

200TB Data Warehouse Fast Track Reference Architecture for Microsoft SQL Server 2017 using Dell EMC PowerEdge R940 and SC9000

Configuration and performance results

Abstract

This paper describes the design principles and guidelines used to achieve an optimally balanced 200TB Data Warehouse Fast Track reference architecture for SQL Server 2017 using Dell EMC™ PowerEdge™ R940 servers and Dell EMC SC9000 arrays.

November 2018

Revisions

Date	Description
November 2018	Initial release

Acknowledgements

Author: Doug Bernhardt

Special thanks to Jamie Reding at Microsoft.

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.

© 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA. [10/30/2018] [Reference Architecture] [4032-RA-SQL]

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	5
1 Microsoft SQL Server 2017: Data warehousing with improved column-store technology	6
1.1 Dell EMC Data Warehouse Fast Track reference architectures for SQL Server 2017	6
2 Recommended reference architectures	7
2.1 Single-server reference architecture	7
2.2 Highly available reference architecture	9
3 Hardware components	11
3.1 PowerEdge R940 server	11
3.2 QLogic QLE2662 16Gbps host bus adapter	11
3.3 Brocade 6505 SAN switch.....	11
3.4 SC9000 array.....	11
4 Storage configuration	13
4.1 Cabling.....	13
4.2 I/O ports	14
4.3 Disk folder.....	14
4.4 Storage type	14
4.5 SC Series storage profiles.....	14
4.6 Read and write cache.....	15
4.7 SC Series server object.....	15
4.8 SC Series volumes	15
5 Server configuration	17
5.1 System BIOS	17
5.2 Host bus adapters	17
6 Windows Server 2016 configuration	18
6.1 Power plan.....	18
6.2 Lock pages in memory	18
6.3 Windows volumes.....	18
6.4 MPIO.....	20
7 SQL Server 2017 Enterprise Edition configuration	21
7.1 Grant perform volume maintenance task privilege.....	21
7.2 SQL Server maximum memory	21

Table of contents

7.3	Max degree of parallelism (MAXDOP)	21
7.4	Resource governor	21
7.5	Database configuration.....	21
7.6	Tempdb configuration.....	21
8	Additional considerations for the highly available reference architecture	22
9	DWFT for SQL Server 2017 certification.....	23
10	Summary	24
A	Technical support and resources	25

Executive summary

Dell EMC™ and Microsoft®, in cooperation, provide guidelines and principles to assist customers in designing and implementing a balanced configuration for Microsoft SQL Server® data warehouse workloads to achieve out-of-the-box scalable performance. These database reference architectures enable each of the components in the database stack to provide optimal throughput to match the database capabilities of the specific setup. Innovative Dell EMC PowerEdge™ 14th-generation (14G) servers, along with robust and cutting-edge Dell EMC SC9000 storage arrays, form efficient candidates for a high-performing data warehouse solution.

This paper describes the design principles and guidelines used to achieve an optimally balanced 200TB Data Warehouse Fast Track (DWFT) reference architecture for SQL Server 2017 using PowerEdge R940 servers and SC9000 arrays. The configuration used to achieve the performance numbers for the reference configuration is presented in detail.

The target audience for this paper includes database administrators, business intelligence architects, storage administrators, IT directors, and data warehousing users seeking sizing and design guidance for business intelligence solutions with SQL Server 2017.

1 Microsoft SQL Server 2017: Data warehousing with improved column-store technology

Microsoft SQL Server 2017 has made significant improvements in data warehousing technologies and performance, including column-store features as well as many other improvements.

Column-store indices offer great advantages over traditional row stores for analytics and data warehousing queries. They are ideally suited for the star schemas, and tables with billions of rows which are commonly seen. Some of their advantages for analytics include the following:

- **Up to 10X compression in data size:** Data warehouses are very large by nature and the compression offered by column-store index technologies offers both space and cost savings as well as significantly increased performance. These benefits are possible due to the dramatically reduced I/O requirements given by the compression and coupled by the ability to only scan the specific columns required by each query. This compression also reduces the amount of memory required to hold a given number of rows from the source data warehouse.
- **Additional indices:** SQL Server 2017 adds the capability to add B-Tree indices to column-store-based tables, which enables efficient single-row lookup.

In addition to these architectural features, Microsoft has further optimized the processing of queries in column-store indices in the following ways:

- **Operator pushdown:** Pushdown refers to moving both filter and aggregation query operations closer to the data, so that many of the filters and calculations can be done in the scan operators, dramatically reducing the volume of data that needs to be handled further on in-query processing.
- **Batch-mode processing:** SQL Server 2017 includes enhancements in batch-mode processing that handles many rows at a time rather than serially performing calculations on each individual row. These batch operations are further optimized by leveraging Single Instruction Multiple Data (SIMD) vector processing CPU instructions in the Intel® architectures.

1.1 Dell EMC Data Warehouse Fast Track reference architectures for SQL Server 2017

Dell EMC and Microsoft have refreshed the DWFT reference architecture offerings with the latest technology advancements in database, server, and storage technology. PowerEdge 14G servers with Intel® Xeon® processors, SC9000 arrays, and SQL Server 2017 are the latest additions to the list of reference architecture components.

The Dell EMC DWFT reference architectures for SQL Server 2017 are engineered jointly by Dell EMC and Microsoft. The hardware and software optimizations are tested by Dell EMC and the performance results are crosschecked by Microsoft. This approach presents a fast time to value using integrated, balanced, and verified architectures.

2 Recommended reference architectures

The following subsections describe the two DWFT reference architectures for SQL Server 2017, comprised of PowerEdge R940 servers and SC9000 arrays.

2.1 Single-server reference architecture

Figure 1 illustrates the single-server reference architecture with the major elements. Table 1 lists the component details.




Figure 1 Single-server reference architecture

Table 1 Single-server reference architecture details

Component	Description	
Server	PowerEdge R940	
	Processors	Four Intel® Xeon® Platinum 8168 Processors (2.7GHz, 24 cores, 48 threads)
	Total cores	96
	Total logical processors	192 (Intel Hyper-Threading Technology enabled)
	Total memory	3072GB
	Network adapters	Minimum of one network adapter (1Gbps or 10Gbps based on requirements) Recommended to have more than one network adapter with load balancing configured
	Host bus adapters	Eight QLogic® QLE2662 dual-port 16Gbps FC HBAs
Software	Operating system	Microsoft Windows Server® 2016 Enterprise Edition
	Database software	SQL Server 2017 Enterprise Edition
Storage	Storage array	SC9000 (SCOS v7.3.2)
	Disk drives	48 x 3.84TB read-intensive SSDs (2.5-inch SAS)
	SAN switches	Two Brocade® 6505 with 16Gbps SFPs

Table 2 Server firmware versions

Hardware component	Firmware version
BIOS	1.4.9
Backplane 1	3.56
Dell™ 64 Bit uEFI Diagnostics	4301A18
Dell OS Driver Pack	18.04.07,A00
Integrated Dell Remote Access Controller (iDRAC)	3.21.23.22
Broadcom® Gigabit 4P X520/I350 rNDC	20.8.4
Lifecycle Controller	3.21.23.22
OS Collector	3.0
Power supply	00.32.55
QLogic QLE2662	14.07.05
System CPLD	0.4.0
PowerEdge RAID Controller (PERC) H730P adapter	25.5.5.0005

2.2 Highly available reference architecture

For database high availability, Microsoft Windows® failover clustering is recommended. Using Microsoft clustering services, one database server is configured as the primary (active) server and the second server is configured as the secondary (passive) server. The secondary server should have exactly the same configuration as the primary server. Since the database is only active on a single server at any point of time, the performance of the database on the primary server (active) is comparable to the single-server configuration (discussed in section 2.1).

Figure 2 illustrates the highly available reference architecture with the major elements. Table 33 lists the component details.

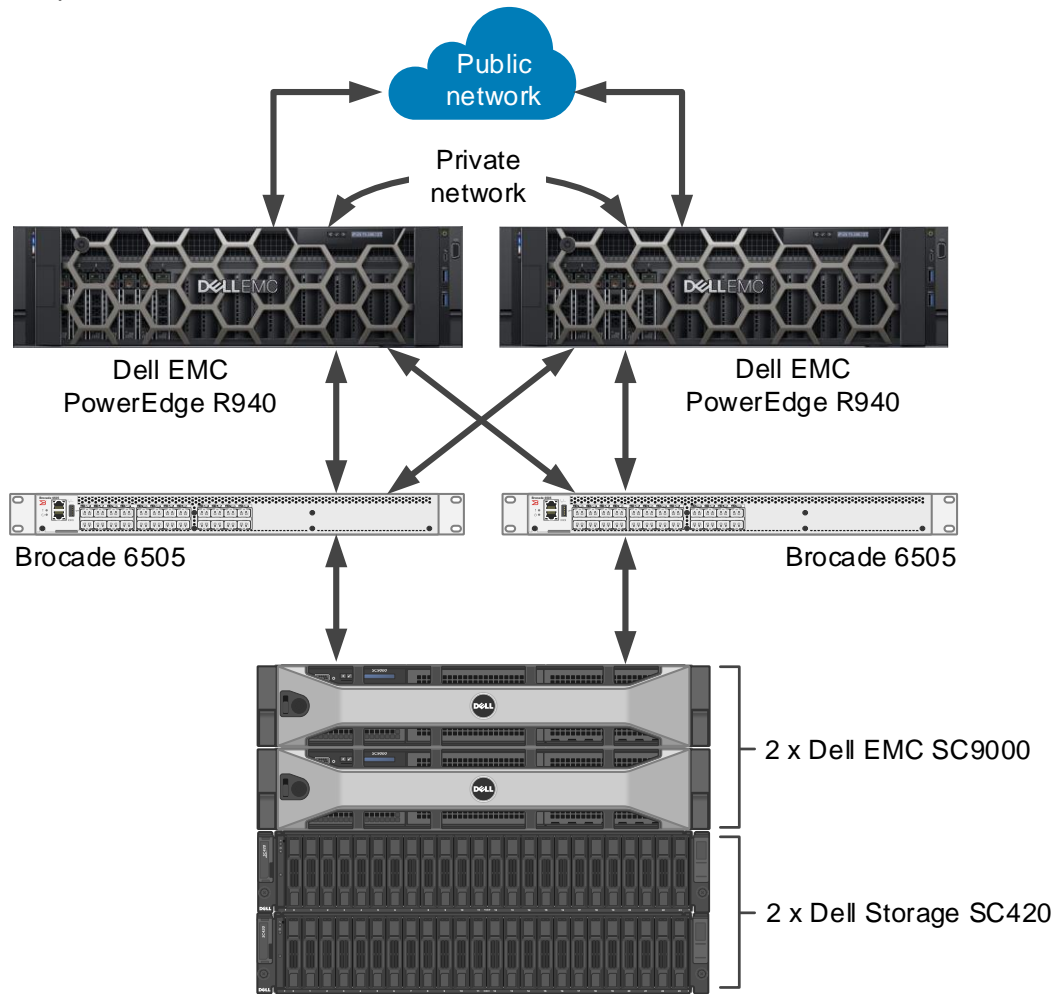


Figure 2 Highly available reference architecture

Table 3 Highly available reference architecture details

Component	Description	
Server	Two PowerEdge R940	
	Processors	Four Intel® Xeon® Platinum 8168 Processors (2.7GHz, 24 cores, 48 threads)
	Total cores	96 per server
	Total logical processors	192 per server (Hyper-Threading Technology enabled)
	Total memory	3072GB per server
	Network adapters	Public network: Minimum of two network adapters (1Gbps or 10Gbps based requirements) per server with load balancing configured Private network (cluster): Minimum of one network adapter with 1Gbps (or more) connectivity per server
	Host bus adapters	Eight QLogic QLE2662 dual-port 16Gbps FC HBAs per server
Software	Operating system	Windows Server 2016 Enterprise Edition with Windows Failover Clustering
	Database software	SQL Server 2017 Enterprise Edition configured as a Failover Cluster Instance
Storage	Array	SC9000 (SCOS v7.3.2)
	I/O cards	Four QLogic QLE2694 quad-port 16Gbps FC HBAs (two per controller)
	Disk drives	48 x 3.84TB read-intensive SSDs (2.5-inch SAS)
	SAN switches	Two Brocade 6505 with 16Gbps SFPs

3 Hardware components

3.1 PowerEdge R940 server

The PowerEdge R940 server is a highly versatile, 4U rack server with up to four Intel® Xeon® Scalable processors, a large memory footprint, extensive I/O options, and a choice of dense, high-performance storage or low-cost, high-capacity storage. The R940 server offers simplified management, purposeful design, and energy efficiency with support for the latest Intel® Xeon® series processors as well as NVDIMM memory. The R940 server supports up to thirteen Gen3 PCIe expansion cards.

For more information, see the [PowerEdge R940 server product page](#).

3.2 QLogic QLE2662 16Gbps host bus adapter

The reference architecture makes use of eight QLogic QLE2662 dual-port 16Gbps Fibre Channel (FC) to PCI Express host bus adapters.

For more information, see the QLogic FAQ document, [QLogic 2600 Series 16Gb Gen 5 FC HBAs for Dell PowerEdge Servers](#).

3.3 Brocade 6505 SAN switch

The Brocade 6505 switch is a 1U, 24-port, rack-mountable Fibre Channel switch providing up to 16Gbps of bandwidth per port. This switch enables organizations to simplify IT infrastructures, improve system performance, maximize the value of virtual server deployments, and reduce overall storage costs.

For more information, see the [Brocade 6505 switch product page](#).

3.4 SC9000 array

The SC9000 array makes storage cost savings automatic with a modern architecture that optimizes the data center for economics while delivering transformational SSD, HDD, or hybrid performance.

SC Series storage provides the lowest effective cost per GB for flash and hybrid flash¹, giving organizations of any size the technology advantage needed to compete in the current fast-changing markets. Highlights include:

- Data Progression: Achieve IOPS goals with the least-expensive mix of storage media, even as your performance needs evolve.
- Deduplication and compression: Dramatically reduce the raw capacity required to store data.
- RAID tiering: Eliminate manual RAID provisioning, and increase efficiency and utilization.
- Federation: Simplify multi-array environments with quick and seamless data movement, plus proactive load balancing assistance using Live Migrate and Volume Advisor.

¹Net usable capacity of Dell array with 5 years of support, after 4:1 data reduction, vs. major competitors net of data reduction. Street price analysis is based on a variety of sources including analyst data, price sheets when available, and public information as of January 2017.

- Dell ProSupport™ services: Reduce deployment costs with remote installation options that ensure the project is successful the first time.
- Persistent software licensing: Future-proof the investment, and minimize the cost of upgrades and expansions.

With up to 1,024 total drives and 6PB² raw capacity in a single array, the SC9000 is massively scalable. All data is managed through the primary storage controller, deployed in redundant pairs. To expand capacity, you can mix and match from any of the five supported expansion enclosures, depending on your specific performance, cost, or density needs.

Options are available for ultra-dense (84 drives in 5U) and high-performance 12Gb SAS enclosures. Adding a new enclosure is simple — the new drives can be automatically assimilated into the existing virtualized pool, or set up as their own standalone pool. Data-in-place upgrades² make it easy to change controllers and other key components, or migrate from legacy SC Series arrays.

²For specifications, see the [SC Hybrid Arrays spec sheet](#) and [SC9000 product page](#).

4 Storage configuration

4.1 Cabling

The hardware components were connected using Dell EMC best practices.

Port 1 on each HBA in the server and ports 1 and 2 on each HBA in the SC9000 were connected to the same Brocade 6505 switch. Port 2 on each HBA in the server and ports 3 and 4 on each HBA in the SC9000 were connected to the other Brocade 6505 switch.

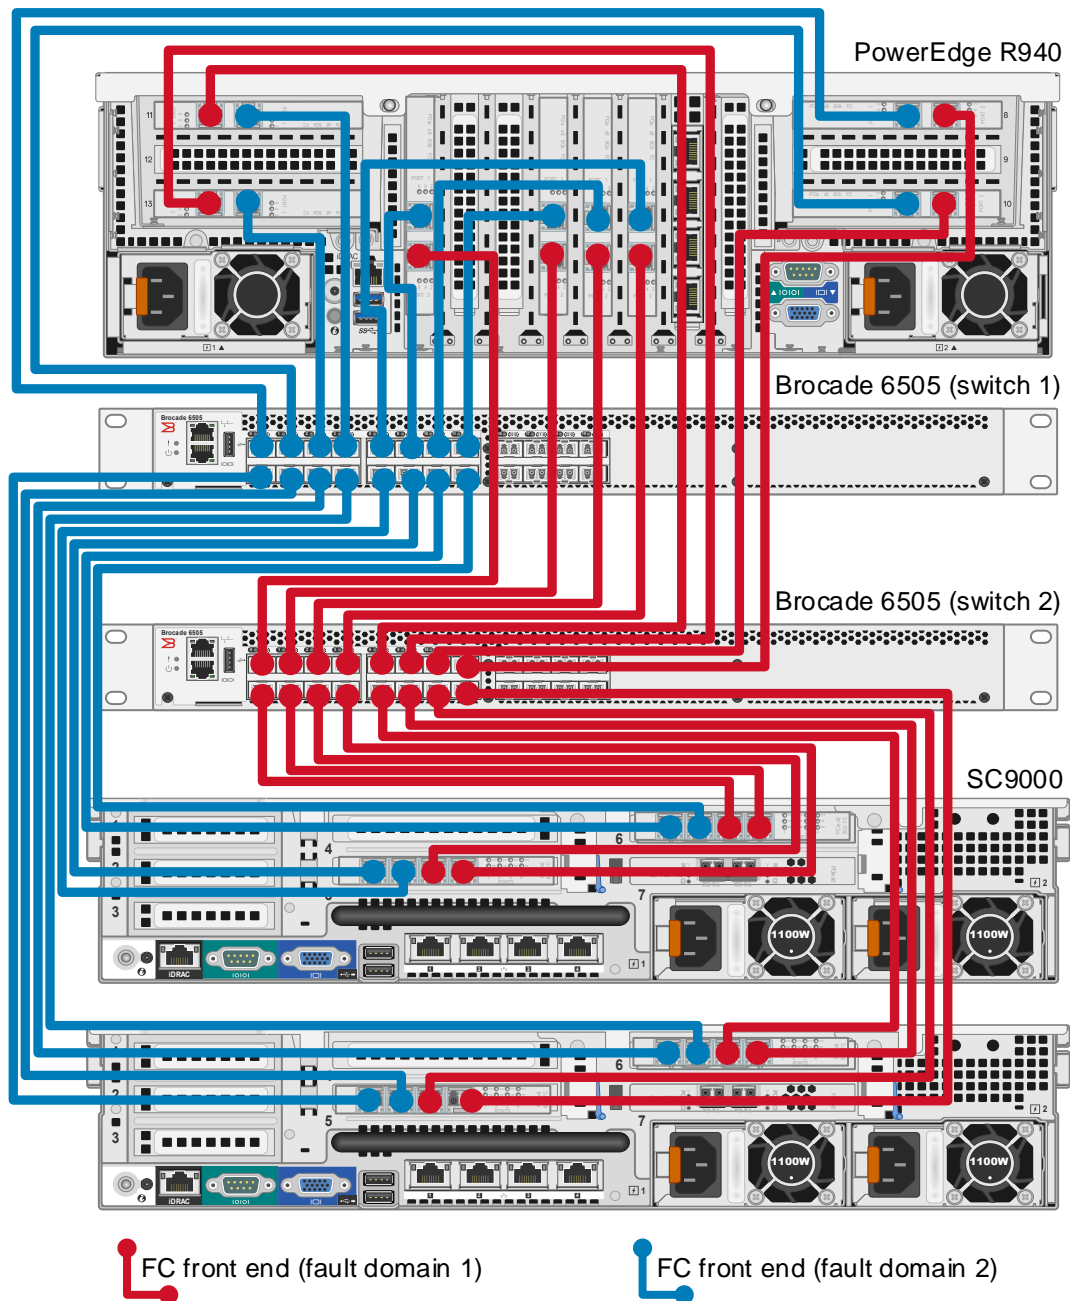


Figure 3 Single-server configuration cabling diagram

4.2 I/O ports

The FC ports were configured to use two fault domains in virtual port mode. Ports 1 and 2 from each HBA on each controller were put into fault domain 1 and ports 3 and 4 from each HBA on each controller were put into fault domain 2.

4.3 Disk folder

All 48 disks were added to the **Assigned** disk folder. The SC Series distributed spare feature introduced with SCOS 7.3.1 makes all disks active, so none are reserved as dedicated spares.

4.4 Storage type

The disk folder storage type defines the page size on the array and how the disks are configured for redundancy (single or dual). The disk folder can be configured to use a 512KB, 2MB (default), or 4MB page.

The **Assigned** disk folder was configured for single redundancy, using a page size of 512KB. The 512KB page size provides the best performance in all-flash configurations.

4.5 SC Series storage profiles

The array uses storage profiles to define the RAID level and tiers where the data is stored for a given volume. By default, newly created volumes use the **Recommended (All Tiers)** storage profile. This storage profile uses RAID 10 for active (writeable) data and RAID 5 for snapshots (replays).

To increase capacity, the reference architecture uses RAID 5 for all data stored on the SQL Server data volumes. A custom storage profile named **RAID 5 All Tiers** was created. SQL Server data volumes on the array use this custom storage profile to ensure that RAID 5 is always used for those volumes. Because the workload is read-only, there is no performance penalty in a RAID 5 configuration for data volumes. All other volumes use the **Recommended (All Tiers)** storage profile.

The RAID 5 stripe width is left at the default setting of **9 Wide (89% Efficient)** to maximize capacity on the array.

Create Storage Profile

General
 Name: RAID 5 All Tiers
 Notes:

Data Writes
 Select the write tier and RAID type to be used for writing data.
 Write Tier: Tier 1
 Write RAID Type: Parity (RAID 5 / RAID 6)

Snapshot Data
 Select the RAID type to use for snapshot data in each tier.
 Tier 1 Snapshot RAID Type: Parity (RAID 5 / RAID 6)
 Tier 2 Snapshot RAID Type: Parity (RAID 5 / RAID 6)
 Tier 3 Snapshot RAID Type: Parity (RAID 5 / RAID 6)

Data will be written to Tier 1 Parity (RAID 5 / RAID 6). Snapshots will remain in Tier 1 Parity (RAID 5 / RAID 6) when they are taken. Data Progression will move data amongst all selected tiers based upon frequency of access.

? Help × Cancel ↑ OK

Figure 4 Creating the custom storage profile, RAID 5 All Tiers

4.6 Read and write cache

Read and write cache were both enabled on the array.

4.7 SC Series server object

The process of mapping a volume to a server object creates the I/O path (or paths) between a volume and a server. When creating a server object, the operating system of the server is defined. The SC Series array includes both single path and multi-path definitions for each version of Windows.

In this configuration, the server object on the array was created using the **Windows Server 2016 MPIO** operating system definition. Since the MPIO version of the Windows 2016 definition was chosen, each volume mapping will contain eight paths.

4.8 SC Series volumes

The **Recommended (All Tiers)** storage profile is assigned to all volumes, with the exception of the SQL Server data volumes. The SQL Server data volumes use the **RAID 5 All Tiers** custom storage profile. This uses RAID 5 for the SQL Server data volumes and RAID 10 for all other volumes.

None of the volumes were configured to use snapshots. The reference architecture does not include hardware resources to support the use of snapshots.

Eight volumes (LUNs) were created to store the SQL Server data files for the data warehouse, with four volumes assigned to each controller. Two volumes were created to store the tempdb data files, with one

volume assigned to each controller. If the volumes are created and mapped to the server, one right after the other, the array should automatically distribute the volumes evenly across the controllers. Otherwise, the controller can be manually selected using the advanced options when mapping the volume to the server. To achieve optimal results, the SQL Server data and tempdb volumes must be evenly distributed across the controllers.

Table 4 SC Series volumes created for the reference architecture

Volume name	Storage profile	Volume purpose
DWFT-4S - MPHost	Recommended (All Tiers)	Mount point host
DWFT-4S - SQLSystem	Recommended (All Tiers)	SQL Server system databases and files
DWFT-4S - SQLLog	Recommended (All Tiers)	Transaction log files for tempdb and the data warehouse
DWFT-4S - SQLData01	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData02	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData03	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData04	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData05	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData06	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData07	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLData08	RAID 5 All Tiers	Data files for the data warehouse
DWFT-4S - SQLTempdb01	Recommended (All Tiers)	Data files for tempdb
DWFT-4S - SQLTempdb02	Recommended (All Tiers)	Data files for tempdb

5 Server configuration

5.1 System BIOS

The system profile is set to **Performance**. All other options, outside of iDRAC configuration, were left at their factory default settings. The **Logical Processor** option, under **Processor Settings**, is left at its default setting of **Enabled**. This enables Intel® Hyper-Threading Technology, which maximizes the number of logical processors available to SQL Server.

5.2 Host bus adapters

After each HBA port was reset to the factory defaults, several parameters were changed in accordance with Dell EMC best practices for SC Series storage arrays. Those changes are detailed in Table 55.

Table 5 HBA parameter changes

Parameter menu	Parameter	Setting
Adapter Settings	Connection Options	1 (Point to Point only)
Advanced Adapter Settings	Login Retry Count	60
	Port Down Retry Count	60
	Link Down Timeout	30
	Enable LIP Reset	Yes

6 Windows Server 2016 configuration

6.1 Power plan

To maximize performance, the server was configured to use the **High performance** power plan as shown in Figure 5.




Figure 5 Windows power plan

6.2 Lock pages in memory

To prevent Windows from paging SQL Server memory to disk, the **Lock pages in memory** option was enabled for the SQL Server service account.

For information on enabling this option, see the article [Enable the Lock Pages in Memory Option](#).

6.3 Windows volumes

A single Windows volume was created on each SC Series volume. All volumes were formatted with the NTFS file system. The boot volume and the mount-point host volume used the default allocation unit. All other volumes used an allocation unit of 64KB.

For DWFT reference architectures, Dell EMC recommends using mount points for the volumes instead of drive letters. It is highly recommended to assign appropriate volume and mount-point names to simplify troubleshooting and performance analysis. Ideally, the mount-point names should be assigned in a way that makes it easy to identify the SC Series volume for a given Windows volume.

Table 6 shows the volume labels and access paths used for the reference configuration.

Table 6 Windows volume details

SC Series volume name	Windows volume label	Access path
DWFT-4S - MPHost	MPHost	M:\
DWFT-4S - SQLSystem	SQLSystem	M:\ft\SQLSystem
DWFT-4S - SQLLog	SQLLog	M:\ft\SQLLog
DWFT-4S - SQLData01	SQLData01	M:\ft\SQLData01
DWFT-4S - SQLData02	SQLData02	M:\ft\SQLData02
DWFT-4S - SQLData03	SQLData03	M:\ft\SQLData03
DWFT-4S - SQLData04	SQLData04	M:\ft\SQLData04
DWFT-4S - SQLData05	SQLData05	M:\ft\SQLData05
DWFT-4S - SQLData06	SQLData06	M:\ft\SQLData06
DWFT-4S - SQLData07	SQLData07	M:\ft\SQLData07
DWFT-4S - SQLData08	SQLData08	M:\ft\SQLData08
DWFT-4S - SQLTempdb01	SQLTempdb01	M:\ft\SQLTempdb01
DWFT-4S - SQLTempdb02	SQLTempdb02	M:\ft\SQLTempdb02

6.4 MPIO

MPIO was configured using Dell EMC best practices. MPIO best practices for the SC Series array are documented in the best-practices paper [Dell EMC SC Series Storage and Microsoft Multipath I/O](#).

The MPIO policy for all volumes is left at the default setting of **Round Robin**, allowing volumes to use all paths between the server and the storage array.

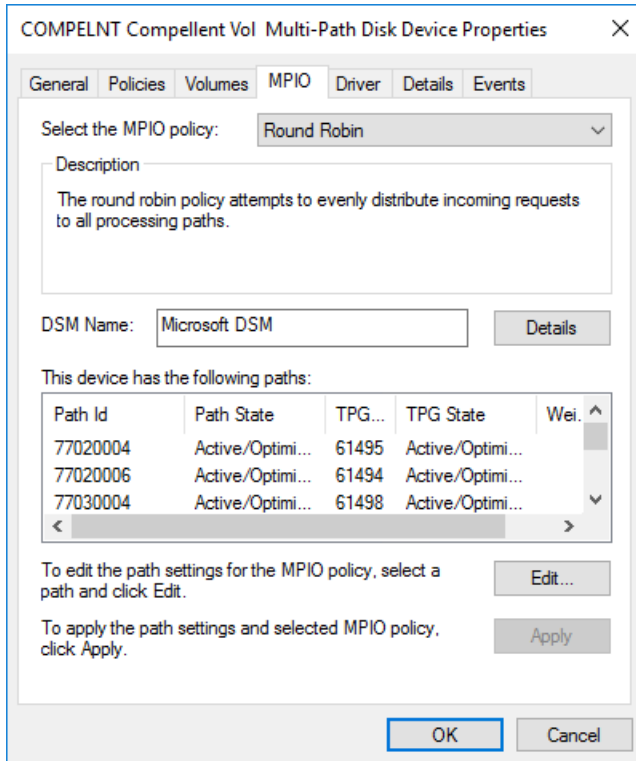


Figure 6 MPIO policy for all volumes

7 SQL Server 2017 Enterprise Edition configuration

7.1 Grant perform volume maintenance task privilege

During installation of SQL Server 2017, the option to grant the **SQL Server Database Engine Service** the **Perform Volume Maintenance Task** privilege was selected.

7.2 SQL Server maximum memory

The maximum server memory for this reference architecture should be set to 3040GB which leaves 32GB for the operating system. If additional applications share the server, adjust the amount of memory left available to the operating system accordingly.

7.3 Max degree of parallelism (MAXDOP)

During rowstore tests, MAXDOP=16 was used for I/O testing and MAXDOP=72 was used for CPU testing. For columnstore tests, MAXDOP=96 was used.

For more information, see the Microsoft article [Configure the max degree of parallelism Server Configuration Option](#).

7.4 Resource governor

The resource governor was used to limit the maximum memory grant. During rowstore I/O tests, resource governor was set to 5, and for all other tests it was set to 12.

For information about the resource governor, see the Microsoft article [Resource Governor](#).

7.5 Database configuration

The data warehouse database was configured to use multiple file groups, each containing four files distributed evenly across the four data volumes. All files were allowed to grow automatically. The file groups were configured with the **AUTOGROW_ALL_FILES** option to help ensure that all files within a given file group remain the same size.

7.6 Tempdb configuration



The tempdb database was configured to use eight data files of equal size. The data files were evenly distributed across the two tempdb data volumes, with four files stored on each volume. The tempdb transaction log file was placed on the log volume. All files were expanded to the appropriate size and auto grow was enabled.

8 Additional considerations for the highly available reference architecture

The highly available (HA) reference architecture leverages Windows Failover Clustering to achieve high availability. When configuring a Windows failover cluster, there are additional storage considerations:

- The recommended quorum configuration is to allow all cluster nodes to have quorum votes and use a disk witness. An additional volume needs to be created and configured as the disk witness. Dell EMC recommends using a 2GB volume for the disk witness. For more information on quorum and voting configurations in a failover cluster, see the Microsoft article [Configure and manage quorum](#).
- All volumes need to be mapped to each node of the cluster. It is recommended to use a cluster server object in the SC Series array when mapping volumes to the cluster.
- All volumes need to be configured as a cluster resource and added to the SQL Server cluster resource group.

9 DWFT for SQL Server 2017 certification

DWFT Certification #2017-020	Dell PowerEdge R940 with Dell EMC SC9000 DWFT Reference Architecture			Report Date: 10/23/2018	
DWFT Rev. 5.4					
System Provider	System Name	Processor Type		Memory	
	Dell PowerEdge R940	Intel Xeon Platinum 8168 2.7 GHz (4S/24C/48T)		3072 GB	
Operating System			SQL Server Edition		
Windows Server 2016			SQL Server 2017 Enterprise Edition		
Storage Provider	Storage Information				
	Dell EMC SC9000 48 x 3.84TB SSD for data, log and tempdb RAID 10 for log and tempdb, RAID 5-5 for data				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
200	582	1,118	584		
Row Store					
Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
582	597	15,854	18,415	17,134	70
Column Store					
Relative Throughput ²	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
1,118	7,264	4,830	N/A	N/A	92
<p>The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology</p> <p>¹ Assumes a data compression ratio of 5:1</p> <p>² Percent ratio of the throughput to the row store throughput of the reference configuration.</p> <p>³ Percent ratio of the throughput to the column store throughput of the reference configuration.</p> <p>* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

10 Summary

Dell EMC, in partnership with Microsoft, enables customers to deploy tested and validated data warehouse solutions using Data Warehouse Fast Track reference architectures for SQL Server 2017. These uniquely designed architectures ensure optimal business intelligence solutions. The end-to-end best practices and recommendations enable the customer to achieve enhanced return on investment and faster time to value with a balanced data warehouse environment that can perform better than traditional data warehouse systems.

Dell EMC DWFT reference architectures provide the following benefits:

- Deliver a tested and validated configuration with proven methodology and performance behavior
- Deliver outstanding performance on the PowerEdge 14G server platform with blazing processor speeds and leading-edge, flash-based Dell EMC storage arrays
- Achieve a balanced and optimized solution at all levels of the stack by following best practices for both hardware and software components, achieving faster time to value, and lower total cost of ownership
- Avoid over provisioning of hardware resources
- Offer high availability at all levels of setup (host, switches, and storage)
- Offer a single point of contact and accountability for purchases, services, and support; SQL Server is available to purchase from Dell EMC worldwide
- Help customers avoid the consequences of an improperly designed and configured system
- Reduce future support costs by limiting solution rearchitect efforts due to scalability challenges

This paper describes a reference architecture using a PowerEdge R940 server with an SC9000 storage array. By implementing Data Warehouse Fast Track for SQL Server 2017 design principles, this configuration achieved a 200TB rating.

A Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage Solutions Technical Documents](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

Additional resources:

- [Dell SQL Server Solutions](#)
- [Dell Data Warehouse Fast Track for SQL Server Advisor](#)
- [SC Series technical documents](#)