



WHITE PAPER

HP ProLiant DL580 Gen8 and HP LE PCIe Workload Accelerator 90TB Microsoft® SQL Server® Data Warehouse Fast Track Reference Architecture

Based on Microsoft SQL Server 2014 Data Warehouse Fast Track (DWFT) Reference Architecture



Western Digital Technologies, Inc.
951 SanDisk Drive, Milpitas, CA 95035

www.SanDisk.com

Table of Contents

Executive Summary	3
About the HP ProLiant DL580 Gen8	3
Key Benefits and Features.....	3
New Data Warehouse Features in Microsoft SQL Server 2014	4
About the Data Warehouse Fast Track (DWFT) Reference Architecture	4
Storage Configuration	4
Storage Layout	4
Power Override	5
Database Configuration	5
TempDB Configuration.....	5
SQL Server Settings	5
Resource Governor	5
Max Degree of Parallelism (MDOP).....	6
Memory Configuration	6
Trace Flags.....	7
Windows Server 2012 R2 Configuration	7
Power Settings	7
BIOS Configuration	7
SQL Server Data Warehouse Fast Track Reference Architecture Results	8
Summary	9
Bill of Materials	10

Executive Summary

The Microsoft SQL Server Data Warehouse Fast Track (DWFT) reference architecture is designed to eliminate the complexity of properly sizing hardware, which helps reduce unnecessary scale-out of storage and servers. The sizing techniques used in SQL Server DWFT will properly size servers, based on I/O and CPU consumption. This consumption-based approach ensures your data warehouse can fully take advantage of your hardware investment.

This document is for individuals (BI Architects, DBAs, Report Developers, and IT Directors) involved in decision making who are looking for guidance when designing enterprise, business intelligence applications.

About the HP ProLiant DL580 Gen8

The HP ProLiant DL580 Gen8 Server is an enterprise-grade, four-socket (4S) x86 server offering breakthrough performance, rock-solid reliability, and compelling consolidation and manageability efficiencies. It is the ideal choice for mission-critical enterprise, business intelligence, and database applications.

Featuring Intel® Xeon® E7-4800/8800 v2 processors, the HP ProLiant DL580 Gen8 offers 2.3X performance with enhanced processor performance, more memory slots (96 DIMMs), greater I/O bandwidth (9 PCIe Gen3.0 slots), and increased storage performance (12 Gb/s SAS). In addition, HP ProLiant DL580 Gen8 has security and data protection features for system resiliency that your business can depend on. Unique features such as HP Advanced Error Recovery and HP Memory Quarantine increase memory and processor reliability by over 30 percent. With intelligent manageability provided by embedded HP Integrated Lights-Out (iLO 4) capabilities, integrated management through HP OneView, and an array of customer-inspired features, your infrastructure management becomes faster and less expensive.

Key Benefits and Features¹

Breakthrough 4S performance and scalability for blazing-fast results on your most demanding applications

- Up to 2.3X system performance and 50 percent more cores with Intel Xeon E7-4800/8800 v2 processors to speed up your mission-critical enterprise, business intelligence, and database applications
- 1.5X memory slots (96 DIMMs) for large-scale, in-memory computing and virtualization
- 2.7X I/O bandwidth with 9 PCIe Gen3 I/O slots, including 5 x 16 FL/FH slots for greater PCIe flash storage or co-processor cores

Leading x86 availability and rock-solid reliability your business can depend on

- 30 percent fewer memory- and processor-based outages with HP Advanced Error Recovery and HP Memory Quarantine
- Industry-leading performance, uptime, and productivity integrated into a personalized, simplified support experience with HP Proactive Care Support Services

Compelling efficiencies with intelligent manageability and scale-up consolidation

- 45 percent lower total cost of ownership with scale-up consolidation efficiencies, and options for fewer cores and high-frequency processors to drive down software licensing costs
- Faster, lower-cost infrastructure management and a single integrated view of your IT infrastructure with HP OneView
- 3X faster administrative updates with 93 percent less downtime and 69 percent less operator time using industry-leading maintenance tools with HP Smart Update Manager

¹ Source: http://h20195.www2.hp.com/v2/GetDocument.aspx?docname=4AA5-1019ENW&doctype=data_sheet&doclang=EN_US&searchquery=&cc=us&lc=en

New Data Warehouse Features in SQL Server 2014

Microsoft added clustered column store indexes (CCI) in SQL Server 2014, which are designed to decrease query response times and deliver deeper levels of data compression. CCI eliminates the need to build summary tables, thus further reducing ETL run times.

- CCI is optimized for query performance. Our solution deliver 7x better query performance when using CCI. CCI accomplishes this by using a columnar format to compress the data by 10x or more, processing a set of rows in batches, and reading only the columns that are referenced in the query.
- CCI is updateable allowing concurrent insert – both bulk import and trickle insert – of new data while query workload is running. This reduces the data latency from the time data is born to when it is available for querying.

About the DWFT Reference Architecture

The SQL Server DWFT reference architecture provides a scalable framework centered on balancing I/O to achieve maximum performance from SMP-based servers. SQL Server DWFT eliminates the complexity of sizing servers with data warehouses by providing a set of data consumption rates that properly balances performance between the disk subsystem, CPU, and memory.

This architecture is based on the HP ProLiant DL580 Gen8 and HP PCIe LE Workload Accelerator storage controller. This configuration is optimized for data warehouse (scan I/O) workloads and is rated by Microsoft for up to 90TB of compressed data.

More information on SQL Server DWFT reference architecture can be found here:

<http://www.microsoft.com/en-us/server-cloud/data-warehouse-fast-track.aspx>

Storage Configuration

Storage Layout

Slot Number	Device	Capacity	Mount Point	Allocation	Notes
1	HP PCIe LE WA PX600-5200	5.2TB	IOM01	Data Files	JBOD
2	HP PCIe LE WA PX600-5200	5.2TB	IOM02	Data Files	JBOD
3	HP PCIe LE WA PX600-5200	5.2TB	IOM03	Data Files	JBOD
4	HP PCIe LE WA PX600-5200	5.2TB	IOM04	Data Files	JBOD
5	HP PCIe LE WA PX600-5200	5.2TB	IOM05	Data Files	JBOD
6	HP PCIe LE WA PX600-5200	5.2TB	IOM06	Data Files	JBOD
8	HP PCIe LE WA PX600-1300	1.3TB	IOMLOG	Log and Staging Files	Mirrored
9	HP PCIe LE WA PX600-1300	1.3TB	IOMLOG	Log and Staging Files	Mirrored

Power Override

Enabling the power override setting on the HP PCIe LE Workload Accelerator product line is required to achieve the performance results below. A server reboot is required for the setting to be active and persist.

Example:

```
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE <device serial number>:<power (milliwatts)>
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE 1234Z5678:40000, 9876Z-5432:40000
```

Database Configuration

File Group	# of Data Files
FT_Demo_Base	6 (1 data file per data volume)
FT_Demo_stage_part_ci1	6 (1 data file per data volume)
FT_Demo_stage_part_ci2	6 (1 data file per data volume)
FT_Demo_stage_part_ci3	6 (1 data file per data volume)
FT_Demo_stage_part_ci4	6 (1 data file per data volume)
FT_Demo_stage_part_ci5	6 (1 data file per data volume)
FT_Demo_stage_part_ci6	6 (1 data file per data volume)
FT_Demo_stage_part_ci7	6 (1 data file per data volume)
FT_Demo_LOG	1 (transaction log on mount volume)

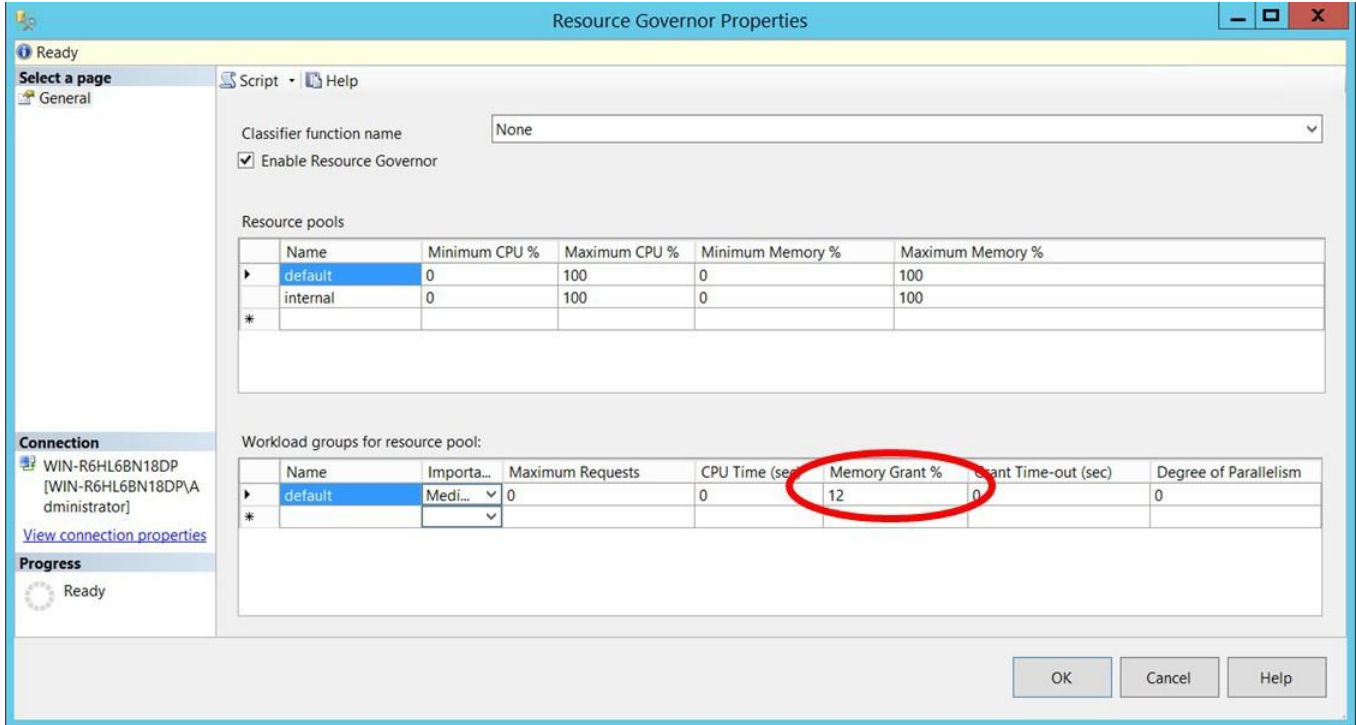
TempDB Configuration

One 20GB tempdb file was stored on each ioMemory volume. In total, six tempdb data files were spread across six ioMemory volumes. The tempdb transaction log file was stored on the volume designated from log and staging files.

SQL Server Settings

Resource Governor

The Memory Grant % value was set to 12% of the memory allocated for row store runs and 25% for column store runs. The settings were changed in the Default resource pool.



Max Degree of Parallelism (MDOP)

MDOP was set to 30 for row store and 120 for column store. This provided the best scan rates for the respective runs.

Example:

```
--for row store runs
EXEC sp_configure 'max degree of parallelism', 30
GO
RECONFIGURE WITH OVERRIDE
GO
--for column store runs
EXEC sp_configure 'max degree of parallelism', 120
GO
RECONFIGURE WITH OVERRIDE
GO
```

Memory Configuration

- SQL Server was allocated 90% of the server memory.
- The SQL Server service account was assigned the Lock Pages in Memory policy.

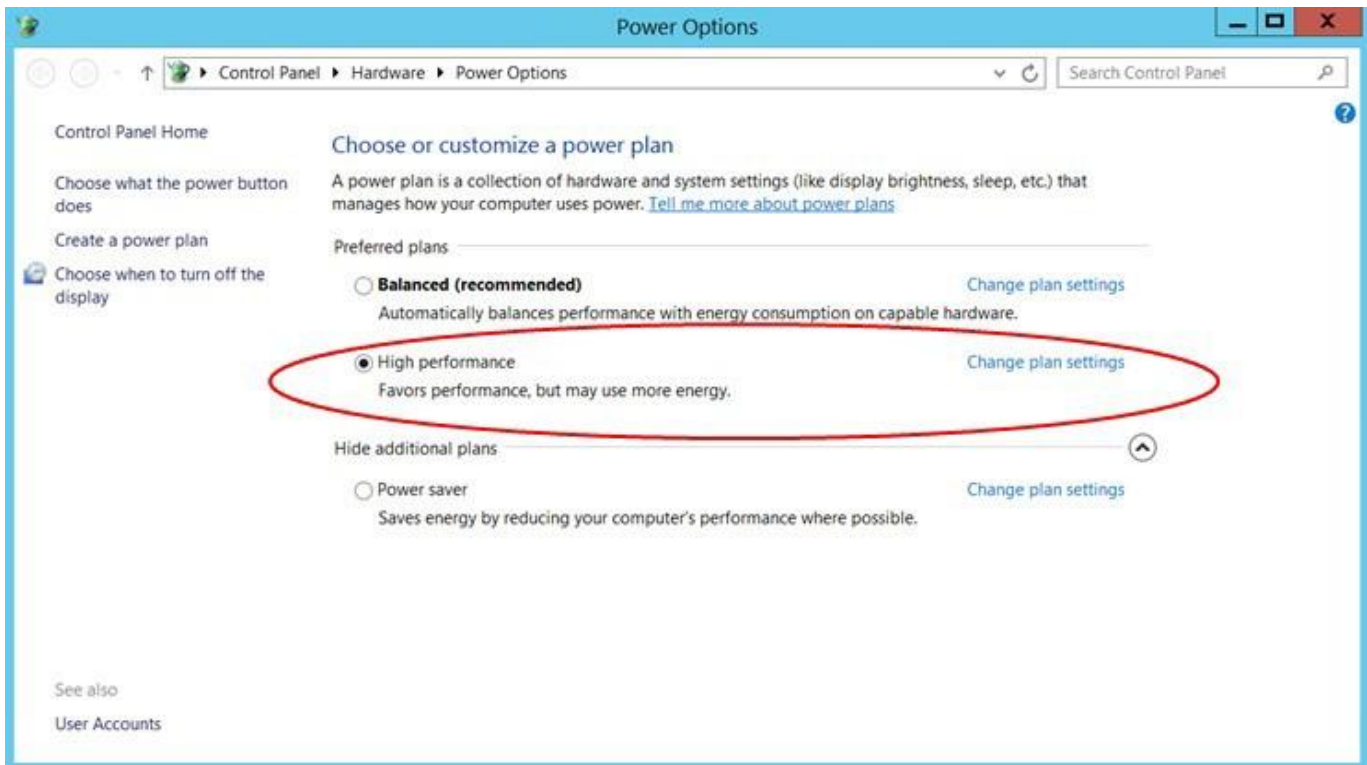
Trace Flags

Trace flag --T1117 was used to increase performance. This flag forces all data files in a file group to grow at once, which reduces “hot spots” of data pages. This ensures that all databases with more than one data file will grow properly across all the data mounts, which in turn ensures maximum physical I/O performance. Trace flag “-E” was omitted, as testing revealed a sharp increase in queue depth and latency which decreases scan performance. Therefore trace flag “-E” is not recommended with our flash technology.

Windows Server 2012 R2 Configuration

Power Settings



The High Performance plan was chosen to reduce CPU throttling.



BIOS Configuration

- Hyper-threading was enabled.
- System Profile was set to High Performance mode.
- Fan Offset was set to “Increased Cooling”.

SQL Server Data Warehouse Fast Track Reference Architecture Results

DWFT Certification #2014-002	HP ProLiant DL580 Gen8 with HP PCIe LE Workload Accelerator 90TB reference architecture for Microsoft SQL Server 2014 Data Warehouse Fast Track			Report Date 9/3/2014	
DWFT Rev. 5.4					
System Provider	System Name	Processor Type	Memory		
	HP ProLiant DL580 Gen8	Intel Xeon E7-4890 v2 2.8 GHz (4/60/120)	1536GB		
Operating System		SQL Server Edition			
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition			
Storage Provider	Storage Information				
	6 x 5.2TB HP PCIe LE Workload Accelerator Flash devices for data and tempdb 2 x 500GB 7200RPM HDD for OS (RAID 1) 2 x 1.3TB HP PCIe LE Workload Accelerator Flash devices for log (RAID 1)				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity (TB)		
90	314	419	108		
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.) (%)
314	366	8,351	10,139	9,245	92
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.) (%)
419	2,721	3,392	N/A	N/A	96

The reference configuration is a 2 socket system rated for 25TB using the FTDW 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.

Summary

This solution went through hundreds of hours of testing and engineering to provide the most optimal and reliable configuration for the HP and SanDisk® SQL Server Data Warehouse Fast Track Reference Architecture. The HP PCIe LE Workload Accelerator simplifies storage configuration by reducing the importance of sequential I/O, as evangelized in previous Data Warehouse Fast Track Reference Architectures.

With 9000+ MB/s of database read throughput and a rated user data capacity 90TB, the HP ProLiant DL580 Gen8 with the HP PCIe LE Workload Accelerator delivers industry-leading, breathtaking performance with the ability to host extremely large data warehouses or a consolidation of data warehouses.

Bill of Materials

90TB Certification

SKU	Description	Quantity
728551-B21	HP ProLiant DL580 Gen8 Configure-to-order Server	1
728955-L21	HP ProLiant DL580 Gen8 Intel® Xeon® E7-4890v2 (2.8GHz/15-core/37.5MB/155W) FIO	4
732411-B21	HP ProLiant DL580 Gen8 12 DIMM Slots Memory Cartridge	8
708641-B21	HP 16GB (1x16GB) Dual Rank x4 PC3-14900R (DDR3-1866) Registered CAS-13 Memory Kit	96
700752-B21	HP FlexFabric 10Gb 2-port 534FLR-SFP+ FIO Adapter	1
698537-B21	HP 4GB FIO Flash Backed Write Cache	1
775672-B21	HP 5.2TB FH/HL PCIe Light Endurance (LE) Workload Accelerator	6
775668-B21	HP 1.3TB HH/HL PCIe Light Endurance (LE) Workload Accelerator	2
652611-B21	HP 300GB 6G SAS 15K rpm SFF (2.5-inch) SC Enterprise 3yr Warranty Hard Drive	4
755996-B21	Microsoft Windows Server 2012 R2 Standard Edition FIO	2
	Microsoft SQL Server 2014 Enterprise Edition License (60 cores) – Available through Microsoft	

Specifications are subject to change. ©2014 - 2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk and the SanDisk logo are trademarks of Western Digital Corporation or its affiliates, registered in the U.S. and other countries. Fusion ioMemory is a trademark of Western Digital Corporation or its affiliates. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s). HP DL580 Gen8 20160621

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.

