



Technical Report

Reference Architecture for Microsoft SQL Server 2014 Data Warehouse Fast Track:NetApp EF550

Wei Liu, NetApp

June 2015 | TR-4387

Abstract

This technical report describes a Microsoft SQL Server 2014 Data Warehouse Fast Track (DWFT) reference architecture that was built with a Fujitsu Primergy RX200 S8 Server (1U) and a NetApp® EF550 flash array (2U), and the configuration and performance results of the 3U solution.

TABLE OF CONTENTS

1	Solution Overview	3
1.1	Target Audience.....	3
2	EF550 DWFT Reference Architecture Design and Implementation	3
2.1	Fujitsu RX200 S8 Server.....	3
2.2	Microsoft SQL Server 2014 Enterprise	3
2.3	NetApp EF550 Flash Array	4
2.4	Database and Storage Layout.....	4
3	EF550 DWFT Reference Architecture Certification	6
4	Bill of Materials	7
5	Conclusion	7
	References.....	7

LIST OF TABLES

Table 1)	Bill of materials.....	7
----------	------------------------	---

LIST OF FIGURES

Figure 1)	Database and storage layout details.....	5
Figure 2)	EF550 DWFT reference architecture certification.....	6

1 Solution Overview

Data Warehouse Fast Track (DWFT) for SQL Server 2014 is a program administered by Microsoft to produce efficient, purpose-built, and out-of-box balanced reference configurations for SQL Server data warehouse workloads. In August 2014, Microsoft released the DWFT validation kit for SQL Server 2014. This kit is the fifth generation of the Microsoft DWFT validation kit. This latest DWFT validation kit, or DWFT V5, is comprehensive. It contains all of the components needed to build the database, populate the data, and run the queries.

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

The DWFT V5 test takes advantage of the recently introduced clustered column store index (CCI) feature introduced in SQL Server 2014. CCI enables 10x data compression, resulting in reduced I/O, smaller in-memory footprint, and significantly improved query performance. 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 the EF-Series platform 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
- DBAs
- IT planners

For more information about Microsoft SQL Server Data Warehouse Fast Track, go to <http://www.microsoft.com/en-us/server-cloud/data-warehouse-fast-track.aspx>.

2 EF550 DWFT Reference Architecture Design and Implementation

The design of this DWFT reference architecture calls for a Fujitsu RX200 S8 server directly attached to a NetApp EF550 flash array with four 6Gbps SAS links.

2.1 Fujitsu RX200 S8 Server

The Fujitsu Primergy RX200 S8 server is a one-rack (1RU) server powered by the Intel Xeon processor E5-2680 V2 product family, up to 1536GB of RAM, up to eight hot-plug internal HDDs or SSDs, and two 1Gb Ethernet LAN interfaces built into the motherboard, delivering outstanding levels of density and performance in a compact package.

The Fujitsu RX200 S8 server under test was configured with dual-socket 10-core Xeon E5-2680 V2 2.80GHz CPUs, for a total of 20 cores, and 40 logical CPUs with hyperthreading enabled. The server was equipped with 256GB of RAM and two internal HDDs for OS.

The Fujitsu RX200 S8 server has three PCIe slots. Two of the slots were occupied by two LSI SAS HBA cards. The third slot was installed with an LSI MegaRAID card.

2.2 Microsoft SQL Server 2014 Enterprise

Microsoft SQL Server 2014 Enterprise was installed, and SQL Server 2014 was configured with the following settings.

Startup Options

- -E must be added to the startup options to increase the contiguous extents in database files and to enable large transfer sizes in sequential disk access.

- -T1117 should also be added to the startup options to make sure of even growth of all files in a file group if autogrow is enabled.
- -T834 was added to the startup options to enable large pages. This trace flag was used for column store testing only.

SQL Server Maximum Memory

For ease of benchmarking, SQL Server 2014 maximum memory should be set based on the number of sockets of the reference architecture implementation. Because the RX200 S8 server under test has two sockets, 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 = 20, for row store testing
- MAXDOP = 40 (that is, the system's logical processor count), for column store testing

Resource Governor (RG)

The SQL Server 2014 resource governor was enabled and set to 12%.

2.3 NetApp EF550 Flash Array

The NetApp EF550 is an all-flash storage array that brings together extreme performance and enterprise-grade reliability. The EF550 is a 2U dual controller storage array and for this test was populated with 24 800GB SSDs. The dual controllers were configured with SAS host interface cards (HICs) and 24G of controller cache. The EF550 runs on enterprise-proven NetApp SANtricity® management software version 91.20 and controller firmware version 98.20.

For additional information about NetApp EF-Series solutions, visit www.netapp.com.

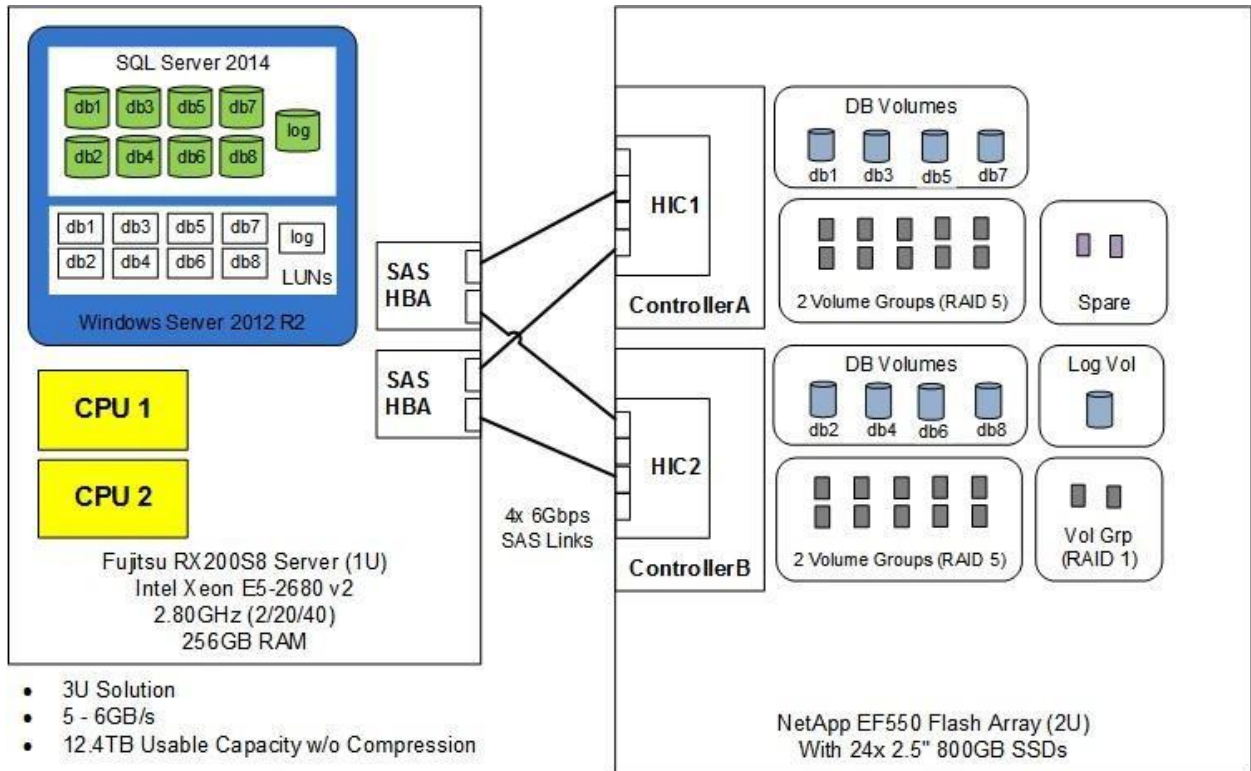
2.4 Database and Storage Layout

Figure 1 shows the detailed storage layout and database layout for the reference configuration. Except for ensuring the dynamic read prefetch option was disabled, all volumes created used the standard default settings. Of the 24 SSDs, four RAID 5 volume groups were created for SQL Server databases, with 5 SSDs per RAID 5 group. A RAID 1 volume group was created with 2 SSDs for transaction log files. The remaining 2 SSDs were configured as global hot spares.

As shown in Figure 1, eight DB volumes were created (2 volumes per RAID 5 volume group) and then mapped to the host as eight database LUNs. One log volume was created on the RAID 1 volume group for database and tempdb log files.

The 1TB target database was equally distributed among the 8 database LUNs. Also, 32 tempdb files (10GB per file) were created and evenly distributed among the 8 database LUNs.

Figure 1) Database and storage layout details.





- 3U Solution
- 5 - 6GB/s
- 12.4TB Usable Capacity w/o Compression

3 EF550 DWFT Reference Architecture Certification

Warning: The data in this section is preliminary and pending Microsoft review and approval. Do not quote.

Figure 2) EF550 DWFT reference architecture certification.

DWFT Certification #2014-027	NetApp EF550 with Fujitsu RX200S8 DWFT Reference Architecture			Report Date: 4/29/2014	
DWFT Rev. 5.4					
System Provider	System Name	Processor Type		Memory	
	EF550	Intel Xeon E5-2680 v2 2.80 GHz (2/20/40)		384 GB	
Operating System			SQL Server Edition		
Windows Server 2012 R2			SQL Server 2014 Enterprise Edition		
Storage Provider	Storage Information				
	20x 880 TB SSD for data and tempdb 2x 146 GB 15K SAS for OS (RAID 1) 2x 800 GB SSD for log (RAID 1)				
Primary Metrics					
Rated User Data Capacity ¹	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹		
(TB)			(TB)		
38	147	204	44		
Row Store					
Relative Throughput ²	Measured Throughput	Measured Scan Rate Physical	Measured Scan Rate Logical	Measured I/O Throughput	Measured CPU (Avg.)
	(Queries/Hr/TB)	(MB/Sec)	(MB/Sec)	(MB/Sec)	(%)
147	172	3,769	4,881	4,325	95
Column Store					
Relative Throughput ²	Measured Throughput	Measured Scan Rate Physical	Measured Scan Rate Logical	Measured I/O Throughput	Measured CPU (Avg.)
	(Queries/Hr/TB)	(MB/Sec)	(MB/Sec)	(MB/Sec)	(%)
204	1,323	1,955	N/A	N/A	99
<p>The reference configuration is a 2 socket system rated for 25TB using the DWFT V4 methodology</p> <p>¹Assumes 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>					

4 Bill of Materials

Table 1 lists the bill of materials for this reference architecture, including both server and storage components.

Table 1) Bill of materials.

Part Number	Description	Quantity
S26361-K1455-V101	PY RX200S8 4x2.5	1
S26361-F3802-E280	Intel Xeon E5-2680v2 10C/20T 2.8GHz 25MB	2
S26361-F3694-E2	Performance mode installation	2
S26361-F3781-E516	16GB (1x16GB) 2Rx4 L DDR3-1600 R ECC	16
S26361-F3269-E2	DVD-RW supermulti slimline SATA	1
S26361-F3628-E201	SAS ctrl 6G 8ext PCIe LP	2
S26361-F4610-E202	PLAN CP 2x1Gbit Cu Intel I350-T2 LP	1
S26361-F1452-E130	Region kit America	1
S26361-F1790-E243	iRMC S4 advanced pack	1
S26113-F574-E12	Modular PSU 800W platinum HP	2
S26361-F1386-E120	Fan upgrade kit 2nd CPU	1
S26361-F4482-E514	146GB 15krpm, <4.5ms, 32MB cache	2
S26361-F3554-E8	RAID ctrl SAS 6G 8-port internal (LSI MegaRAID)	1
EF-X551203A-R6-C	EF550A, 12GB controller, 6Gb SAS, -C	2
E-X5681A-R6	Empty 2U-24 DE5600 enclosure	1
E-X4041A-R6	SSD, 800GB, non-FDE, DE5600	24
X-52197-00-R6	Power cord, in cabinet, 2m, C14-C13, E-Series	1

5 Conclusion

The NetApp reference architecture for Microsoft SQL Server 2014 DWFT presented in this report was designed to show the ease of setting up the EF array and 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 is capable of delivering 500,000 IOPS and 12GB/sec throughput. In this test with a single 1U server, the 3U solution (server and storage) delivers 4325MB/sec row store measured I/O throughput and 1323 queries/hr/TB column store measured throughput.

References

This report references the following documents and resources:

- Microsoft SQL Server Data Warehouse Fast Track:
<http://www.microsoft.com/en-us/server-cloud/data-warehouse-fast-track.aspx>

- NetApp EF-Series solutions:
<http://www.netapp.com>

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 © 1994–2015 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4387-0615