten disks and was fully populated with 1.92 TB SSD Toshiba disks. From the available 40 disks within the storage pool, different volumes were carved out.

The VxFlex compute nodes had SDCs installed on ESXi hosts. From the volumes in VxFlex OS, datastores were assigned to ESXi 6.7. One SQL VM was deployed on each of the datastores, thus there were three SQL VMs deployed on to the VxFlex cluster. All these three VMs could communicate with each other.

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

3.4 SQL database layout

The SQL Server 2019 was running on Windows Server 2019 guest VMs.

The SQL Server virtual machine consists 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 disk 1
102	E	Database disk 2
104	F	Database disk 3
106	G	Database disk 4
150	H	DB Log 1
150	L	DB Log 2
120	T	Temp DB Log

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 Best practices

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

4.1 VxFlex integrated rack

- Minimum number of nodes for 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 VxFlex nodes.
- Ensure VxFlex integrated rack is compliant to an RCM.

4.2 VxFlex OS

- Configure high-performance profile for MDM, SDS, and SDC.
- Disable Read Flash cache and Read RAM cache for all flash clusters.
- Check with VxFlex 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.

4.3 VxFlex integrated rack network

- Confirm with VxFlex integrated rack platform team to enable Jumbo frames for Windows VM, SVM, and at ESXi host.
- 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.
- Separate management and control traffic from production application traffic. You can provide this separation by using VLANs.
- Separate VMware vSphere vMotion traffic from production traffic as per VxFlex standard. For more information, see [Network architecture](#).

4.4 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.

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.

- Put TempDB in NVMe persistent storage volumes.
- If NVMe storage is available, put TempDB into it.
- Create drives of slightly different sizes.

SQL Server datafiles

- Split each database into multiple files: one file per vCPU.
- Size database files equally.
- 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.

Temp database

- Use multiple TempDB files, all the same size.
- Do not autogrow TempDB files.
- If cores < 8, the number of TempDB files = cores.
- If cores > 8, start with eight TempDB files and monitor for performance.
- Do not drop the TempDB database and size TempDB appropriately.
- One TempDB drive should be enough for most environments.

RAM

- More RAM can increase SQL database read performance.
- From the total Windows OS memory, reserve 4 GB to the OS 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.

5 Conclusion

This paper demonstrates how you can deploy SQL Server 2019 on VxFlex integrated rack to meet performance, resiliency, and scale. In addition, it states the best practices for deployment of VxFlex integrated rack, and SQL Server 2019.

Very common workloads like SQL are best suited for VxFlex integrated rack to achieve optimal and fast performance.

Table 4 Benefits of SQL Server 2019 VM on VxFlex integrated rack

Feature	Description
FG Storage Pools	<ul style="list-style-type: none"> The user-data size on disk is reduced due to inline data compression. FG Snapshots are more space efficient. Persistent checksum enabled improved data integrity.
Volume Migration	<ul style="list-style-type: none"> Volumes tied with SQL Database can be migrated between storage pools and protection domains. This can be done not disrupting ongoing I/O operations happening. Migration policy (I/O priority) can be set similar to Rebuild/Rebalance balance.
Snapshot Policy Management	<ul style="list-style-type: none"> Volume level snapshots can be taken on schedule. Retention policy for snapshots can be configured. <p>Note: These two policies can be highly useful in SQL disaster recovery solution.</p>
MG Performance improvements	VxFlex 3.0 ESXi HCI MG performance improvement of 20-30% seen because of SVM change from SLES 12.2 to CentOS 7.5.
Modular incremental scale	The VxFlex integrated rack has four nodes to begin with and can easily scaled as needed. The VxFlex integrated rack can scale unprecedentedly up to 1000 nodes with 100s-PB of Storage.
High performance	For a common workload like SQL Server, VxFlex integrated rack storage delivers low latency with very high number of TPMs when tested in a controlled environment. This shows SQL Server performs at its best when deployed on to VxFlex integrated rack.
VxFlex OS GUI	The VxFlex OS enables ease of creating volumes, shows the different statistics of the SDS and SDC during real time making it possible for the users to monitor performance of SQL VMs and take any relevant decisions.
VxFlex Manager	This tool enables users to create and deploy a VxFlex integrated rack cluster along with VxFlex OS deployment very quickly. SQL VMs can then be provisioned, and overall deployment time is drastically reduced.
VxFlex integrated rack configuration options	Since there are different configuration options to set up VxFlex integrated rack, it gives a great flexibility to the user to set up based on his choice and get the performance as per need.

6 Technical support and resources

6.1 Related resources

See the following referenced or recommended resources that are related to this document:

Note: The following links are open to customers although some may require registration for access.

- [SQL Server 2019](#)
- [SQL Server 2019 on VMware](#)
- [VxFlex OS blog](#)
- [Microsoft Diskspd](#)
- [Hammer DB](#)

6.2 Additional resources

Referenced or recommended publications:

[Dell.com](#) is focused on meeting customer needs with proven services and support.

[Dell EMC Technical Resource Center](#) on DellEMC.com provides expertise that helps to ensure customer success on Dell EMC VxFlex integrated rack platforms.

A Appendix

A.1 VxFlex integrated rack cluster details

The following table summarizes the software resources that were used to carry out the performance tests.

Table 5 Software resources

Component	Details
VxFlex integrated rack	<p>Compute Nodes</p> <p>3 x VxFlex Nodes (R640 servers):</p> <ul style="list-style-type: none"> • VxFlex OS version: R3_0.1 • ESXi version: 6.7 U2 • CPU: 2 x Intel(R) Xeon(R) Gold 6140 CPU @ 2.30 GHz, 18 cores • Memory: 384 GB RAM ((12 x 32 GB DIMMs) <p>4 x VxFlex Nodes (R640 servers):</p> <ul style="list-style-type: none"> • VxFlex OS version: R3_0.1 • VMware ESXi version: 6.7 U2 • CPU: 2 x Intel(R) Xeon(R) Gold 6126 CPU @ 2.60 GHz, 12 Cores <p>Memory: 192 GB RAM (12 x 16 GB DIMMs)</p>
Network	2 NIC cards, each having 2 ports 25 GbE connection
SQL Database VM	<ul style="list-style-type: none"> • Microsoft Windows Server 2019 • Microsoft SQL Server 2019 • 64 GB RAM • 16 vCPUs