

# Dell EMC Ready Solutions for Microsoft SQL: Dell EMC Unity 650F All Flash Storage

December 2018

H17444

## Design Guide

### Abstract

This design guide describes the architecture, design, configuration best practices, and sizing guidelines for a SQL Server 2017 solution with Windows Server 2016, VMware vSphere 6.7 on PowerEdge R740 servers, and Unity 650F All Flash storage.

Dell EMC Solutions

## Copyright

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 © 2018 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. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 12/18 Design Guide H17444.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



# Contents

<b>Chapter 1</b>	<b>Overview</b>	<b>4</b>
	Introduction.....	5
	About this guide.....	6
	We value your feedback.....	6
<b>Chapter 2</b>	<b>Solution Architecture</b>	<b>7</b>
	Overview.....	8
	Architecture.....	8
	Key components.....	10
<b>Chapter 3</b>	<b>Solution Design and Configuration</b>	<b>15</b>
	Solution design.....	16
	Server configuration.....	17
	Network configuration.....	17
	Storage configuration.....	18
	Windows Server, ESXi, and SQL Server configuration.....	20
	Solution sizing.....	23
<b>Chapter 4</b>	<b>Sample Implementation and Use Cases</b>	<b>24</b>
	Introduction.....	25
	Architecture design.....	25
	Hardware components.....	26
	Infrastructure configuration.....	26
	OLTP use case.....	29
	Decision-support use case.....	37
	Summary.....	40
<b>Chapter 5</b>	<b>References</b>	<b>42</b>
	Dell EMC documentation.....	43
	Microsoft documentation.....	43

# Chapter 1 Overview

This chapter presents the following topics:

<b>Introduction</b> .....	<b>5</b>
<b>About this guide</b> .....	<b>6</b>
<b>We value your feedback</b> .....	<b>6</b>

## Introduction

### Microsoft SQL Server

Microsoft SQL Server has developed into a market-leading data management platform in the 25 years since the first release of version 4.2 for Windows NT in 1993. In a [2018 survey of professional software developers](#), Stack Overflow found that SQL Server was the second most popular database platform among professional developers and only 6 percentage points behind the leader MySQL. The early versions of SQL Server, however, were slow to gain widespread adoption for mission-critical applications that needed to be highly scalable and achieve reliability of at least 99.99 percent. By 2000, however, any lingering doubts regarding full Atomicity, Consistency, Isolation, and Durability (ACID) compliance and disaster recoverability were disappearing. By 2005, the analyst community was widely reporting on the ability of SQL Server to scale to levels comparable to or higher than the more established offerings from IBM and Oracle.

The SQL Server 2017 database platform includes a broad range of technologies, features, and services, supporting mission-critical applications such as analytics, in-memory databases, business intelligence (BI), and reliable and scalable online transaction processing (OLTP). The SQL Server platform has acquired capabilities to handle data integration, data warehousing, reporting, high-speed advanced analytics, data replication and programmability features including hosting in-database common language runtimes, service broker hosting, and semi-structured datatype management. Microsoft realizes that with such a breadth of services available in the platform, not all customers or applications need all features enabled on every SQL Server instance. Therefore, Microsoft packages different combinations of features in each of the four SQL 2017 editions—Enterprise, Standard, Express, and Developer.

### Dell EMC Ready Solutions for Microsoft SQL

Because of the vast capabilities of the SQL Server platform, availability of multiple versions, and popularity among developers, customers' varied needs require tailored solutions. Dell EMC Ready Solutions for Microsoft SQL are designed and validated for superior performance, significant cost savings, and future-ready scalability. Our solutions offer customers choices of server form factors, server virtualization versus bare metal, hypervisors including Microsoft Hyper-V and VMware ESXi, and storage flexibility including hyper-converged, storage array, and NVDIMM for high transaction throughput applications.

This Ready Solution is a pre-architected, validated, end-to-end solution for implementing a virtualized infrastructure for SQL Server 2017. Based on Microsoft, VMware, and Dell EMC best practices, the solution provides exceptional performance and scalability, delivers faster time to value, and increases return on investment.

Dell EMC's enterprise products provide high performance, energy efficiency, and high availability. Fourteenth-generation Dell EMC PowerEdge servers, along with Dell EMC storage and networking products, help customers virtualize, consolidate, and migrate their distributed legacy environments. Dell EMC enterprise storage and servers enable the solution to deliver high performance with low latency while maintaining high availability for its components.

## About this guide

### Audience and purpose

This guide describes the architecture, design principles, configuration best practices, and sizing considerations for the Dell EMC Ready Solution for Microsoft SQL Server 2017. It is intended for database administrators, system engineers, IT managers, system administrators, storage administrators, and architects who design and maintain database infrastructures. Readers should have some knowledge of Microsoft Windows Server, Microsoft SQL Server, VMware virtualization, Dell EMC PowerEdge servers, Dell EMC storage, and Dell EMC networking products.

### Scope

This guide includes all the setup and configuration details for implementing this two-tier SQL Server solution, including important best practices to enable you to make decisions for tailoring a configuration for your needs. We demonstrate the solution performance by simulating workload against six virtual machines (VMs) with a standard OLTP workload. This guide also provides information about SQL Server data protection using a Data Domain DD6300 appliance for a backup target. The solution uses Dell EMC Data Domain Boost for Enterprise Applications (DDBEA) to perform SQL Server integrated data protection. DDBEA has an enterprise application data-protection framework that uses Data Domain Boost (DD Boost) to integrate the SQL Server native backup application functionality with the Data Domain backup appliance.

## We value your feedback

Dell EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the [Dell EMC Solutions team](#) with your comments.

**Authors:** Sanjeev Ranjan, Vaani Kaur, Phil Hummel, Mahesh Anand Reddy, Sam Lucido

The following Dell EMC Communities web page provides links to additional documentation for this solution: [Microsoft SQL Info Hub for Ready Solutions](#).

## Chapter 2 Solution Architecture

This chapter presents the following topics:

<b>Overview</b> .....	<b>8</b>
<b>Architecture</b> .....	<b>8</b>
<b>Key components</b> .....	<b>10</b>

## Overview

This Ready Solution uses a two-tier design with Dell EMC PowerEdge R740 rackmount servers for compute and the all-flash Dell EMC Unity 650F model for the storage tier. Dell EMC Connectrix DS 6610B Fibre Channel (FC) switches manage the communication between the compute and storage tiers. Dell EMC Networking S-Series S4048-ON ultra-low-latency 10/40 GbE top-of-rack (ToR) switches provide Ethernet connectivity. All SQL Server 2017 workloads are virtualized on the compute nodes with VMware ESXi 6.7 software. The solution's modular and expandable storage system provides advanced storage management, and the network architecture provides end-to-end I/O connectivity.

## Architecture

### Server architecture

In a database environment, reducing cost while achieving high performance is important. Infrastructure equipment is inarguably the biggest cost component of a data center, apart from the power and cooling that are required to run the infrastructure. This solution uses server virtualization to consolidate different SQL Server roles and reduce equipment cost and footprint. The solution also balances database performance and resource utilization by optimally assigning the processors, memory, and network bandwidth.

### Network architecture

This Ready Solution for SQL Server 2017 provides end-to-end I/O connectivity by using multiple port channels and Virtual Link Trunking (VLT) configurations for LANs. Dell EMC Networking S4048-ON switches provide 10 GbE network connectivity between the compute cluster and the rest of the data center.

In accordance with SQL Server 2017 best practices and Dell EMC infrastructure design principles, each application network is deployed as a separate workload VLAN that is defined in the data-center core network. All the workload VLANs are created as virtual network adapters on the converged virtual switch across four network connections. To provide high availability for network connections, multiple network cards within each server connect to redundant network switches.

Emulex LightPulse LPE31002 FC adapters provide FC connectivity in the host operating system, and the VMs that require in-guest FC connectivity use virtual FC adapters. The Emulex LightPulse host bus adapters (HBAs) are connected to the Connectrix DS-6610B FC switches and then to the Unity 650 All Flash array, forming a highly available, dual-fabric storage area network (SAN) architecture.

The following figure illustrates the network architecture.



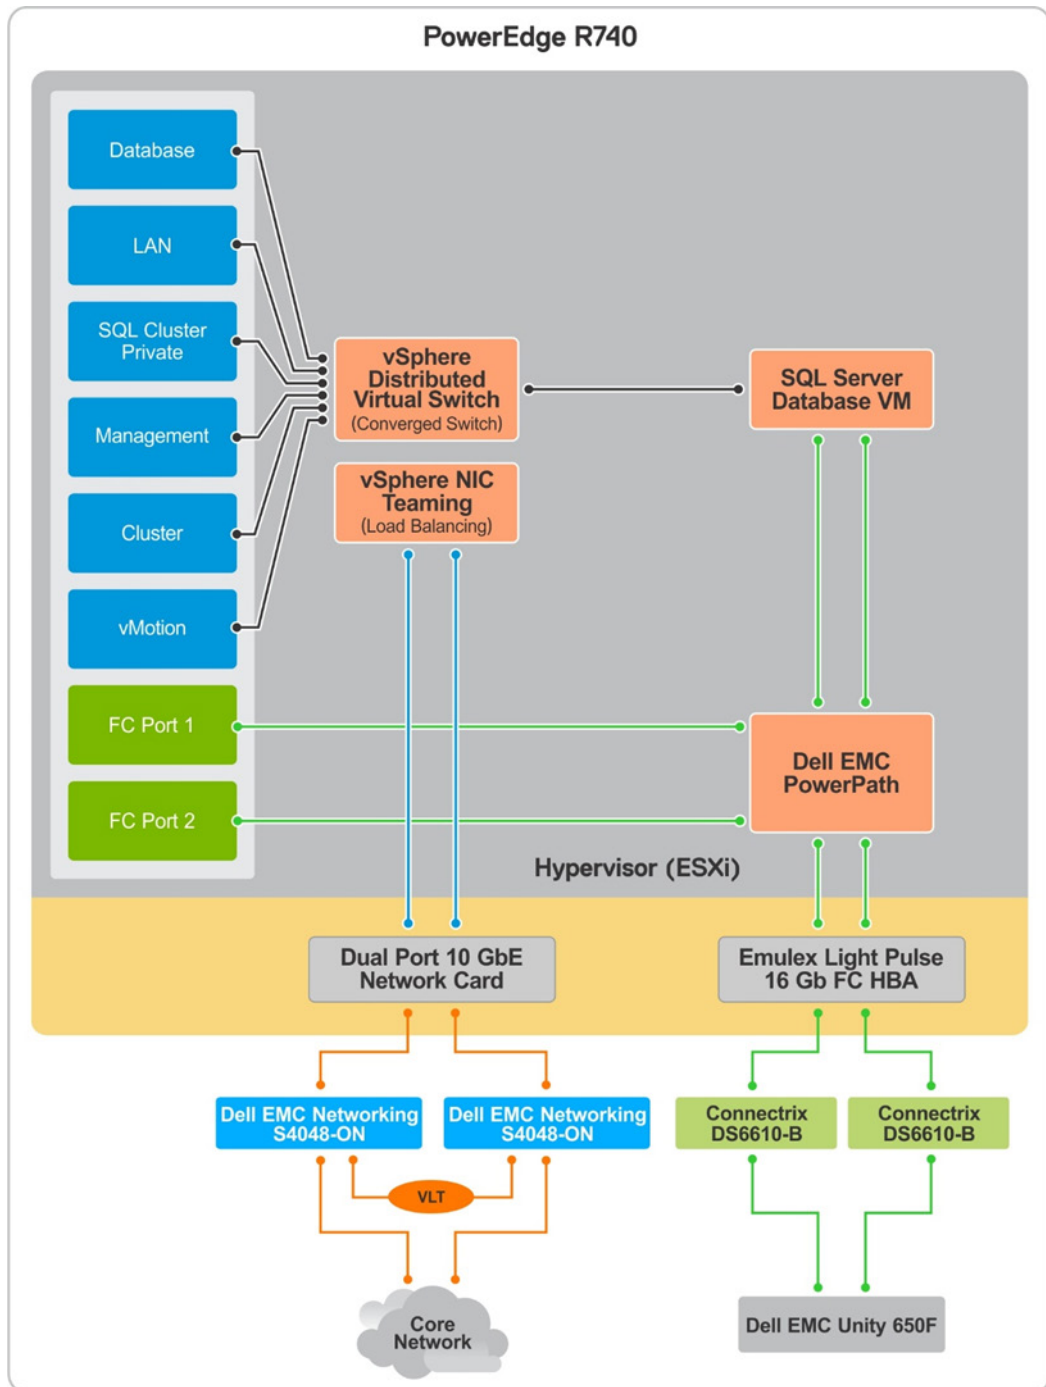


Figure 1. Network architecture

## Storage architecture

Unity 650F storage offers modular expansion to approximately 16 PB of raw SAN capacity with read and write flash solid-state drives (SSDs). By using Dell EMC Unisphere storage management, you can allocate separate volumes for SQL Server data and transaction logs. The Unity 650F array uses high-density SSDs to provide consistent performance and low response times. Further balancing cost and performance are features such as inline compression, deduplication and zero detection.

## Key components

This Ready Solution includes the following key components:

- Dell EMC PowerEdge R740 server
- Dell EMC Unity 650F storage array
- Dell EMC Networking S4048-ON ToR switches
- Connectrix DS-6610B FC switches
- VMware vSphere 6.7 virtualization platform
- Microsoft Windows Server 2016
- Microsoft SQL Server 2017
- Dell EMC Data Domain DD6300 backup appliance
- Dell EMC Data Domain Boost for Enterprise Applications

### Dell EMC PowerEdge R740

The PowerEdge R740 2U rack server is designed to run complex workloads in midsize and large enterprises. The 2-socket R740 server:

- Features the Intel Xeon Scalable processor family with up to 28 cores
- Uses highly scalable memory, I/O, and network options
- Provides up to 80 TB of internal storage with sixteen 2.5-inch drives or eight 3.5-inch drives—ideal for database and BI applications

The following table lists the technical specifications of the PowerEdge R740.

**Table 1. PowerEdge R740 technical specifications**

Component/feature	Specifications
Form factor	2U
Processor	Intel Xeon Scalable processor family with up to 28 cores per processor
Processor sockets	2
Chipset	Intel C620 series
Memory	Up to 3 TB (24 DIMM slots): 8/16/32/64/128 GB DDR4 with up to 2,666 MT/s
I/O slots	Riser options with up to 8 PCIe Gen 3 slots; maximum of 4 x 16 slots
Storage controllers	<ul style="list-style-type: none"> <li>• Internal controllers: PERC H330, H730p, H740p, HBA330, Software RAID (SWRAID) S140</li> <li>• Boot Optimized Storage Solution (BOSS) subsystem: Hardware RAID 2 x M.2 SSDs—120 GB, 240 GB</li> <li>• External PERC (RAID): H840</li> <li>• External HBA (non-RAID): 12 Gb/s SAS</li> </ul>
Network interface	4 x 1 GB, 4 x 10 GB, 2 x 10 GB + 2 x 1 GB, or 2 x 25 GB Network Daughter Cards

## Dell EMC Unity 650F storage

Dell EMC Unity is a modern, powerful, and flexible storage array that is ideal for deploying mission-critical, latency-sensitive applications such as SQL Server 2017. The Unity 650F array is the fastest Unity All Flash midrange storage system. The array offers a combination of features including high performance, large storage capacity, and low latency in a midrange system for the broadest range of SAN and NAS use cases. The Unity 650F system includes the following key features:

- Unification of block, file, and VMware VVols support
- A maximum of 16 PB of raw storage
- Enhanced virtualization and deduplication
- The capability to maximize all-flash storage with new inline data reduction

The following table lists the technical specifications of the Unity 650 All Flash array.

**Table 2. Unity 650F technical specifications**

Component/characteristic	Specifications
Controllers per array	2 (active/active)
Operating system	Unity OE 4.4.0.1534750794
Maximum drives supported	1,000 per array
Maximum raw capacity*	16.0 PB per array
Storage media	Write-intensive, read-intensive SSDs with various capacities: 400 GB, 800 GB, 1.6 TB, 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB
Array enclosures	Disk processor enclosure (DPE) with two form factors: <ul style="list-style-type: none"> <li>• 2U with 25 x 2.5-inch drives</li> <li>• 2U with 8 x 3.5-inch drives</li> </ul>
Drive enclosures	Support for: <ul style="list-style-type: none"> <li>• 2U 25-drive trays for 2.5-inch drives</li> <li>• 3U 80-drive trays for 2.5-inch drives</li> </ul>
RAID options	1/0, 5, 6
Embedded SAS I/O ports per array	4 x 4-lane 12 Gb/s SAS ports for back-end connectivity
Embedded CNA ports per array	4 ports: 8/16 Gb FC, 10 Gb IP/iSCSI, or 1 Gb RJ45
Embedded 10GBASE-T ports per array	4
Maximum FC ports per array	20
Maximum SAN hosts	2,048
Maximum LUN size	256 TB

\* Maximum raw capacity varies based on available drive sizes at time of purchase.

Unity arrays are part of Dell EMC's Future-Proof Loyalty Program, which provides:

- Guaranteed satisfaction for 3 years
- Investment protection through the trade-in of existing or competitive systems for credit toward new Unity storage

For more information, see [The Future-Proof Loyalty Program](#) on the Dell EMC website.

### **Dell EMC Networking S4048-ON**

The Dell EMC Networking S-Series S4048-ON is an ultra-low-latency 10/40 GbE ToR switch that is built for applications in the high-performance data center and computing environment. The compact S4048-ON high-density switch features 48 dual-speed 1/10 GbE (SFP+) ports as well as six 40 GbE QSFP+ uplinks. The S4048-ON maximizes network performance with a nonblocking switching architecture that provides line-rate L2 and L3 forwarding capacity with ultra-low latency.

### **Connectrix DS-6610B FC switch**

The Connectrix DS-6610B switch delivers up to 32 Gb/s FC performance, with additional optic options available for 4, 8, or 16 Gb/s. The DS-6610B is a 1U switch that supports up to 24 ports through a ports-on-demand (PoD) model. The PoD kits enable expansion in 8-port increments.

Management tasks can be performed using the Connectrix Manager Converged Network Edition (CMCNE). Optional software license key features include Fabric Vision for proactive monitoring and management, Extended Fabric for FC connectivity greater than 10 km, and Inter-Switch Link (ISL) Trunking for aggregation of multiple physical links.

### **VMware vSphere 6.7**

VMWare vSphere 6.7 is an efficient and secure platform for virtualization and cloud. It includes components such as VMware vCenter Server and the VMware ESXi hypervisor:

- vCenter Server is a unified platform for managing VMWare vSphere environments. vCenter Server features ease of deployment, extensibility onto the public cloud, and proactive optimization.
- ESXi is a bare-metal hypervisor that is installed directly onto physical servers. ESXi is more efficient than hosted architectures because it has direct access to and control of underlying resources. The ESXi hypervisor has a small footprint of 150 MB, minimizing space requirements and security threats. It provides reliable performance for different application workloads by configuring VMs with up to 128 vCPUs, 6 TB of RAM, and 120 devices.

### **Microsoft Windows Server 2016**

Windows Server 2016 is a cloud-ready operating system that provides enhanced security, built-in containers, and support for new software-defined capabilities for modern data centers. Key features include:

- Resilient File System (ReFS), which enables faster placement of VMs on the file system
- Software-defined networking that includes enhanced policies to control both physical and virtual networks

For a complete list of new features of Windows Server 2016, see the Microsoft article [What's new in Windows Server 2016](#).

## Microsoft SQL Server 2017

SQL Server 2017 delivers high-performance OLTP, new encryption features, greater support for in-memory databases, and a new end-to-end BI solution. Key features include:

- **SQL Server Machine Learning Services**—Supports Python and R languages
- **Support for multiple platforms**—Runs on Linux, Linux-based Docker containers, and Windows
- **Resumable online index rebuild**—Resumes an online index rebuild operation from where it stopped after a failure, or pauses and later resumes an online index rebuild operation
- **Automatic database tuning**—Provides insight into potential query performance problems and recommends solutions, and can automatically fix identified problems

For more details about Microsoft SQL Server 2017, see the Microsoft article [What's New in SQL Server 2017](#).

## Dell EMC Data Domain DD6300 appliance

Data Domain DD6300 is a backup appliance that is designed for the modern software-defined data center. Offering simplicity, efficiency, and high reliability, the DD6300 can natively tier deduplicated data to any cloud.

The DD6300 appliance includes comprehensive data protection software that enables organizations of all sizes to protect, archive, and recover mission-critical workloads, regardless of where the data resides. Dell EMC data protection software protects a broad range of applications ranging from VMs to databases (high I/O OLAP and low-latency OLTP), and VMware workloads running on AWS. In some cases, Data Domain solutions accelerate the backups up to 20 times and the recovery up to 10 times for mission-critical applications.

The following table lists the technical specifications of the DD6300 appliance.

**Table 3. DD6300 technical specifications**

Feature		DD6300 All-in-One base configuration	DD6300 All-in-One expanded configuration
Rack height		2U	2U
Processor		E5-2620 V3	E5-2620 V3
Memory configuration (nonextended retention)		48 GB	96 GB
DIMMs		6 x 8 GB	12 x 8 GB
Supported capacity (nonextended retention)		76 TB (28 TB internal + 48 TB external)	180 TB (36 TB internal + 144 TB external)
HDDs in 3.5-inch bays		7 or 7+5	12
SSDs in 2.5-inch bays		1	2
NVRAM		NVRAM 8g Model 3	NVRAM 8g Model 3
SAS I/O modules (quad-port 6 Gb/s SAS)		> 0 for internal storage only; > 1 with external storage	> 0 for internal storage only; > 1 with external storage
SAS string depth (max)	ES30	1	4
	DS60	0	1

## Dell EMC Data Domain Boost for Enterprise Applications

Dell EMC Data Domain Boost for Enterprise Applications (DDBEA) provides application-centric data protection. DDBEA uses an enterprise application data-protection framework that uses Data Domain Boost (DD Boost) to integrate the application with the Data Domain backup appliance.

The Microsoft application agent for DDBEA enables database and application administrators to efficiently back up and restore applications by using the applications' native tools.

Features and capabilities of the Microsoft application agent for SQL Server include:

- Instance-level backups
- Multiple database backups and restores
- Transact-SQL (T-SQL) scripting to perform backups and restores
- Scheduled backups by using SQL Server Agent jobs
- For transaction log backups, a choice of either skipping the simple model databases from the backups or promoting the simple model databases to the full backups
- Full backup copies of the Always On availability group databases
- Support for IPv6

For more information about DDBEA for Microsoft applications, see the [Dell EMC Microsoft Application Agent Installation and Administration Guide](#).

## Chapter 3 Solution Design and Configuration

This chapter presents the following topics:

<b>Solution design</b> .....	<b>16</b>
<b>Server configuration</b> .....	<b>17</b>
<b>Network configuration</b> .....	<b>17</b>
<b>Storage configuration</b> .....	<b>18</b>
<b>Windows Server, ESXi, and SQL Server configuration</b> .....	<b>20</b>
<b>Solution sizing</b> .....	<b>23</b>

## Solution design

Because databases are mission-critical for organizations, high availability and service-level agreements are essential elements of OLTP database design. Today’s customers also demand faster response time and better optimization of hardware resource utilization.

This Ready Solution provides high availability through vSphere High Availability (vSphere HA). The design uses the cluster to manage host resources and provide high availability for VMs. RAID for storage drives, dual controllers, and multiple I/O paths provide storage high availability. Multiple FC HBAs and FC switches between the SQL Server infrastructure and the storage network build resiliency in network connectivity.

The following figure illustrates the solution design.

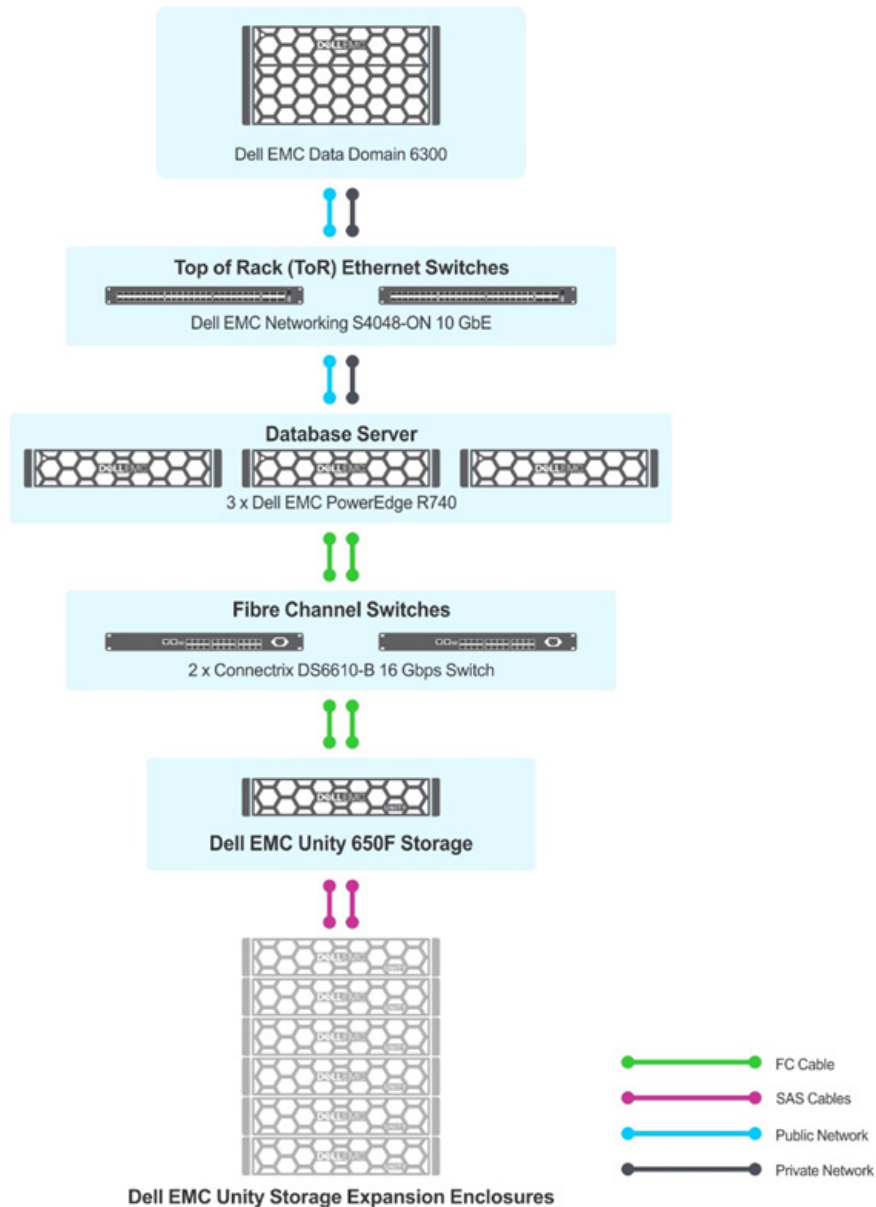


Figure 2. Solution design



## Server configuration

Follow these Dell EMC-recommended practices for configuring the servers:

- Use the latest BIOS and firmware for the PowerEdge servers.
- Use the latest vendor-subscribed HBA drivers.
- Enable Intel Hyper-Threading, which enables higher utilization of the available physical cores to deliver the best SQL Server performance from the compute resources.
- Configure the power plan as follows:
  - In the system firmware:
    - i Go to **System Setup > System BIOS > System Profile Settings**, and select **Performance Per Watt Optimized (DAPC)** from the list.
    - ii Go to **System Setup > System BIOS > Processor Settings > Dell Controlled Turbo**, and select **Enabled**.
  - In the operating system: Go to **Control Panel > Power Plans Scheme**, and select **High Performance Plan**.

## Network configuration

This section outlines the best practices for configuring the network in the Ready Solution for SQL Server.

### Local area network

Best practices for configuring the LAN include:

- Use two 10 GbE S4048-ON network switches for end-to-end client traffic to prevent a single point of failure.
- Use VLAN tagging to sustain switch failure, and use NIC teaming to provide multitier network high availability.
- Use NIC teaming to provide high availability at the NIC port level.
- Use NIC teaming for virtual and physical network connections.
- Use a minimum of two network cards to provide high availability and prevent a single point of failure.
- Use two Intel Ethernet 10 GbE dual-port X710 adapter cards to enable connectivity between the SQL Server host and applications on the external network.

### Storage area network

Best practices for designing the SAN include:

- Use redundant Connectrix 6610 switches to prevent a single point of failure.
- Use two Emulex LightPulse LPE31002 cards per server for high availability.
- Use named zoning of end devices (FC HBAs) and Dell EMC Storage Center virtual ports while creating zones. Named zoning provides better flexibility because it does not bind the devices and virtual ports to specific physical ports on either the switch or the storage array.

- Make crisscross cable connections between redundant FC HBA cards and redundant FC switches to enable high availability, as shown in the following figure.




Figure 3. Connections between PowerEdge R740 and Connectrix switches

## Storage configuration

Dell EMC Unisphere storage management is a web-based solution that provides an interface for storage management and configuration tasks. The configuration depends on performance and capacity needs, preferred type of server connectivity, performance in terms of IOPS and MB/s, and future growth needs for both performance and capacity. This section provides guidance and best practices for storage management and configuration.

### Storage pools

Dell EMC recommends that you use fewer storage pools within Unity storage to reduce complexity and increase flexibility. However, you might want to configure multiple storage pools to do the following:

- Separate workloads with different I/O profiles.
- Dedicate resources to meet specific performance goals.
- Separate resources for multitenancy.
- Create smaller failure domains.

Storage pools must maintain free capacity to operate properly. By default, Unity storage raises an alert if a storage pool has less than 30 percent free capacity. It also automatically begins to invalidate snapshots and replication sessions if the storage pool has less than 5 percent free capacity. Dell EMC recommends that storage pools always have at least 10 percent free capacity.

### Drive selection

Dell EMC storage systems use multiple drive types, with varying capacities and performance levels, and attempt to match data with the optimal drive type based on storage pools. Typically, mission-critical data is stored on high-performance drives, while less-critical data is moved to high-capacity drives.

While every environment is unique, an OLTP workload generally consists of small, random reads and writes. The required number of IOPS primarily determines the sizing of a storage system that services an OLTP workload. Different drive types have different performance capabilities. Faster-spinning drives can provide more IOPS with lower

latency than slower-spinning drives. For best performance, Dell EMC recommends SSDs for OLTP workloads.

The following factors affect drive selection:

- Type of workload
- Expected IOPS, latency, and throughput
- Database size

You can calculate the number of required disks as follows, depending on the target IOPS requirements:

Number of disks = (Read IOPS + [Write IOPS x RAID performance overhead]) / IOPS per disk

### I/O path settings

FC zoning enables the partitioning of the FC fabric into multiple subsets called zones. A zone includes FC initiators and targets, and enables communication between the initiators and targets.

Best practices for setting the I/O paths include:

- **Single-initiator, multiple-target zoning**—Creates each FC zone with a single initiator (HBA port) and multiple targets (front-end ports). Therefore, each HBA port requires a separate FC zone with front-end ports. Create independent zones for each HBA that is installed in the host.
- **World Wide Name (WWN) zoning**—Contains only the host HBA port and the front-end primary ports. In most cases, including the front-end reserve ports is unnecessary because they are not used for volume mapping. Dell EMC recommends creating zones by using a single initiator host port and multiple Storage Center ports.
- **Multipathing**—Designates how many of the Dell EMC storage front-end ports that the system allows the volume to be mapped through. Dell EMC recommends the use of the Dell EMC PowerPath/VE Multipathing utility for designating the ports.

### Volume provisioning and mapping

Create a server cluster for high availability and map the volumes to the server cluster. Map the volumes by using Storage Center Console or Enterprise Manager Console.

For optimal performance and better monitoring resolution, Dell EMC recommends that you:

- Place tempdb, data, and log files on separate LUNs.
- Leave room for data growth and avoid exceeding 80 percent capacity of the LUNs for database files.
- Configure end-to-end multipathing by using PowerPath/VE software for availability/redundancy and throughput optimization.
- Use multiples of 64 KB as allocation units for SQL Server volumes.
- Ensure that I/O is distributed across both controllers of the storage array. For best results, evenly distribute the heavy-I/O LUNs and try to use the same number of LUNs or volumes on each controller.

## ESXi storage design

For ESXi storage design, Dell EMC recommends that you:

- Create a virtual SAN by grouping WWNs of virtual FC adapters that are hosted in PowerEdge servers and Dell EMC storage. Configure the VMs that are hosting SQL Server with direct access to the storage LUNs.
- Create a single large pool with RAID 10 configured, rather than many smaller pools.
- Use multiple LUNs or volumes for the database data files that have low latency and high IOPS requirements.
- Use PowerPath/VE software for multipathing.

## Windows Server, ESXi, and SQL Server configuration

This section outlines the best practices for configuring Windows Server 2016, ESXi 6.7, and SQL Server 2017.

### Windows Server 2016

Best practices for configuring the Windows operating system include:

- Use an allocation unit size of 64 KB to format the volume that stores the database files.
- Enable the Windows Lock Pages in Memory (LPIM) policy by adding an account with privileges to run `sqlservr.exe`. The LPIM policy determines which accounts can use a process to keep data in physical memory. The policy prevents the Windows operating system from paging out a significant amount of data from physical memory to virtual memory on disk.
- Configure the High Performance power profile.

### ESXi

Best practices for configuring ESXi and deploying VMs include:

- Do not overcommit resources. ESXi allows overcommitting of resources such as CPU, memory, and network; however, overcommitting leads to performance degradation when the resource usage exceeds the available resources.
- Calculate host memory as follows:  
$$\text{HostMem} \geq \text{Sum of VM memory} + \text{overhead}$$
- Set the memory reservation equal to the provisioned memory to eliminate the possibility of ballooning or swapping. When calculating the amount of memory to provision for the VM, use the following formulas:
  - $\text{VM Memory} = \text{SQL maximum server memory} + \text{ThreadStack} + \text{operating system memory} + \text{VM overhead}$
  - $\text{ThreadStack} = \text{SQL maximum worker threads} \times \text{ThreadStackSize}$
  - $\text{ThreadStackSize} = 1 \text{ MB on x86} = 2 \text{ MB on x64 OS Mem: } 1 \text{ GB for every } 4 \text{ CPU cores}$
- Provision compute resources according to your workload requirements. Ensure that the CPU utilization at the VM level is less than 75 percent. A virtual CPU (vCPU) is

a representation of the physical core of a processor or threads (logical processors) in the core.

- Use fixed-size virtual hard disks for the operating system. Dell EMC recommends using fixed virtual hard disks (VHD/VHDX) for production workloads. Using dynamic virtual hard disks can result in occasional pauses during disk resizing. Use dynamic disks for noncritical test environments or nonproduction environments.
- Use the High-Performance power management profile. To change power profiles, see [Host Power Management Policies](#).
- Use separate storage LUNs for data and log files to rule out disk contention. For best performance, create distinct LUNs for SQL Server data files and SQL Server log files. These strategies also help to create different access patterns on the storage system—sequential access for log files and random access for data files. The access patterns enable the Unity 650F array to manage the data most effectively.
- Use Raw Device Mapping (RDM) for LUNs. For more information about the benefits of RDM and how to configure it, see [About Raw Device Mapping](#).
- Use VMware Para-Virtual SCSI (PVSCSI) adapters. For information about how to map LUNs using PVSCSI adapters, see VMware KB 1010398, [Configuring disks to use VMware Paravirtual SCSI \(PVSCSI\) adapters](#).
- Adjust the **Disk.reqCallThreshold** value for your VMs. For instance, for OLTP, we found that the lowest latency occurred at the VM level with the parameter set to 1. For decision support, we achieved the best results with the default value of 8. For more information, see [Virtualizing Performance—Critical Database Applications in VMware vSphere 6.0 Performance Study](#).
- Adjust the VM latency setting value from normal to high.

**SQL Server 2017** This section provides guidelines and best practices to optimize a virtualized SQL Server environment.

### Memory settings

For SQL Server 2017, you can assign memory for a VM either dynamically or statically. To choose between dynamic and static memory, consider the following factors:

- How frequently you monitor your database
- The VM size (size of the processors, memory, and other resources) as compared to the size of a single non-uniform memory access (NUMA) node on the host's physical architecture
- Preference between performance and levels of scalability

In general, choose dynamic memory when VMs are unmonitored and are relatively small, and when scalability is preferred over performance. For better and more consistent performance for larger production VMs that are reasonably monitored, choose static memory.

Best practices for dynamic memory include:

- Determine and implement the VM startup RAM and minimum memory based on your needs. Microsoft recommends leaving **max server memory** at its default setting, which enables SQL Server to manage memory dynamically. However, Dell EMC recommends that you change this value, as needed, if one or both of the following are true:
  - You are running multiple applications on the VM.
  - You can reasonably ascertain the maximum amount of memory that you want to assign to SQL Server.
- Set **min server memory** based on usage and performance considerations because dynamic memory is enabled for the VM. The default value for **min server memory** is 0.
- When using dynamic memory, set the reserves for the VM by using the **Memory Buffer** option. Base the amount of reserved memory on the **min server memory** setting for SQL Server and the memory that is required for any other applications and the operating system.

Recommendations for static memory allocation include:

- To maximize performance, assign a particular amount of memory to a VM based on its virtual and physical NUMA architecture.
- Set **max server memory** and **min server memory** values based on the amount of memory you want to reserve for the operating system, typical requirements of your SQL Server, and other performance considerations. For more information about memory settings, see [Blitz Result: Memory Dangerously Low or Max Memory Too High](#).

### Parallelism settings

For SQL Server parallelism settings, consider the following recommendations.

---

**Caution:** Changing the parallelism settings can have extreme unanticipated consequences. Dell EMC recommends that you do not change these settings without testing and that you are prepared to closely monitor the system in production after making changes.

---

- The SQL Server configuration option **max degree of parallelism** controls the number of processors that are used for the parallel execution of a query. If the SQL Server VM in question migrates from hosts with different core configurations, leave this setting at 0; however, if that is not the case, use another value.
- The SQL Server configuration option **cost threshold for parallelism** specifies a threshold at which query plans run in parallel. You can change the setting for this option from the default value of 5 to 50. You can adjust the setting further, depending on your requirements and testing.

For more information about parallelism settings, see [MAXDOP of Confusion](#).

For more information about Max Worker Threads, see [Configure the max worker threads Server Configuration Option](#).

## Solution sizing

Each SQL Server environment has distinct system requirements and must be sized according to the instance-specific demands and the workload deployed. Workload characteristics in a SQL Server deployment include the nature of transactions, the expected user load, the required number of transactions per second (TPS), and the permissible levels of average query response time. SQL Server is a storage-intensive workload and requires high availability. The significant sizing considerations are server and storage.

For server sizing:

- Determine the type of processor that is best suited for handling the SQL Server workload requirements. The number of cores per processor and the frequency at which the processors operate are significant factors in selecting processors for any database solution.
- Determine the amount of required memory and allocate the DIMMs to the processor memory channels to take advantage of full memory bandwidth. Also, the amount of required memory for the database solution depends on database size, query traffic on the database, and average query response time requirements. The more memory that the SQL Server instance has the more data it can keep in its buffer pool and the less physical I/O it has to perform.
- Select the appropriate host network adapters based on the switches that are part of the solution. Networks can sometimes be a bottleneck in delivering results from the database server to the application client. For better compatibility and performance, select network adapters with a speed that matches the speed of the switches.
- Consider using two FC adapters for high availability and better performance, and ensure that the speed of the FC switches match the speed of the FC adapters.

For storage sizing and disk selection requirements, see [Storage configuration](#).

# Chapter 4 Sample Implementation and Use Cases

This chapter presents the following topics:

<b>Introduction</b> .....	<b>25</b>
<b>Architecture design</b> .....	<b>25</b>
<b>Hardware components</b> .....	<b>26</b>
<b>Infrastructure configuration</b> .....	<b>26</b>
<b>OLTP use case</b> .....	<b>29</b>
<b>Decision-support use case</b> .....	<b>37</b>
<b>Summary</b> .....	<b>40</b>



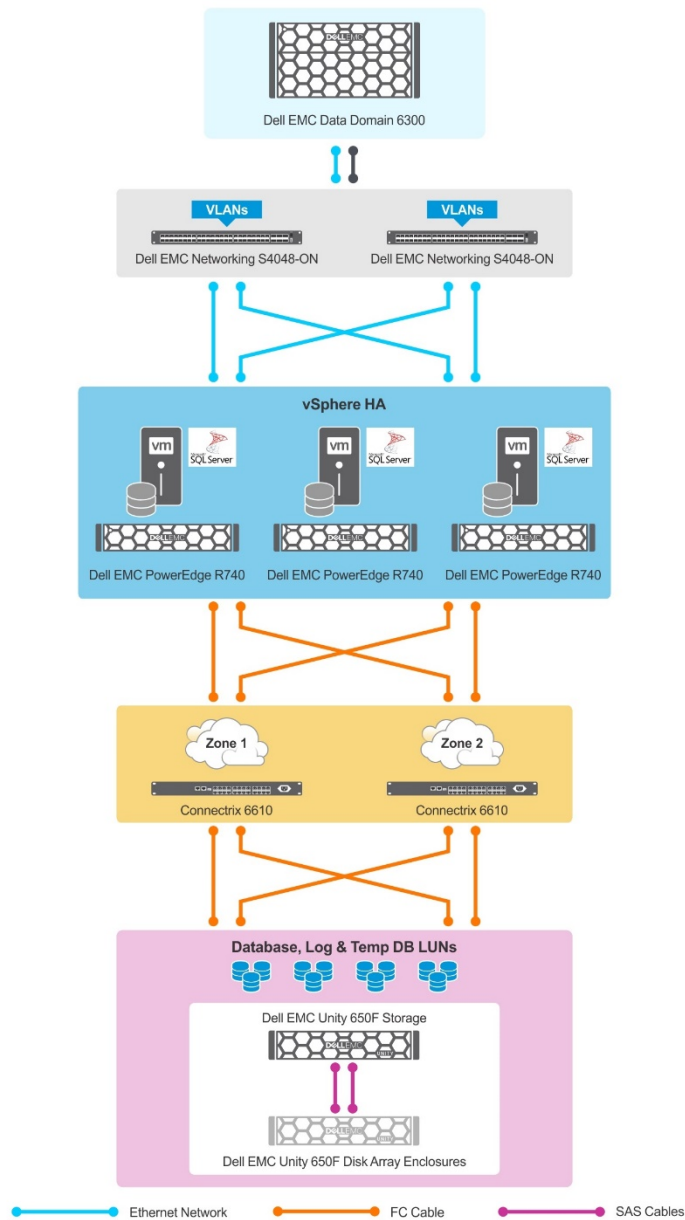
## Introduction

This chapter provides the architecture design and component configurations for a sample implementation of this solution. It includes implementation details and test results for the following use cases on a database size of up to 3.6 TB:

- OLTP workload supporting up to 24,533 application-defined TPS
- Decision-support workload supporting up to 603,146 queries per hour

## Architecture design

The following figure shows the architecture for this sample implementation.



**Figure 4. Solution architecture for sample implementation**

## Hardware components

The following table lists the hardware components for this sample implementation.


**Table 4. Workload solution hardware components**

Hardware components	Specifications
Servers	3 x Dell EMC PowerEdge R740
Network switches	2 x Dell EMC Networking S4048-ON
FC switches	2 x Connectrix DS-6610B
Backup appliance	1 x Dell EMC Data Domain DD6300
External storage	1 x Dell EMC Unity 650 flash array: <ul style="list-style-type: none"> <li>• Storage pools: 1</li> <li>• Drive type: Caching, 4 x 800 GB SAS flash</li> <li>• Extreme Performance tier: 116 x 400 GB SAS flash RAID 1/0</li> </ul>
	6 x Dell EMC Unity storage expansion enclosure
<b>Components per server</b>	
Processor	2 x Intel Xeon 8180 Scalable family; 2.5 GHz, 28 cores
Memory (RAM)	768 GB (24 x 32 GB, 2,666 MT/s DIMM)
rNDC	Broadcom NetXtreme BCM5720 GbE PCIe
Add-on NICs	2 x Intel Ethernet 10 GbE dual-port X710 adapter
FC HBA	2 x 16 Gb FC (GFC) Emulex LightPulse LPE31002

## Infrastructure configuration

### PowerEdge R740 servers and SQL Server VMs

The PowerEdge R740 server is designed to maximize compute density for enterprise applications and databases with a small 2U, 2-socket rack configuration. For this SQL Server with Unity testing, we used three identical PowerEdge R740 servers to provide a highly available VMware vSphere cluster that could sustain the loss of one server with minimal impact to the SQL Server database workloads. Our final testing configuration had six SQL Server VMs, as shown in the following figure, with two VMs running on each server.



**Figure 5. PowerEdge R740 servers**

We enabled vSphere HA so that two VMs could automatically restart on the surviving servers if a single server failed. A VMware administrator using vSphere vMotion can also proactively migrate the VMs on a server when a server requires patching or maintenance. During a vMotion migration, users experience only a temporary and minor impact on database performance.

We configured each VM with 24 vCPUs and 192 GB of memory. VMware recommends that the allocation of vCPUs not exceed physical CPUs for production workloads. In this implementation, we assigned 24 vCPUs per VM, which is less than half the number of 112 logical cores (hyper-threading was enabled) that are available in each server. We were well under the guideline of not exceeding the number of physical cores, thus preventing a possible processor performance bottleneck.

Each VM also had 192 GB of reserved memory. SQL Server and other databases use memory efficiently and are dependent on memory for optimal performance. Using a memory reservation prevents memory contention or pressure that might affect database performance. In this case, we reserved a total of 384 GB of vMem across two VMs. The physical server memory was 768 GB, leaving 384 GB (768 – 384) free memory headroom on the server.

### PowerPath/VE Multipathing

Dell EMC PowerPath/VE Multipathing software integrates with the VMware ESXi hypervisor to intelligently manage data paths between servers and storage. PowerPath/VE software has patented algorithms that provide improved performance under ever-increasing workloads. Because our use-case approach was to incrementally increase the SQL Server workload on the solution infrastructure, we used PowerPath/VE software to optimize path management.


### Unity 650F storage

The Unity 650F storage array is Dell EMC's fastest midrange storage system. Unity All Flash systems can provide the high-throughput, low-latency IOPS that consolidated SQL Server database environments need. Our Unity test configuration used an extreme performance layer of 116 x 400 GB SAS flash drives to deliver high performance and low latency. All-flash storage is ideal for mixed OLTP and decision-support system consolidation. SSDs provide low-latency performance for small-block random reads and writes as well as larger-block table and range scans.

Each VM and SQL Server database used the same storage configuration. For the operating system and VM files, we created a 200 GB VMFS datastore. For each SQL Server database, we created RDM files for the following storage:

- **Data**—700 GB
- **Log**—200 GB
- **tempdb1**—300 GB
- **tempdb2**—200 GB

Adhering to a best practice, we separated files into different LUNs for performance analysis, as shown in the following figure. Having separate LUNs for data, log, and tempdb files enabled us to monitor throughput and latency for each type of I/O profile. Having separate LUNs also provided the capability to report on storage performance at a granular level.



**Figure 6.** Unity 650F all-flash storage array with LUN mapping

Using Quest Benchmark Factory, we configured workloads to simulate a small to medium environment. The Unity array can scale to support large SQL Server workloads, and the DBAs can use multiple data LUNs to reduce the potential for I/O queuing managed by Windows for each device. Although having multiple data LUNs does not increase performance, it does enable I/O distribution across more devices and paths to the storage array so that performance bottlenecks are less likely to develop. Dell EMC has SQL Server specialists and consulting services that can help customers design a Unity storage configuration that is based on workload analysis to optimize manageability and performance.

The Unity 650F array simplifies storage design and enables fast provisioning of storage for SQL Server databases. The all-flash array offers the performance that is needed for small- and medium-size businesses to consolidate their SQL Server database ecosystem. Key Unity features that apply to SQL Server and other databases include:

- Multiprotocol support for block, iSCSI, file (such as CIFS), and VVols.
- Advanced data reduction technology to maximize the value of all-flash.
- Up to 55:1 data protection deduplication when you use recommended best practices. For more information about Dell EMC’s data protection deduplication guarantee, see [The Future-Proof Loyalty Program](#) on the Dell EMC website.

In validating and testing SQL Server 2017 on this Ready Solution with Unity storage, we demonstrate how well the infrastructure can support your databases.

## OLTP use case

### OLTP workload implementation

The following table lists the OLTP workload implementation details.

**Table 5. OLTP workload implementation details**


Characteristic	Value
Number of VMs	6
Number of data file LUNs	6
Number of log file LUNs	6
Number of tempdb data file LUNs	6
Number of tempdb log file LUNs	6
<b>Per VM configuration</b>	
vCPUs	24
Memory	192 GB
Memory for database	160 GB
Database LUN size	700 GB
Log file LUN size	200 GB
tempdb data files LUN size	300 GB
tempdb log file LUN size	200 GB

We used Benchmark Factory to simulate users and transactions for a TPC-E-like workload configuration. The TPC-E benchmark is an OLTP workload that was designed by the Transaction Processing Council (TPC) to test database performance with a mixture of read-only and update-intensive transactions. Our TPC-E-like database workload enables customers to objectively measure performance of the database infrastructure by simulating transactions that are found in enterprise OLTP applications. Our modified TPC-C workload, which is not directly comparable to official benchmarks submitted to and validated by the TPC organization, has the following characteristics:

- User load of 250
- Embedded wait time of 40 ms
- Scale factor of 41
- Database size of 500 GB

Our performance and manageability use case demonstrates how to quickly provision SQL Server databases on a single-vendor compute and storage platform that scales. The Benchmark Factory configuration that we used simulated a medium transactional workload profile to create an incrementally increasing resource demand on the SQL Server infrastructure, as shown in the following figure. Each VM hosted one SQL Server 2017 instance and database with requests coming from the Benchmark Factory OLTP workload. Incremental load tests are useful because they demonstrate scalability—the

capability of the infrastructure to support increased workloads with no configuration changes or impact to the system.



**Figure 7. Total TPC-E-like user workload per test**

The OLTP workload was evenly distributed across the servers through a simple round-robin placement at the time of addition to prevent overloading any one server and creating an artificial performance bottleneck. An evenly distributed load simulates a dedicated virtual cluster that uses VMware Distributed Resource Scheduler (DRS) or manually administered clusters for SQL Server. In this use case, we placed the VMs as shown in the following table to ensure efficient use of the PowerEdge R740 servers.

**Table 6. Incremental workload placement for OLTP tests**

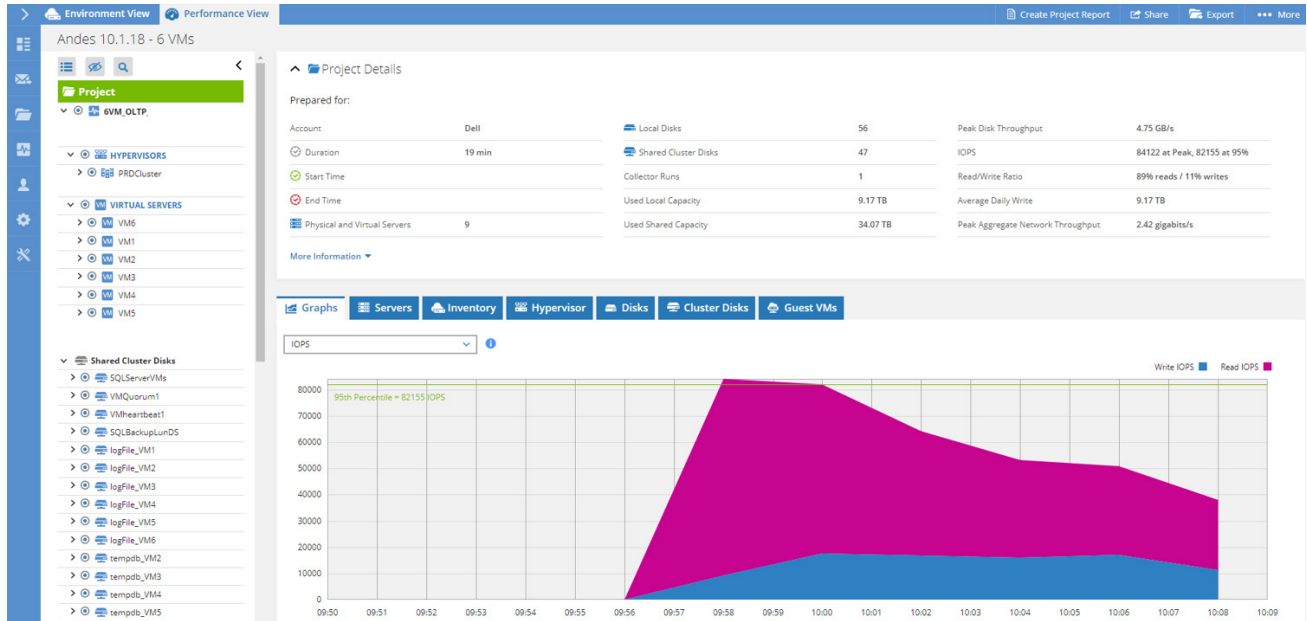
Test	Servers and VMs					
	R740 – 1		R740 – 2		R740 – 3	
	VM 1	VM 2	VM3	VM 4	VM 5	VM 6
Test 1	•					
Test 2	•		•			
Test 3	•		•		•	
Test 4	•	•	•		•	
Test 5	•	•	•	•	•	
Test 6	•	•	•	•	•	•

**Data collection for performance testing**

One of the primary tools that Dell EMC uses for data collection during validation and use case tests is Dell EMC Live Optics, which is a free, agentless software for collecting data from PowerEdge servers. In just minutes, an engineering team can set up Live Optics to collect a wealth of information for configuration and resource utilization analysis. The Live

Optics dashboard is intuitive and enables DBAs to monitor and collect data across the server and VMware virtualization layers.

The following figure shows an example of a Live Optics dashboard. The left side shows performance at the project, hypervisor, virtual server, and shared disk levels. The right side shows the collected data and graphs that can assist with quick visual analysis.



**Figure 8. Live Optics dashboard from our tests**

Benchmark Factory provides detailed Microsoft Excel reports for analysis. For example, we used Benchmark Factory reports to develop the test findings on TPS and Batch Requests/sec. The Benchmark Factory user manual fully explains Benchmark Factory reports, making the data easily understandable for any DBA who is interested in running workload tests.

Unity storage arrays provide multiple ways to review storage-related data for analysis, including a detailed dashboard that enables monitoring and reporting. Administrators can right-click a graph in the Unity dashboard and export the displayed performance data. In our tests, we extracted Unity archive files for performance analysis; for example, we generated our overview of storage I/O latency and IOPS by using the Unity archive files. The Unity array saves up to 7 days of performance data to enable historical analysis. The Unity reports provide the capability for the IT team to perform detailed analysis.

## OLTP workload results

This section details the following findings:

- Peak CPU utilization
- Transactions per second
- Batch Requests/sec
- Average read latency for data
- Average write latency for data and log
- Average IOPS per SQL Server database



### Peak CPU utilization

Peak CPU utilization indicates the high-water mark of CPU usage as a percentage during the incremental use case tests. Each VMware virtualized database used 24 vCPUs for the OLTP workload tests. The following figure shows the peak CPU utilization across all tests. Most virtualized databases had a peak CPU utilization of 71.5 percent. The two Intel Xeon 8180 processors (28 cores per CPU) in our PowerEdge servers handled all the incremental workload without exceeding a peak CPU utilization of 75 percent. The X-axis shows the test and VM numbers, as defined in Table 6 on page 30.




Figure 9. Peak CPU utilization

VMware recommends starting with fewer vCPUs to avoid CPU scheduling conflicts, which can affect the performance of databases. Add vCPUs as the database demands more processing power. For this Benchmark Factory workload, our findings show that the number of vCPUs was more than sufficient for our tests.


### Transactions per second

TPS measures the number of transactions that the SQL Server database can complete in 1 second. SQL Server is an ACID-compliant database, which means that all transactions must be fully completed or completely failed and rolled back to maintain data integrity. Transactions can create, read, update, and delete (CRUD) data in the database; thus, all these operations need transactional consistency support even in the absence of explicit BEGIN and END transaction statements.

TPS is a relative workload indicator that DBAs can use to evaluate database activity for a given application and time period. In this scalability use case, the goal is to show how the infrastructure scales, as new work is added, with minimal reduction in TPS for currently running workloads. In the first test run with one virtualized SQL Server database, the solution achieved 4,206 TPS. The second test, with two virtualized databases, nearly doubled the TPS by generating 8,037 TPS. With each incremental load test, the solution demonstrated strong near-linear scalability, as shown in the following figure. In the final test, with all six virtualized databases running in parallel, the solution achieved 24,533



TPS. Each virtualized database in test 6 generated approximately 4,088 TPS, which is 2.9 percent fewer TPS than observed in test 1 with just one virtualized database, showing minimal deviation from linear scalability.



**Figure 10. Total TPS per test**

When evaluating performance, considering the entire system behavior is important. For example, in test 1, the PowerEdge R740 and Unity 650F systems were dedicated to one virtualized SQL Server database. Comparatively, in test 6, each PowerEdge server ran two virtualized databases, and the Unity array supported all six databases. Between test 1 and test 6, the doubling of the workload on each PowerEdge server and the six-times increase of workload on the array resulted in a minor difference of 2.9 percent TPS.


The average TPS is another indicator of performance sustainability across increasing workloads. The average TPS across all virtualized databases in all the tests was 4,086. Our test findings show that the Ready Solution for SQL Server 2017 with Unity storage scales well under increasing workloads. Delivering consistent performance under increased incremental workload means that this solution can be used as a consolidation platform for your new or existing SQL Server ecosystem.

### Batch Requests/sec

Batch Requests/sec is a SQL Server implemented performance counter that captures the number of received batch requests per second. Batch requests consist of multiple SQL Server statements in a group or in a stored procedure. Batch Requests/sec is useful to DBAs as a comparative metric to determine the state of the SQL Server database. In this use case, Batch Requests/sec is used to measure scalability of the SQL Server database on the Unity array. The goal is to incrementally increase workload to evaluate the impact on batch requests per second.

The following figure shows that the Ready Solution for SQL Server demonstrated strong scalability from test 1 through test 6. For example, test 2 with 161,584 Batch Requests/sec nearly doubled the performance of test 1 with 84,527 Batch Requests/sec.

Test 6, which, with six virtualized databases, was the most demanding, generated 493,666 Batch Requests/sec, meaning that each of the six databases generated approximately 82,278 Batch Requests/sec —2.7 percent less than test 1. The average Batch Requests/sec across all virtualized databases in all the tests was 82,057.




**Figure 11. Total Batch Requests/sec per test**

Because DBAs use Batch Requests/sec as a key indicator of system performance, achieving near-linear scalability under increasing workloads is critical. Our tests show that the Ready Solution for SQL Server can scale in terms of Batch Requests/sec with a strong degree of performance consistency.

### Average read latency for data

OLTP workloads are characterized predominately by small reads and writes to storage. A key indicator of storage performance for OLTP workloads is physical read and write latencies. The lower the latency the less time the database waits for reads and writes from storage. The gold standard for all-flash storage arrays is an average of 1 millisecond (ms) or less for all physical database operations. Modern storage arrays from Dell EMC have improved upon the gold standard by accelerating storage operations, frequently achieving latencies of .75 ms or less.


One of the benefits that is derived from incremental load methodology is the development of a performance trend line. A flat horizontal trend line in the following figure would indicate that latency did not increase as the workload increased. Our initial goal was to show no more than a slight increase in read latency in moving from one to six virtualized databases and a trend line with a slight slope, indicating good scalability. Findings show that read latency remained under the .75 ms threshold through test 5, with a five-database workload. As indicated in the following figure, only in test 6, with all six databases, did the read latency exceed the .75 ms threshold, but it remained under .8 ms. The value of lower latencies depends on the percentage of database reads. Across our test cases, average read percentage for data was 74.5 percent, meaning that the fast .65 ms average read latency accelerated the majority of database storage operations.



**Figure 12. Storage average read latency, in milliseconds, for data LUNs**


### Average write latencies for data and log

Write latency is critically important because it supports the durability aspect of ACID-compliant databases. In terms of writes, our findings are separated into writes to the data LUNs and writes to the log LUNs. By default, beginning with SQL Server version 2016, the database issues a checkpoint every minute. A checkpoint is a process in which all in-memory modified pages are saved to storage and the active portion of the transaction log is updated for persistence. The trend line in the following figure shows a slight increase in write latency as our workloads increased. Write latency remained under the .75 ms threshold from test 1 through test 4. In tests 5 and 6, write latency exceeded the .8 ms threshold; however, all average writes across all databases remained under the gold standard of submillisecond performance. Across our test cases, the average write percentage for data was 25.5 percent, meaning that the fast .76 ms average write latency accelerated one-fourth of database storage operations.



**Figure 13. Storage average write latency, in milliseconds, for data LUNs**

Transaction logs in an SQL Server database record logical and physical changes to data, enabling the database to be recovered in the case of an unplanned outage. During normal database operations, most of the transaction log activity is weighted toward writes as the database is recording changes to data. This changes during database recovery as the database engine reads from the transaction log to bring the database to a consistent state. In our tests, the database operated normally; thus, for logs, we reviewed only write latency. As shown in the following figure, the average write latency for logs across all the tests did not exceed the .6 ms threshold.




**Figure 14. Storage average write latency, in milliseconds, for log LUNs**

## Average IOPS per SQL Server database

IOPS is a metric that indicates the load on a storage array. Incremental workload tests are not designed to maximize load and drive the highest IOPS values; instead, the goal is to show that IOPS increase predictably as database loads increase. The capability of scaling database storage workload enables the use of more databases without significantly affecting performance.

In test 1, the solution achieved 10,317 IOPS, which was our high-water mark for any database throughout the incremental workload tests. Thus, our goal was to demonstrate that the Unity 650F array can consistently deliver near 10,317 IOPS for each database as the workload increased. Findings show that the average IOPS across all the tests for all the databases was 10,264. The percentage difference between the IOPS in test 1 and the average IOPS across all the tests is 1 percent.



**Figure 15. Average IOPS per incremental test**

In summary, test findings show:

- Our incremental load tests did not saturate the Unity 650F array; IOPS increases were approximately linear.
- The Unity 650F array demonstrated strong IOPS scalability as workloads increased, with a minor 1 percent delta between the first test and the total average IOPS across all tests.

## Decision-support use case

### Decision-support workload implementation

The following table lists the decision-support workload implementation details.

**Table 7. Decision-support workload implementation details**

Characteristic	Value
Number of VMs	3
Number of data file LUNs	3
Number of log file LUNs	3
Number of tempdb data file LUNs	3
Number of tempdb log file LUNs	3
<b>Per VM configuration</b>	
vCPUs	48
VM memory	352 GB
Memory for database	160 GB
Database data file LUN size	500 GB
Database log file LUN size	200 GB
tempdb data file LUN size	200 GB
tempdb log file LUN size	200 GB

### Decision-support test methodology

In addition to running OLTP workloads, we ran a decision-support workload by using the HammerDB TPC-H test. The decision-support test simulates parameterized business reporting queries to assist with decision analysis. The test specification also includes concurrent data modifications that change multiple sets of data in parallel. In addition, the test includes complex queries; the database must join and aggregate (filter or group by) large sets of data that are typical of business decision analysis. Our findings are for a TPC-H-like test, which means that they are not certified or comparable to published and validated results.

We used a scale factor of 1,000 for the decision-support testing. Scale factor defines the generated test database size. A scale factor of 1 generates approximately 1 GB of test data, so our scale factor of 1,000 generated a database with 1,000 GB of data. The scale factor also defines the minimum number of query streams—that is, a set of queries that must be run serially, one after another. The TPC-H specification requires a minimum of seven query streams for a scale factor of 1,000.

Our focus in running the decision-support workload was to report on the throughput test results. Throughput tests measure the ability of the Ready Solution to process a set number of queries in as little time as possible. We ran our tests without a refresh stream. Refresh streams run update procedures to periodically refresh the dataset in use. Each decision-support test consisted of 22 queries. The throughput findings are an indicator of the solution's capability of supporting decision workloads consisting of multiple users.

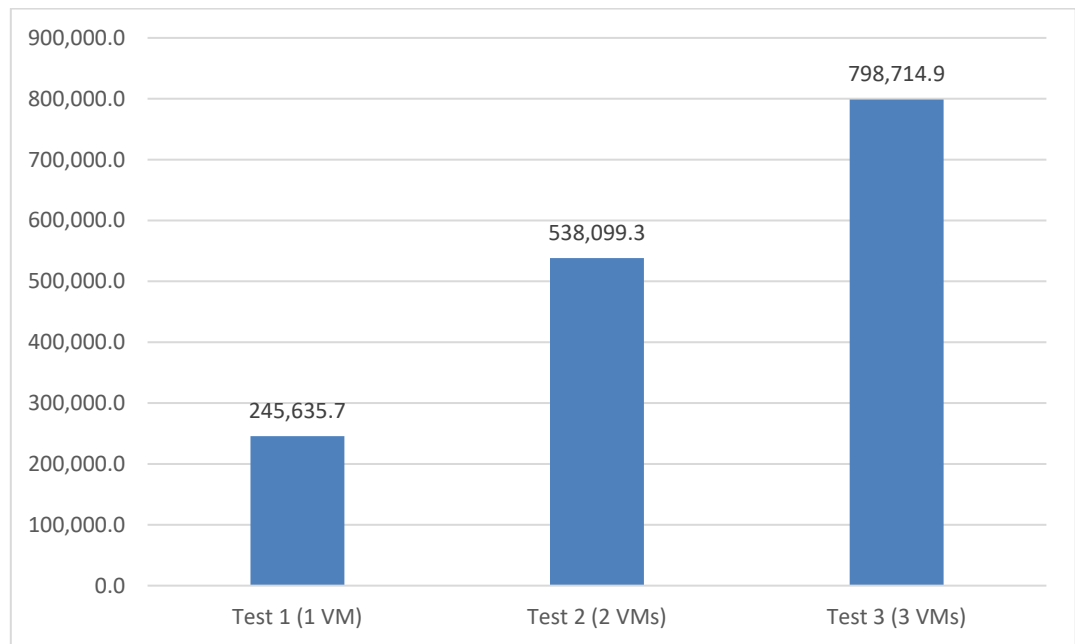
We used the same incremental workload test methodology across three virtualized SQL Server databases, with one database hosted on each PowerEdge server. As shown in the following table, we added one incremental decision-support workload for each test, with the third and final test running all three workloads.

**Table 8. Incremental workload placement for decision-support tests**

Test	Servers and VMs		
	R740 – 1	R740 – 2	R740 - 3
	VM 1	VM 2	VM 3
Test 1	•		
Test 2	•	•	
Test 3	•	•	•

### Decision-support workload results

The following figure shows the Throughput @ Size totals for each incremental test. The Throughput @ Size indicates the number of queries per hour multiplied by the scale factor. The Throughput @ Size value indicates the capability of the system to support more decision-support queries; thus, the higher the value the better. Because this test is a noncertified TPC-H test, we make no comparisons to other systems; instead, we analyze the findings to determine if the solution scales well with decision-support workloads.

**Figure 16. Throughput @ Size totals**

In test 1, the Ready Solution achieved a Throughput @ 1,000 value of 245,635.1, which we used as a measure for tests 2 and 3 to evaluate if the system was scaling with the incremental workloads. In test 2, the two decision-support workloads exceeded the Throughput @ 1,000 of test 1: VM 1 achieved 276,095.6 and VM 2 achieved 262,003.7, for a total of 538,099.3. Test 3 also exceeded test 1: VM 1 achieved 253,150.6, VM 2 achieved 279,154.0, and VM 3 achieved 266,410.3, for a total of 798,714.9. We ran all tests without a refresh stream.

The test findings for this solution show that the system delivered linear scalability across the three incremental workload tests. These decision-support tests did not challenge the servers and Unity system to the extent that performance was affected. However, the test

results indicate that a properly sized solution enables the scaling of decision-support workloads with minimal impact to the system.

## Summary

### OLTP findings summary

Testing of this Ready Solution for SQL Server with the Unity 650F array demonstrated that the solution has the capability to scale with incremental increases in workloads. Key metrics such as TPS, Batch Requests/sec, physical I/O latencies, and IOPS showed that this solution is a proven platform for SQL Server ecosystems because of its capability to scale with workload demands. The solution demonstrated strong performance results as well. Average latency across all the OLTP tests did not exceed the 1 ms threshold, and average latency for the majority of tests was under the .75 ms threshold. With consistently low latency, enterprise applications respond faster to the end user, accelerating business processing.

In summary, test findings include:

- The PowerEdge R740 servers with Intel Xeon 8180 processors and 24 vCPUs per VM seamlessly managed the SQL Server workloads, never exceeding 71.5 percent peak CPU utilization.
- TPS is a good indicator that DBAs can use to evaluate database activity. The Ready Solution demonstrated strong TPS scalability by delivering consistent TPS across all the incremental tests. The top TPS score was achieved in test 1 at 4,206, and, as the workload increased, there was a minor difference between the top TPS score and test 6. With six workloads running in parallel, the average TPS score was 4,088, which is 2.9 percent difference from the top TPS score. The average across all the tests and databases, 4,086 TPS, was similar to the TPS that was achieved during test 6.
- Batch Requests/sec is a useful comparative metric that DBAs can use to determine the level of activity for the SQL Server database. The Ready Solution showed strength in scalability in terms of Batch Requests/sec. In test 1, the solution achieved 84,527 Batch Requests/sec, and, in test 6, each of the six databases averaged 82,278 Batch Requests/sec—only a 2.7 percent difference from test 1. The solution delivered near-linear scalability across the six incremental workload tests.
- For OLTP workloads, where the database activity is predominantly small reads and writes, a critical performance metric is latency. Through the first five tests, average read latency remained under the .75 ms goal, showing fast storage-response times. Only in test 6 was the .75 ms threshold exceeded, but read latency remained under .8 ms.
- For writes, we captured two key metrics. Average write latency for the data LUNs provides insight into how fast data was modified or written anew to storage. Average write latency for the log LUNs provides insight into how fast data was saved to storage for persistence and recoverability.
  - **Data**—Average write latency for tests 1 through 4 remained under the .75 ms threshold. In tests 5 and 6, average write latency exceeded the .8 ms threshold, but all average writes across all databases remained under the gold standard of submillisecond performance.



- **Log**—The average write latency across all the tests did not exceed the .6 ms threshold.
- IOPS is a metric that indicates the load on a storage array. Incremental workload tests are not designed to maximize load and drive the highest IOPS values; instead, the goal is to show that IOPS consistently scales for the databases as load increases. The solution showed strong IOPS scalability in that the top score was 10,317 in test 1, and the average across all the tests was 10,264 IOPS—a 1 percent difference from test 1. Thus, the solution scaled IOPS near-linearly in our incremental workload tests.

### Decision-support findings summary

We used the HammerDB TCP-H test to simulate a decision-support workload. A decision-support workload is different from an OLTP workload in that it consists of longer-running ad-hoc, complex queries that read and modify multiple datasets. We used a similar incremental workload approach with three virtualized databases. Test analysis focused on Throughput @ Size, which is the number of queries per hour multiplied by the scale factor. The decision-support-workload test results, as shown in the following table, showed linear scalability as the workload increased from one decision-support system through three systems running in parallel.

**Table 9. Decision-support test results**

VM	Throughput @ Size		
	Test 1	Test 2	Test 3
VM 1	245,635.7	276,095.6	253,150.6
VM 2		262,003.7	279,154.0
VM 3			266,410.3
Totals	245,635.7	538,099.3	798,714.9

The Ready Solution for Microsoft SQL Server with Unity 650F has demonstrated strong scalability and performance, and the capability of addressing OLTP and decision-support workloads. With this solution, you can consolidate your databases and gain efficiencies in managing your database ecosystem.

# Chapter 5 References

This chapter presents the following topics:

<b>Dell EMC documentation .....</b>	<b>43</b>
<b>Microsoft documentation.....</b>	<b>43</b>

## Dell EMC documentation

For links to additional Dell EMC documentation for this solution, see [Microsoft SQL Info Hub for Ready Solutions](#).

The following links provide additional and relevant information. Access to some documentation depends on your login credentials. If you do not have access to a document, contact your Dell EMC representative.

- [Dell EMC Connectrix DS-6610B FC Switch Specification Sheet](#)
- [Dell EMC Live Optics](#)
- [Dell EMC PowerEdge R740 Spec Sheet](#)
- [Dell EMC PowerPath/VE Data Protection Software](#)
- [Dell EMC Unity Best Practices Guide](#)
- [The Future-Proof Loyalty Program](#)

## Microsoft documentation

The following Microsoft documentation provides additional and relevant information:

- [Database Checkpoints \(SQL Server\)](#)
- [What's new in SQL Server 2017](#)
- [What's new in Windows Server 2016](#)