



Technical Report

Reference Architecture for Microsoft SQL Server 2016 Data Warehouse Fast Track: NetApp EF560

Pat Sinthusan, NetApp
April 2017 | TR-4586

Abstract

This technical report describes a Microsoft SQL Server 2016 Data Warehouse Fast Track reference architecture that was built with a NetApp® EF560 flash array (2U) and a Fujitsu Server PRIMERGY RX2540 M1 (2U). This report also presents the configuration and the performance results of the 4U solution.

TABLE OF CONTENTS

1	Solution Overview	3
1.1	Target Audience.....	3
2	EF560 DWFT Reference Architecture Design and Implementation.....	3
2.1	NetApp EF560 Flash Array	3
2.2	Fujitsu Server PRIMERGY RX2540 M1	3
2.3	Microsoft SQL Server 2016 Enterprise	4
2.4	Database and Storage Layout	4
3	EF560 DWFT Reference Architecture Certification.....	6
4	Bill of Materials	7
5	Conclusion	8
	References.....	8

LIST OF TABLES

Table 1)	Bill of Materials for the server.....	7
Table 2)	Bill of Materials for the storage array.....	7

LIST OF FIGURES

Figure 1)	Database and storage layout details.....	5
Figure 2)	Sample EF560 DWFT reference architecture certification.....	6

1 Solution Overview

Data Warehouse Fast Track (DWFT) for SQL Server 2016 is a program that Microsoft administers to produce efficient, purpose-built, and out-of-the-box balanced reference configurations for SQL Server data warehouse workloads. The DWFT validation kit for SQL Server 2016 is the fifth generation of the Microsoft DWFT validation kit. This latest DWFT validation kit (DWFT V5.4.5345) is comprehensive. It contains all the components that users need to build the database, populate the data, and run the queries.

The new DWFT V5.4.5345 includes a certification template. Section 3 of this report shows the certification template with the performance data of the reference architecture that is described in this report.

The DWFT V5.4.5345 test takes advantage of the performance improvement and flexibility of the columnstore index (CCI). For more information about the SQL Server 2016 columnstore index, see [Columnstore indexes—what's new](#). The query performance is also improved by the addition of a new batch-mode query execution mechanism, which is closely integrated with and optimized around the column-store storage format.

This report was created to show NetApp EF-Series system interoperability with the data warehouse solution. No specific performance tuning was applied to this environment.

1.1 Target Audience

The target audience for the solution includes the following groups:

- Data warehouse architects
- Business intelligence (BI) users
- Database administrators
- IT planners

For more information about Microsoft SQL Server Data Warehouse Fast Track, go to [Data Warehouse Fast Track](#).

2 EF560 DWFT Reference Architecture Design and Implementation

The design of this DWFT reference architecture calls for a Fujitsu Server PRIMERGY RX2540 M1 that is directly attached to a NetApp EF560 flash array with four 12Gbps SAS links.

2.1 NetApp EF560 Flash Array

The NetApp EF560 is an all-flash storage array that brings together extreme performance and enterprise-grade reliability. The EF560 is a 2U dual controller storage array, and for this test, it was populated with twenty-four 800GB solid-state drives (SSDs). The dual controllers were configured with SAS host interface cards (HICs) and 24GB of controller cache. The EF560 runs on enterprise-proven NetApp SANtricity® 11.30 management software and controller firmware version 8.30.01.00.

For additional information about NetApp EF-Series solutions, visit [EF-Series Flash Storage Arrays](#).

2.2 Fujitsu Server PRIMERGY RX2540 M1

The Fujitsu PRIMERGY RX2540 M1 server was equipped with two of the latest Intel Xeon E5-2699 v3 @2.30GHz processors with up to 36 cores, along with new DDR4 memory technology of up to 1.5TB. The modular design of the server offers excellent expandability with up to 24 drives, high storage density, DynamicLoM technology, and up to 8 Peripheral Component Interconnect Express (PCIe) Gen 3 I/O expansion slots.

The tested server was equipped with 384GB of RAM and 2 internal SSDs for the operating system, and 2 of the 6 PCIe slots were occupied by LSI SAS 9300-8e host bus adapter (HBA) cards.

2.3 Microsoft SQL Server 2016 Enterprise

Microsoft SQL Server 2016 Enterprise was installed and was configured with the following settings.

For start-up options:

- `-E` was added to the start-up options to increase the contiguous extents in database files and to enable large transfer sizes in sequential disk access.
- `-T1117` was also added to the start-up options to facilitate even growth of all the files in a file group if autogrow is enabled.
- `-T834` was added to the start-up options to enable large pages. This trace flag was used for columnstore testing only.

SQL Server Maximum Memory

For ease of benchmarking, SQL Server 2016 maximum memory was set based on the number of sockets of the reference architecture implementation. Because the Fujitsu PRIMERGY RX2540 M1 server under testing had two sockets, the SQL Server maximum memory was set to 118GB.

MAXDOP

MAXDOP, or maximum degree of parallelism, is a tuning option that controls how many SQL Server schedulers a SQL Server query can use per query step. The following settings were used:

- MAXDOP = 32, for row store testing
- MAXDOP = 72 (that is, the system's logical processor count), for columnstore testing

Resource Governor

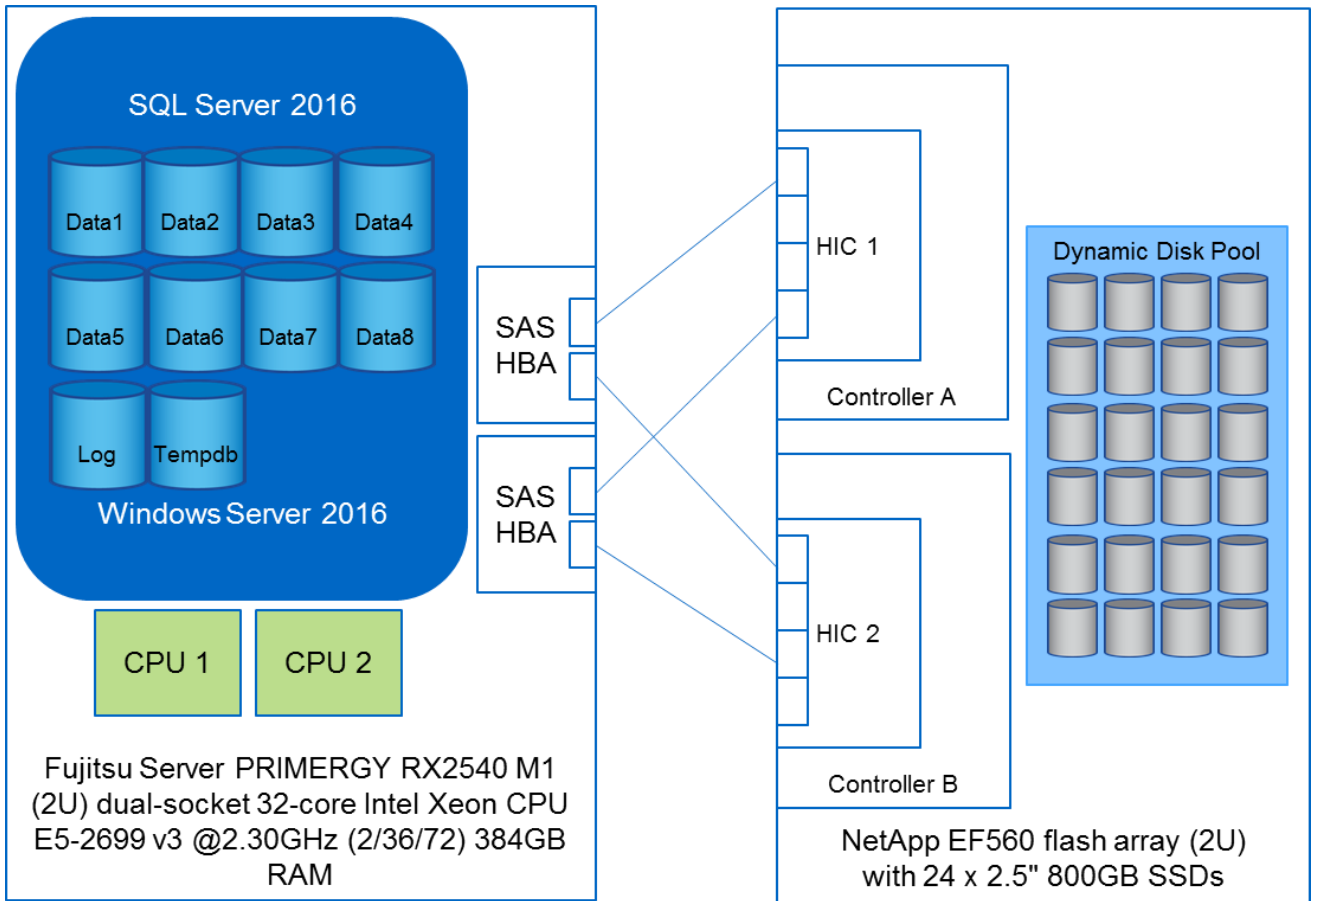
The SQL Server 2016 Resource Governor was enabled and was set to 12%.

2.4 Database and Storage Layout

The combination of NetApp SANtricity Dynamic Disk Pools (DDP) technology and NetApp E-Series storage is optimal for high-performance computing (HPC). DDP technology distributes data, parity information, and spare capacity across a pool of drives. To promote full data protection, its intelligent algorithm (seven patents pending) defines which drives are used for segment placement. DDP dynamic rebuild technology uses every drive in the pool to rebuild a failed drive, enabling exceptional performance under failure. Flexible disk-pool sizing optimizes utilization of any configuration for maximum performance, protection, and efficiency.

Therefore, a single Dynamic Disk Pool pool was created on the EF560 system by using 24 SSDs for SQL Server database data, log, and tempdb files. Figure 1 shows the detailed storage layout and database layout for the reference configuration.



Figure 1) Database and storage layout details.



3 EF560 DWFT Reference Architecture Certification

Figure 2 shows a DWFT reference architecture certification.

Figure 2) NetApp EF560 DWFT reference architecture certification.

DWFT Certification #2016-014	NetApp EF560 with Fujitsu RX 2540 M1 DWFT Reference Architecture			Report Date: 3/29/2017	
DWFT Rev. 5.4					
System Provider	System Name	Processor Type		Memory	
	PRIMERGY RX2540 M1	Intel® Xeon® CPU E5-2699 v3 @2.3GHz (2/36)		384 GB	
Operating System			SQL Server Edition		
Windows Server 2016 Datacenter			SQL Server 2016 Enterprise Edition		
Storage Provider	Storage Information				
	24x800 GB SSD for data, log and tempdb (Dynamic Disk Pool) 2x300 GB SSD for OS (RAID 1)				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
55	217	252	66		
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.) (%)
217	283	5,367	7,404	6,385	98
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.) (%)
252	1,640	3,079	N/A	N/A	85
The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and the DWFT V4 methodology					
¹ Assumes a data compression ratio of 5:1					
² Percent ratio of the throughput to the row store throughput of the reference configuration.					
³ Percent ratio of the throughput to the column store throughput of the reference configuration.					
* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.					

4 Bill of Materials

Table 1 and Table 2 list the Bill of Materials for this reference architecture, including server and storage components, respectively.

Table 1) Bill of Materials for the server.

Part Number	Description	Quantity
S26361-F3843-E516	16GB (1 x 16GB) 2Rx4 DDR4-2133 R ECC	16
S26361-F2495-E440	8 x 2.5" HDD bays	1
S26361-F3849-E100	Cooling kit 2nd CPU	1
S26361-F3778-E1	DVD-RW supermulti-ultraslim	1
S26361-F3629-E202	2-port NIC, Intel 82599 based	1
S26361-F3849-E170	Intel Xeon E5-2699 v3 18C/36T 2.30GHz	2
S26113-F574-E13	Modular SV 800W platinum hp	2
S26361-F4530-E10	Mounting of RMK in symmetrical racks	1
T25139-Y3850-E10	Option "no powercord." for countries without specific cable orderable, for example, China	2
S26361-F5313-E201	PFC EP QLE2670 LP	1
S26361-F5302-E201	PLAN EM 2 x 1Gb T interface card	1
S26361-F5243-E1	PRAID EP400i	1
S26361-F3694-E2	Performance mode installation	2
H3-25460-01	SAS9300-8e dual-port 12GB (8 Phy) PCIe SAS 3 adapt.	2

Table 2) Bill of Materials for the storage array.

Part Number	Description	Quantity
EF-X561204A-R6	EF560-A2R24	2
110-00391	12Gb SAS3 host interface card quad port	2
E-X5681A-R6	Empty 2U-24 DE5600 enclosure	1
112-00379	HDSAS to HDSAS 1M cable	4
PX02SMF080	800GB SSDs	24

Part Number	Description	Quantity
X-52197-00-R6	Power cord, in cabinet, 2m, C14-C13, E-Series	1

5 Conclusion

The NetApp reference architecture for Microsoft SQL Server 2016 DWFT that is presented in this report was designed to show the ease of configuring the EF-Series array by using Dynamic Disk Pools technology. It was also designed to provide a level of confidence in implementing the solution. The EF-Series all-flash array is built to deliver high performance with exceptionally low latency.

The array can deliver 825,000 IOPS and 12GBps throughput. In this test with a single 2U server, the 4U solution (server and storage) delivered 7404MBps row-store measured I/O throughput and 1,640 queries/hr/TB column-store measured throughput.

References

The following references were used in this TR:

- Columnstore indexes—what's new
<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-what-s-new>
- Data Warehouse Fast Track
<https://www.microsoft.com/en-us/cloud-platform/data-warehouse-fast-track>
- EF-Series Flash Storage Arrays
<http://www.netapp.com/us/products/storage-systems/all-flash-array/ef-series.aspx>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.