

# DATA WAREHOUSE FAST TRACK FOR MICROSOFT SQL SERVER 2014

EMC VNX 5600 Storage Array, Intel Xeon Processors,  
HP Proliant DL580 Server

EMC Solutions

Sep 2014



Copyright © 2014 EMC Corporation. All Rights Reserved.

Published Oct 2104

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

The information in this publication is provided as is. EMC Corporation 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.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

Part Number H13566

# Table of contents

<b>Overview .....</b>	<b>5</b>
Document purpose .....	5
Introduction to the VNX 5600 unified storage platform.....	5
Software suites available .....	5
Software packs available.....	5
Solution purpose .....	5
The business challenge.....	6
The technology solution .....	6
<b>Solution architecture.....</b>	<b>7</b>
Architecture diagram.....	7
Hardware resources .....	8
Software resources .....	8
<b>Key components.....</b>	<b>9</b>
Introduction .....	9
EMC VNX 5600 storage system.....	9
Microsoft Data Warehouse Fast Track program.....	9
HP Proliant DL580 G7 servers.....	9
<b>EMC technology overview.....</b>	<b>10</b>
Introduction .....	10
EMC Unisphere .....	10
EMC FAST Cache.....	11
<b>Microsoft Data Warehouse Fast Track overview .....</b>	<b>12</b>
Introduction .....	12
Value Proposition.....	13
SQL 2014 Fast Track .....	13
<b>The EMC-HP Data Warehouse Fast Track hardware and software.....</b>	<b>14</b>
<b>EMC VNX 5600 storage design layout.....</b>	<b>16</b>
<b>Database configuration .....</b>	<b>17</b>
<b>SQL Server settings.....</b>	<b>18</b>
Startup options .....	18
SQL Server memory settings.....	18
Hyper-threading.....	19
SQL server MAXDOP setting (max degree of parallelism) .....	19

Resource Governor.....	20
<b>Upgrade to DWFT for SQL2014 reference architecture .....</b>	<b>21</b>
<b>SQL 2014 DWFT metrics .....</b>	<b>21</b>
<b>EMC SQL 2014 reference architecture certification .....</b>	<b>21</b>
<b>Conclusion .....</b>	<b>22</b>
<b>Bills of materials .....</b>	<b>23</b>
<b>References.....</b>	<b>24</b>
EMC documentation.....	24
Microsoft documentation .....	24

## Overview

**Document purpose** The document describes the reference architecture of Data Warehouse Fast Track for Microsoft SQL Server 2014 using HP Proliant DL580 servers and an EMC® VNX5600™ unified storage system. The result is a balanced and optimized hardware and software configuration designed for SQL Server Data Warehouse deployments.

**Introduction to the VNX 5600 unified storage platform** The VNX5600 delivers industry-leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

### Software suites available

- **FAST™ Suite**—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (not available for the EMC VNXe series or the EMC VNX5100™).
- **Local Protection Suite**—Practices safe data protection and repurposing (not applicable to the EMC VNXe3100™ as this functionality is provided at no additional cost as part of the base software).
- **Remote Protection Suite**—Protects data against localized failures, outages, and disasters.
- **Application Protection Suite**—Automates application copies and proves compliance.
- **Security and Compliance Suite**—Keeps data safe from changes, deletions, and malicious activity.

### Software packs available

- **Total Efficiency Pack**—Includes all five software suites.
- **Total Protection Pack**—Includes local, remote, and application protection suites.
- **Total Value Pack**—Includes all three protection software suites and the Security and Compliance Suite.

**Solution purpose** The purpose of this reference architecture is to build and demonstrate the functionality, performance, and scalability of Data Warehouse Fast Track for Microsoft SQL Server 2014, enabled by the VNX 5600, Intel Xeon processors, and HP Proliant DL580 server.

This reference architecture validates the performance of the solution and provides guidelines to build similar solutions.

This document is not a comprehensive guide to every aspect of this solution.

## The business challenge

Customers require a scalable, tiered, and highly available infrastructure on which to deploy their data warehouse environment. Data warehouses provide the foundation for business intelligence systems. Business Intelligence has been a high growth application and will continue to grow dramatically as organizations realize the importance of mining their accumulated data for key business insights. Enterprise data warehouses have been growing rapidly. The massive data growth is often accompanied by increased complexity, but the data warehouse configuration typically suffers from design principles that were originally intended for online transaction processing systems (OLTP).

The reference configurations presented in this document address the common business challenges that data warehouses face today, and helps customers design and implement balanced configurations specifically for Microsoft SQL server Data Warehouse databases with the goal of providing a hardware-balanced approach and predictable out-of-box performance.

## The technology solution

This solution demonstrates how to use a VNX platform to provide the storage resources for a robust Microsoft data warehouse environment.

Data Warehouse Fast Track for Microsoft SQL Server 2014 provides a framework that allows customers to select a reference architecture that has been designed, tested and validated by hardware vendors. It is designed to let organizations quickly and reliably deploy total system solution stack, including server, storage, and connectivity for data warehouses based on Microsoft SQL Server 2014. This approach helps to dramatically reduce the burden on end users in selecting, sizing, and testing the combination of products and technologies while minimize the risk of failures

# Solution architecture

## Architecture diagram

This solution describes the Data Warehouse Fast Track for Microsoft SQL Server 2014 reference architecture using HP Proliant DL580 servers and the VNX5600 unified storage system.

Figure 1 depicts the physical infrastructure of the solution.

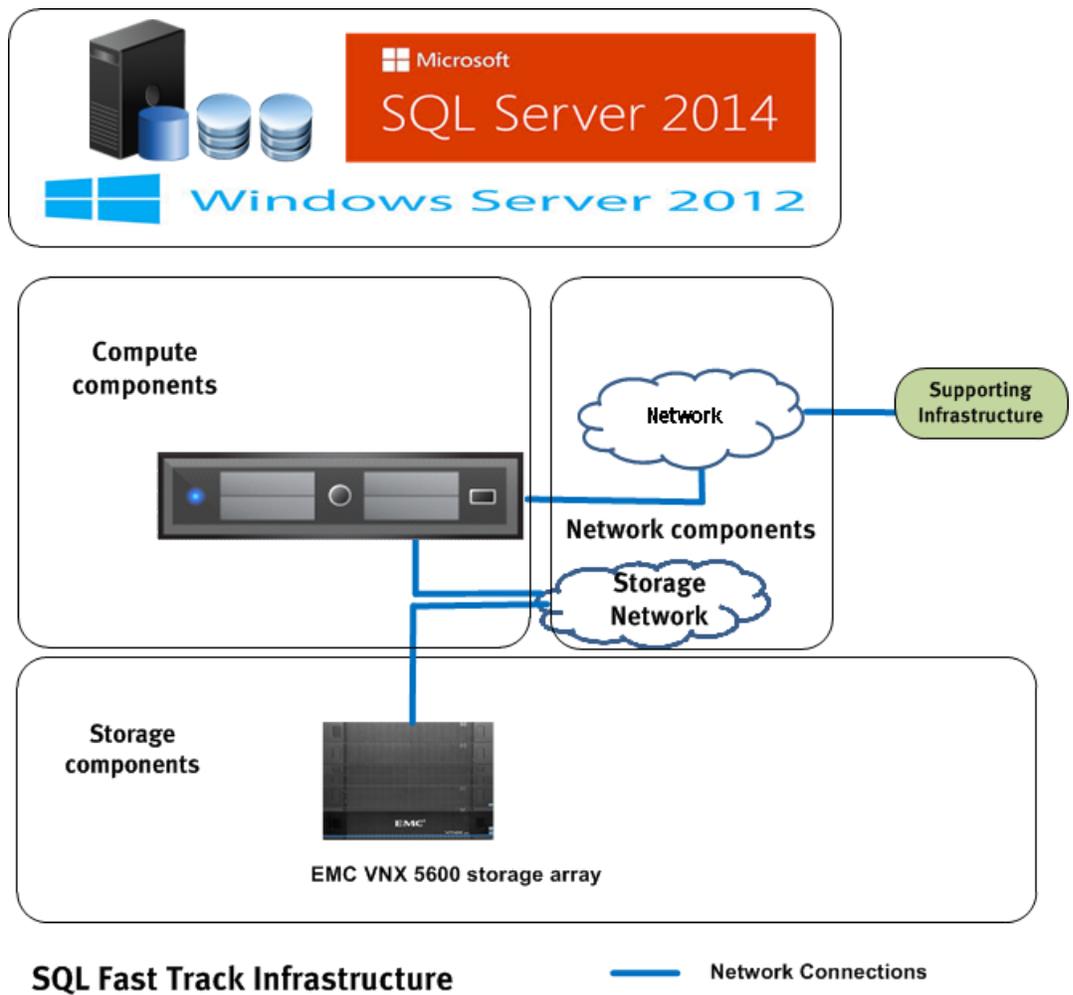


Figure 1. Solution architecture

## Hardware resources

Table 1 lists the hardware used in this solution.

**Table 1. Solution hardware**

Hardware	Quantity
EMC VNX5600 storage array	1
HP DL580R07 (E7) server	1
Emulex LPe12002 Fibre Channel Controllers	4
Fibre Channel Multi-Mode Cables with SF connectors	4

**Software resources** Table 2 lists the software used in SQL Server 2014 solution.

**Table 2. Solution software**

Software	Configuration
Microsoft Windows	Microsoft Windows Server 2012 R2 Standard Edition
Microsoft SQL Server	Microsoft SQL Server 2014 Enterprise Edition
Fast Track testing	Microsoft Reference Point Service Version 5
EMC PowerPath	EMC PowerPath version 5.5 SP1
EMC Unisphere™	Management tool for EMC VNX storage series
EMC VNX5600 VNX OE	VNX OE release 31, OS for block

## Key components

### Introduction

This section briefly describes the key components used in this solution, including:

- VNX 5600 storage system
- Data Warehouse Fast Track for Microsoft SQL Server 2014
- HP Proliant DL580 G7 servers

### EMC VNX 5600 storage system

The new VNX 5600 series storage is designed for mid-tier storage, supports block, file and unified configurations and uses a disk processor enclosure (DPE) chassis. The new MCx architectures enable the array to fully take advantage of the new Intel Xeon E5 multicore CPU architecture and ensure optimum performance at high scale. It allows the storage system to provide flash-drive class performance to data with a high locality of reference, which increases IOPs without placing all of the data onto flash drives. Multicore Fully Automated Storage Tiering (FAST) Cache reduces the load on the hard disk drive (HDD) by absorbing I/O bursts from applications and helps to improve the total cost of ownership (TCO) of the storage solutions.

With industry-leading performance, VNX 5600 storage systems offer advanced availability and flexibility, deliver scalable performance with MCx multicore optimization, provide five-nines reliability, and deliver built-in features for supporting replication, disaster protection, and automatic storage tiering.

### Microsoft Data Warehouse Fast Track program

Optimized for common data warehouse, decision support, and business analytics deployments based on Microsoft SQL Server, Data Warehouse Fast Track for Microsoft SQL Server 2014 program allows organizations to deploy tested configurations using proven technologies that offer balanced solution stacks. By following these reference configurations, new projects can be implemented quickly and confidently, avoiding much of the guesswork, planning choices, and difficult decisions involved with traditional approaches.

### HP Proliant DL580 G7 servers

Powered by Intel Xeon Processor E7-4870 processors, the HP Proliant DL580 G7 Server offers the unprecedented reliability, scalability, flexibility, and performance needed for complex IT environments, in a cost-effective and energy-efficient manner.

## EMC technology overview

### Introduction

This section identifies and briefly describes the major VNX features used in this solution environment, including:

- EMC Unisphere
- EMC FAST Cache
- Block Data Compression

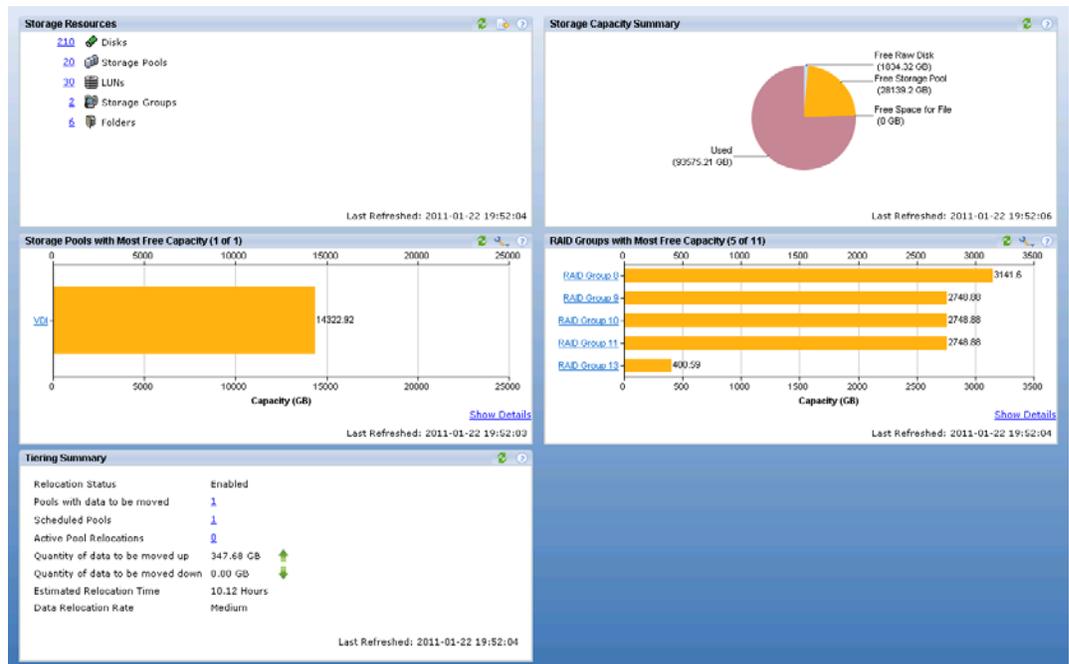
### EMC Unisphere

EMC Unisphere provides a flexible, integrated experience for managing CLARiiON, Celerra, and VNX platforms in a single view. This new approach to midtier storage management fosters simplicity, flexibility, and automation. Unisphere's unprecedented ease of use is reflected in intuitive task-based controls, customizable dashboards, and single-click access to real-time support tools and online customer communities.

Unisphere features include:

- Task-based navigation and controls that offer an intuitive, context-based approach to configuring storage, creating replicas, monitoring the environment, managing host connections, and accessing the Unisphere support ecosystem.
- A self-service Unisphere support ecosystem, accessible with one click from Unisphere, that provides users with quick access to real-time support tools, including live chat support, software downloads, product documentation, best practices, FAQs, online communities, ordering spares, and submitting service requests.
- Customizable dashboard views and reporting capabilities that enable at-a-glance management by automatically presenting users with valuable information in terms of how they manage their storage. For example, customers can develop custom reports up to 18 times faster with EMC Unisphere.
- Common management provides a single sign-on and integrated experience for managing both block and file features.

Figure 2 provides an example of the Unisphere Summary page that gives administrators a wealth of detailed information on connected storage systems, from LUN pool and tiering summaries to physical capacity and RAID group information.



**Figure 2. Unisphere Summary page**

## EMC FAST Cache

VNX FAST Cache, a part of the VNX FAST suite, enables flash drives to be used as an expanded cache layer for the array. FAST Cache has array-wide features available for both file and block storage. FAST Cache works by examining 64 KB chunks of data in FAST Cache enabled objects on the array. Frequently accessed data is copied to the FAST Cache and subsequent accesses to that data chunk are serviced by FAST Cache. This allows immediate promotion of very active data to the flash drives. This dramatically improves the response time for very active data and reduces the data hot spots that can occur within the LUN.

FAST Cache is an extended read/write cache that can absorb read-heavy activities such as boot storms and antivirus scans, and write-heavy workloads such as operating system patches and application updates.

# Microsoft Data Warehouse Fast Track overview

## Introduction

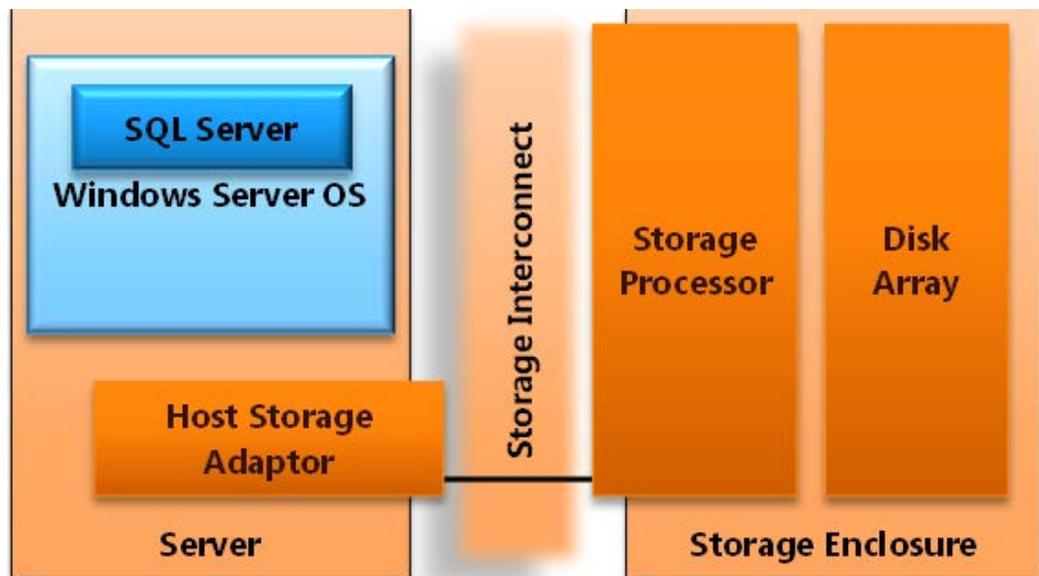
Optimized for common data warehouse, decision support, and business analytics deployments based on Microsoft SQL Server, Data Warehouse Fast Track for Microsoft SQL Server 2014 program allows organizations to deploy tested configurations using proven technologies that offer balanced solution stacks. By following these reference configurations, new projects can be implemented quickly and confidently, avoiding much of the guesswork, planning choices, and difficult decisions involved with traditional approaches. Goals of the program include:

- Accelerated data warehouse projects with pre-tested and validated hardware and software configurations
- Reduced hardware and maintenance costs as a result of purchasing the right balanced hardware solution that is optimized for data warehouse workloads
- Reduced planning and setup costs and risk through choosing the right scale model, configuring the system correctly, and taking advantage of tuning guidance.

Data Warehouse Fast Track reference configurations are successful because of their specific focus on sequential I/O, which are critical to data warehouse performance, as well a balanced approach to hardware configuration. Most system architectures consider the memory, I/O bandwidth, and storage capacity separately when planning a system. Instead, Data Warehouse Fast Track reference architecture balances the overall system so that throughput is matched across the various hardware and software components, specifically for data warehouse performance. Key considerations for balanced data warehousing application performance include:

- Compute power in terms of available CPU cores
- Main memory for the use of the Microsoft SQL Server cache
- I/O channel throughput between the server and the storage system, including the combined throughput of the Host Bus Adapters (HBA) and switching hardware
- Storage system throughput for sequential reads, caching capabilities, data storage capacity, and redundancy

The goal of a Data Warehouse Fast Track reference configuration is to achieve a cost-effective balance between SQL Server data processing capabilities and the capabilities of the underlying hardware components. Figure 3 shows the component architecture.



**Figure 3. SQL Server Data Warehouse Fast Track component architecture**

### Value Proposition

Three principles constitute the value proposition for Data Warehouse Fast Track:

- Predetermined balance across key system components, which minimizes the risk for overspending for CPU or storage resources.
- Predictable out-of-the-box performance. Fast track configuration matches SQL server application capacity for a target workload
- Workload-centric. Fast track reference architecture aligned specifically with data warehouse workload, not a one-size-fits-all scenario.

### SQL 2014 Fast Track

Data Warehouse Fast Track for SQL 2014 program fully takes advantage of the Clustered Column store index (CCI) technology in SQL Server 2014. CCI store data in columnar storage, enabling 10X+ data compression while speeding up the query performance with batch-mode processing.

- CCI provides superior data compression, reduce storage and IO bandwidth requirement as only columns referenced in the query need to be brought into memory.
- CCI also reduces the need for secondary indexes, simplifies the overhead on incremental data load, and allows customer to load data in the maintenance window.
- CCI execute adhoc complex analytic queries efficiently because it is designed to use predicate pushdown and partition elimination to reduce the data to be processed in large sets of data.

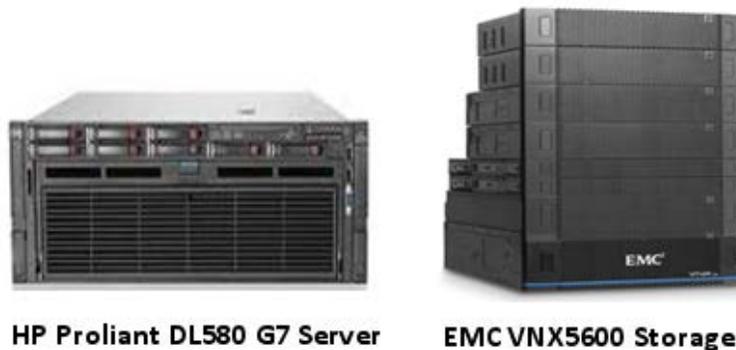
## The EMC-HP Data Warehouse Fast Track hardware and software

The EMC VNX 5600 Data Warehouse Fast Track reference architecture consists of an HP Proliant DL580 G7 server and one EMC VNX5600 storage array. The HP Proliant DL 580 G7 server is based on Intel Xeon Processor 4870 CPUs. Intel Xeon processors at the heart of the system deliver outstanding scalability and performance to support the new business intelligence capabilities available in Microsoft SQL Server 2014. These advancements help enable large-scale consolidation efforts while serving demanding data-intensive workloads.

New features specific to Intel Xeon processor E7 series, such as multi-bit error correction, provide advanced reliability, availability, and serviceability for enterprise database applications. In addition, dramatic TCO reductions are possible with Intel Xeon processors, helping organizations realize cost reductions with automated energy efficiency and consolidation.

The new VNX 5600 includes a 2.4 GHZ, four-core Xeon E5 processor with 24 GB of RAM, PCIe GEN 3 IO modules and a 6 GB/s SAS backend infrastructure with up to eight 6 GB/s high-bandwidth connections to 60 drive DAE and supports up to 500 drives with FX, iSCSI, and FCoE. The VNX5600 offers performance, capacity, protection, and the best TCO to meet diversified business requirements. Powerful multicore optimization with MCx provides scaling performance for mixed workloads. The boosted sustained read-and-write bandwidth it offers is typically crucial for data warehouse deployment.

Figure 4 shows the HP Proliant DL580 G7 Server and the VNX5600 storage.



**Figure 4. HP Proliant DL580 G7 Server and EMC VNX5600 storage**

Table 3 lists the hardware configuration for the solution. Figure 3

**Table 3. Hardware configuration for EMC-HP reference architecture**

Components		Quantity/Description
<b>Server</b>	<b>HP Proliant DL 580 G7 Server</b>	
	CPU	2 x 10 Core Intel Xeon Processors Model X4870 (2.40 GHz, 30 Mb cache, 130 W)
	Number of cores	20

Components		Quantity/Description
	PCI-E slots	2 x 8 PCIe3.0 2 x 16 PCIe3.0
	Internal drives	2 x 146 GB 10k RPM SAS disk drives
	Storage controller	1 x Smart Array P410/1 GB flash-based write cache (FBWC)
	Network Adapters	1 x 1 GbE NC3751 multifunction 4 ports
	RAM	128 GB PC2-5300 Fully Buffered DIMMs (FBD)
	Memory Slot	32 GB DDR3 dual in-line memory modules (DIMMs) with 64 DIMM slots
<b>Storage</b>	<b>EMC VNX 5600</b>	
	Host bus adapters	4 x Emulex LPe12002-E
	Disk drives	75 x 2.5" 300 GB 10k-RPM SAS disk drives
	Data protection	<ul style="list-style-type: none"> <li>• RAID 5 for data and TempDB</li> <li>• RAID 1/0 for log</li> </ul>
<b>Connectivity</b>		
	Protocol	Fibre Channel (FC)

**Table 4. Software configuration for EMC–HP reference architecture**

Software	Configuration
Windows	Microsoft Windows Server 2012 R2 Standard Edition
SQL Server	Microsoft SQL Server 2014 Enterprise Edition
Fast Track testing	Microsoft Reference Point Service Version 5.0
EMC PowerPath	EMC PowerPath version 5.5 SP1

## EMC VNX 5600 storage design layout

The VNX5600 storage system uses a DPE available in a 25 x 2.5 drive form factor. Additional disk drives can be connected using integral disk array enclosures (DAEs) that each provide 25 additional drive slots. The VNX5600 storage platform provides a robust, scalable, and simplified storage infrastructure.

Disk drives within the VNX5600 storage system are configured according to how they are used within the data warehouse system. At least 70 disks are required for this configuration:

- **Storage System**—The first four drives in enclosure 0 bus 0 house the VNX array’s operating system, boot image, and file control LUNs.
- **Transaction Log**—Four drives in a 2+2 RAID 10 group are used for the transaction log.
- **Hot Spares**—Any unassigned drive in the system can operate as a hot spare. The new VNX series does not require disks to be defined as permanent hot spares. When a drive is used in a sparing function, the particular drive becomes a permanent member of the RAID group without the need to re-balance and equalize the spare drive to a new drive. This reduces the exposure and performance overhead of sparing operations.
- **Staging Area**—Two drives are used for staging area such as data landing.
- **Data LUNS**—60 drives are used for database and tempDB files. These 60 drives are organized into 4+1 RAID 5 sets, known as RAID groups.

The RAID 5 configuration consists of four data drives protected by one parity drive. In a RAID 5 configuration, each 256 KB chunk of database data stored on one of the data LUNs is striped in 64 KB increments across four of the five drives in each RAID set. On the fifth drive, 64 KB of parity information is computed and stored based on the data on the other four drives.

RAID 5 was chosen as the storage deployment layout because:

- RAID 5 provides high read performance because all drives contribute to the delivery of data.
- RAID 5 provides lower cost data protection compared to RAID 10, requiring only 20 percent of the total usable capacity for storing parity data.

## Database configuration

The testing for this reference architecture was performed using a 1 TB TPC-H database and a sequential read workload stimulation tool called Reference Point Service. The performance number depicted is output from the simulation tool running both a simple and a complex workload. The test database uses a master file group and additional file groups which represent the 7 partitions, as shown in Table 5.

**Table 5. Test database file groups**

File Group	Files	Notes
Primary	1	Contains master DB file(MDF)
Part_ci1FG	12	1 file on each LUN
Part_ci2FG	12	2 files on each LUN
Part_ci3FG	12	3 files on each LUN
Part_ci4FG	12	4 files on each LUN
Part_ci5FG	12	5 files on each LUN
Part_ci6FG	12	6 files on each LUN
Part_ci7FG	12	7 files on each LUN

Select a page

- General
- Files
- Filegroups
- Options
- Change Tracking
- Permissions
- Extended Properties
- Mirroring
- Transaction Log Shipping

Script Help

Rows

Name	Files	Read-Only	Default
PRIMARY	1	<input type="checkbox"/>	<input type="checkbox"/>
Base	8	<input type="checkbox"/>	<input checked="" type="checkbox"/>
part_ci1FG	8	<input type="checkbox"/>	<input type="checkbox"/>
part_ci2FG	8	<input type="checkbox"/>	<input type="checkbox"/>
part_ci3FG	8	<input type="checkbox"/>	<input type="checkbox"/>
part_ci4FG	8	<input type="checkbox"/>	<input type="checkbox"/>

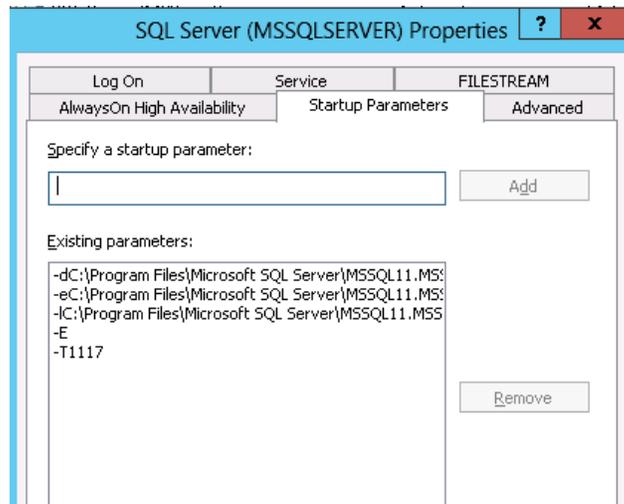
## SQL Server settings

### Startup options

You must configure the startup options as follows:

- Add E to the startup options. This increases the number of contiguous extents in each file that are allocated to a database table as it grows and improves sequential disk access.
- Add T1117 to the start-up options. The trace flag ensure even growth of all files in a file group in the case that auto-grow is enabled.

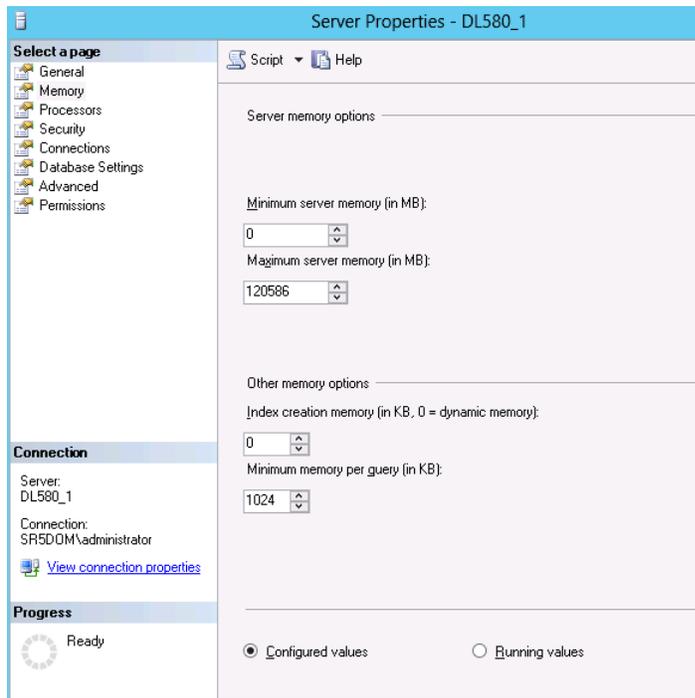
Figure 5 shows the Startup Parameters dialog box.



**Figure 5. SQL Server startup parameters set up**

### SQL Server memory settings

According to SQL server 2014 best practices, you should allocate no more than 92 percent of total available RAM to SQL server. In this configuration, the SQL server memory setting is 117 GB (120,586 MB), as shown in Figure 6.



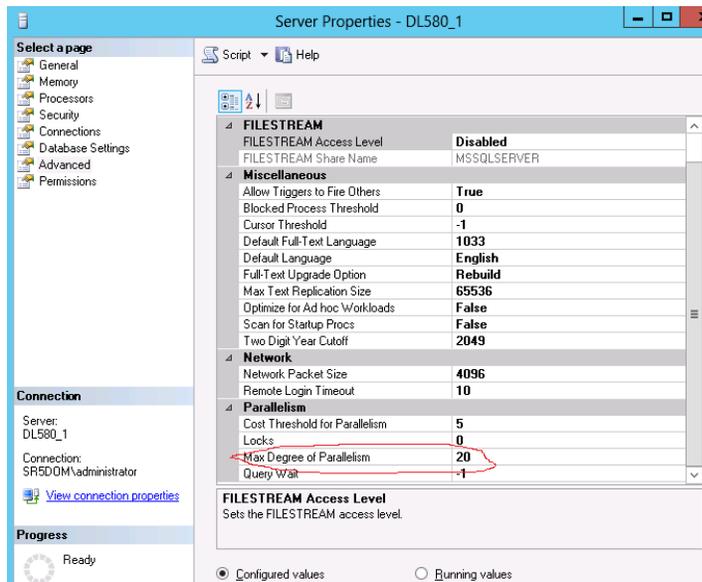
**Figure 6. SQL Server Memory Settings in MB**

## Hyper-threading

The processors in the HP DL580 server use Intel Hyper-threading technology, which allows the server to use more logical CPU cores than what is physically available in the system. When higher intensity workloads use a large number of parallel queries, the hyper-threading technology allows SQL Server to over-subscribe the amount of CPU and memory resources given to queries during run time. For example, in this configuration, the HP DL580 server has 20 physical CPU cores, but SQL Server could allocate up to 40 logical CPU cores to queries after hyper-threading is turned on.

## SQL server MAXDOP setting (max degree of parallelism)

MAXDOP controls the number of logical cores that SQL Server uses for the parallel execution of a query. With hyper-threading technology, SQL Server could dynamically allocate SQL threads. Figure 7 shows the maximum degree of parallelism (DOP) parameter setting in SQL Server.

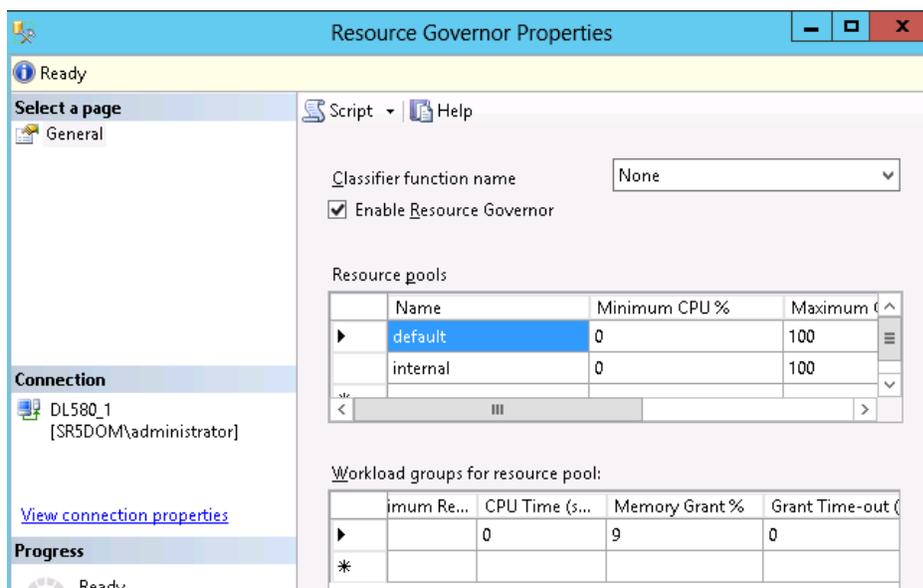


**Figure 7. SQL Server MAXDOP set up**

**Resource Governor** You can use the Resource Governor to reduce the maximum memory consumed by each query. You can set the constraint options to provide the balance you need:

- Lower constraints provide higher performance for individual queries.
- Increased constraints guarantee the number of queries that can run concurrently.

The default setting for SQL Server 2014 is 25 percent of SQL server memory resources for each session. In this configuration, the Resource Governor memory grant was set at 9 percent, as shown in Figure 8.



**Figure 8. SQL Server Resource Governor set up**

## Upgrade for DWFT for SQL Server 2014 reference architecture

To upgrade from Data Warehouse Fast Track 4.0 to Data Warehouse Fast Track for SQL Server 2014, no new hardware is required. Simply upgrade the software to SQL Server 2014. The throughput remains the same.

### SQL 2014 FTDW metrics

SQL 2014 FTDW certification uses the four (4) primary metrics for capacity evaluation process.

- **Rated user data capacity:** it is a calculated value based on the Row Store Relative Throughput, the Column Store Relative Throughput, available storage and the physical system memory. The calculation based on a compression ratio of 5:1
- **Maximum user data capacity:** it is a calculated value based on the total disk capacity of all disks allocated to primary data storage and assume a compression ratio of 5:1.
- **Row Store Relative Throughput:** it is calculated as a percent ratio of the Row Store throughput to the Row Store throughput of the validated sample FTDW reference configuration (A 2 socket system which has 25TB of rated user data capacity)
- **Column Store Relative Throughput:** it is a calculated value as a percent ratio of the column Store throughput to the Column Store throughput of the FTDW RA)

To upgrade from Data Warehouse Fast Track 4.0 to 5.0, no new hardware is required. Simply upgrade the software to SQL Server 2014. The throughput remains the same.

## EMC DWFT for SQL 2014 reference architecture certification

DWFT Certification #2014-005	<b>EMC VNX5600 SQL</b>			Report Date: 7/31/2014	
FTDW Rev. 5.4	<b>DWFT Reference Architecture</b>				
<b>System Provider</b>	<b>System Name</b>	<b>Processor Type</b>	<b>Memory</b>		
	EMC VNX SQL DWFT 28TB Reference Architecture	Intel Xeon E5-4870 2.2 GHz (2/20/40)	256 GB		
<b>Operating System</b>		<b>SQL Server Edition</b>			
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition			
<b>Storage Provider</b>	<b>Storage Information</b>				
	48x 300 GB SAS for data and tempdb (RAID 5) 4x 300 GB SAS for Staging (RAID 5) 2x 300 GB SAS for log (RAID 10)				
<b>Primary Metrics</b>					
Rated User Data Capacity <sup>1</sup>  (TB)	Row Store Relative Throughput <sup>2</sup>	Column Store Relative Throughput <sup>3</sup>	Maximum User Data Capacity <sup>1</sup>  (TB)		
28	92	137	44		
<b>Row Store</b>					
Relative Throughput <sup>2</sup>	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.)  (%)
92	90	2,189	3,204	2,697	88
<b>Column Store</b>					
Relative Throughput <sup>2</sup>	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.)  (%)
137	889	466	N/A	N/A	94
<p>The reference configuration is a 2 socket system rated for 25TB using the DWFT V4 methodology</p> <p><sup>1</sup> Assumes a data compression ratio of 5:1</p> <p><sup>2</sup> Percent ratio of the throughput to the row store throughput of the reference configuration.</p> <p><sup>3</sup> Percent ratio of the throughput to the column store throughput of the reference configuration.</p> <p><sup>4</sup> Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

## Conclusion

The Microsoft SQL Server Fast Track architecture presents a balanced configuration of processor core and disk performance capabilities that are optimized for sequential workloads. The reference configuration presented here is designed, tested, and validated by EMC and Microsoft, and offers simplicity, reliability, scalability, and low TCO for organizations deploying an infrastructure for a data warehouse with predictable performance. Contact EMC or Microsoft today to learn if this offering is right for you.

## Bills of materials

The following tables list the bills of materials for the server, storage, and switch components of the reference architecture.

**Table 6. Server components**

Component	Part number	Quantity
HP DL580R07 (E7) CTO Chassis	643086-B21	1
HP E7-4870 DL580 G7 2P Factory installable option (FIO) Kit	643067-L21	1
HP 16GB 2Rx4 PC3L-10600R-9 Kit	627812-B21	16
HP Slim 12.7mm serial advanced technology attachment (SATA) DVD Optical Kit	481041-B21	1
HP 146GB 6G SAS 15K 2.5in DP Enterprise Hard Disk Drive	512547-B21	2
HP DL580G7 PCI Express Kit	588137-B21	1
HP 1G Flash Backed Cache	534562-B21	1
HP 1.83m 10A C13-UL US Pwr Cord	AF556A	4
HP 1200W Common Slot Silver Hot Plug Power Supply Kit	500172-B21	4
HP PL Foundation Pk Single Rel FIO Software	534516-B21	1
HP IC ML/DL/BL Bundle E-LTU 24x7 SW	TC278AAE	1

**Table 7. Connectivity components**

Component	Model number	Quantity
Emulex LPe12002 Fibre Channel Controllers	LPe12002-E	4
Fibre Channel Multi-Mode cables w/ SF connectors	N/A	8

**Table 8. Storage components**

Component	Model number	Quantity
VNX5600 (Including 4 FLARE drives)	VNX56D253010F	1
25*2.5 drive Disk Array Enclosure (DAE)	V2-DAE-R-25-A	3
300 GB 2.5" 10K SAS drive	VX-2S10-300	75

## References

### EMC documentation

The following white papers, located on [emc.com](http://emc.com), provide more information:

- [\*Introduction to the New VNX Series: VNX 5200, VNX 5400, VNX 5600, VNX 5800, VNX 7600, and VNX 8000—A Detailed Review\*](#)
- [\*EMC Unisphere: Unified Storage Management Solution for the New VNX Series\*](#)

### Microsoft documentation

The following links provide more information about the Microsoft SQL Server Data Warehouse Fast Track solution:

- [\*Microsoft SQL Server 2012 Fast Track Reference Architecture\*](#)
- Data Warehouse Fast Track Reference Guide for SQL Server 2014