HP ProLiant DL380 Gen8 and HP PCIe LE Workload Accelerator 28TB/45TB Data Warehouse Fast Track Reference Architecture

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

SanDisk
a Western Digital brand

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

www.SanDisk.com

Table of Contents

Executive Summary

The Microsoft SQL Server Data Warehouse Fast Track (DWTF) 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/0 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.

Why Choose Our Solutions?

Together, HP and SanDisk dedicated hundreds of hours of testing to engineer the SQL Server DWFT solutions to ensure maximum reliability and performance. These series of tests pushed the HP ProLiant DL380 Gen8 and HP PCIe LE Workload Accelerator to their peak performance without failure of the hardware. The reliability and performance experienced during testing is what can be expected in production environments.

About the HP Proliant Dl380 Gen8

What is your server bottleneck – storage, processing, or expansion? No matter the bottleneck, the HP ProLiant DL380p Gen8 Server series can help. It sets the data center standard for 2U, 2-socket rack servers with the latest in serviceability, enhanced configuration flexibility, customer-inspired design, and unmatched performance.

The HP ProLiant DL380p Gen8 Server offers the perfect solution for the dynamic compute requirements of growing small businesses as well as demanding data centers.

What's New?[1]

- Eliminate infrastructure complexity with HP 0neView 1.10, now available on all HP ProLiant DL Gen8 Servers.
- Increase performance of up to 35%[1] at the same or lower power levels with the new Intel® Xeon® E5-2600 v2 processors.
- GPU support is provided in primary and secondary risers to accelerate compute and enable large-data processing for a wide array of applications.
- New HP Common Slot Power Supplies provide 96% server power efficiency, and new infrastructure power efficiencies are achieved with 277VAC and 380VDC input voltages.
- HP SmartMemory enables 16.6% better memory performance.
- Up to 4x workload performance with HP SmartCache, a controller-based caching solution that caches hot data onto lower-latency SSDs.

[1] Source: http://www8.hp.com/us/en/products/proliant-servers/product-detail.html?oid=5177957#!tab%3Dfeatures

Features

- Improved capacity and performance on compute and storage.
- Intel® Xeon® processor E5-2600 v2 product family offering increased performance, improved security and efficiency; ideal for demanding workloads.
- HP SmartMemory, with speeds up to 1866MHz. This prevents data loss and downtime with enhanced error handling, while improving workload performance and power efficiency.
- GPU compute support to maximize performance.
- HP SmartCache, a controller-based caching solution. This caches hot data onto lower latency SSDs providing up to 4x workload performance.[2]

New Data Warehouse Features in SQl Server 20l4

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. 0ur solutions 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 SQl Server Data Warehouse Fast Track Reference Architecture

The SQL Server DWFT reference architecture provides a scalable framework centered on balancing I/0 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 DL380 Gen8 and HP PCIe LE Workload Accelerator PX600-2600/PX600-5200 storage controllers. These configurations are optimized for data warehouse (scan I/0) workloads and are rated, by Microsoft for up to 45TB of compressed data.

More information on SQL Server DWFT can be found here:

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

[2] Feb. 2013 - Performance and configuration as outlined in Whitepaper TC1211951

Data Warehouse Fast Track Reference Architecture Configurations

Storage Configuration

Storage layout for 28TB Certification

Slot Number	Device	Capacity	Mount Point	Allocation	Notes
1	HP PCIe LE WA PX600-2600	2.6TB	I0M01	Data Files	JB0D
2	HP PCIe LE WA PX600-2600	2.6TB	I0M02	Data Files	JB0D
4	HP PCIe LE WA PX600-2600	2.6TB	I0M03	Data Files	JB0D
5	HP PCIe LE WA PX600-2600	2.6TB	I0M04	Data Files	JB0D
3,6	Empty				

Storage layout for 45TB Certification

Slot Number	Device	Capacity	Mount Point	Allocation	Notes
1	HP PCIe LE WA PX600-5200	5.2TB	I0M01	Data Files	JB0D
2	HP PCIe LE WA PX600-5200	5.2TB	I0M02	Data Files	JB0D
4	HP PCIe LE WA PX600-5200	5.2TB	I0M03	Data Files	JB0D
5	HP PCIe LE WA PX600-5200	5.2TB	I0M04	Data Files	JB0D
3,6	Empty				

A- SAS controller was used for the database transaction log in both configurations, with 8 x 300GB 15K spindles in RAID 10.

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. For instructions on enabling the power override, see the HP PCIe LE Workload Accelerator User Guide for instructions.

Example:

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

Database Configuration

File Group	# of Data Files
FT_Demo_Base	4 (1 data file per data volume)
FT_Demo_stage_part_ci1	4 (1 data file per data volume)
FT_Demo_stage_part_ci2	4 (1 data file per data volume)
FT_Demo_stage_part_ci3	4 (1 data file per data volume)
FT_Demo_stage_part_ci4	4 (1 data file per data volume)
FT_Demo_stage_part_ci5	4 (1 data file per data volume)
FT_Demo_stage_part_ci6	4 (1 data file per data volume)
FT_Demo_stage_part_ci7	4 (1 data file per data volume)
FT_Demo_L0G	1 (transaction log on log volume)

TempDB Configuration

Four tempdb files, 5GB each, were stored on each data volume. In total, 24 tempdb data files were spread across six data mount points. 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.

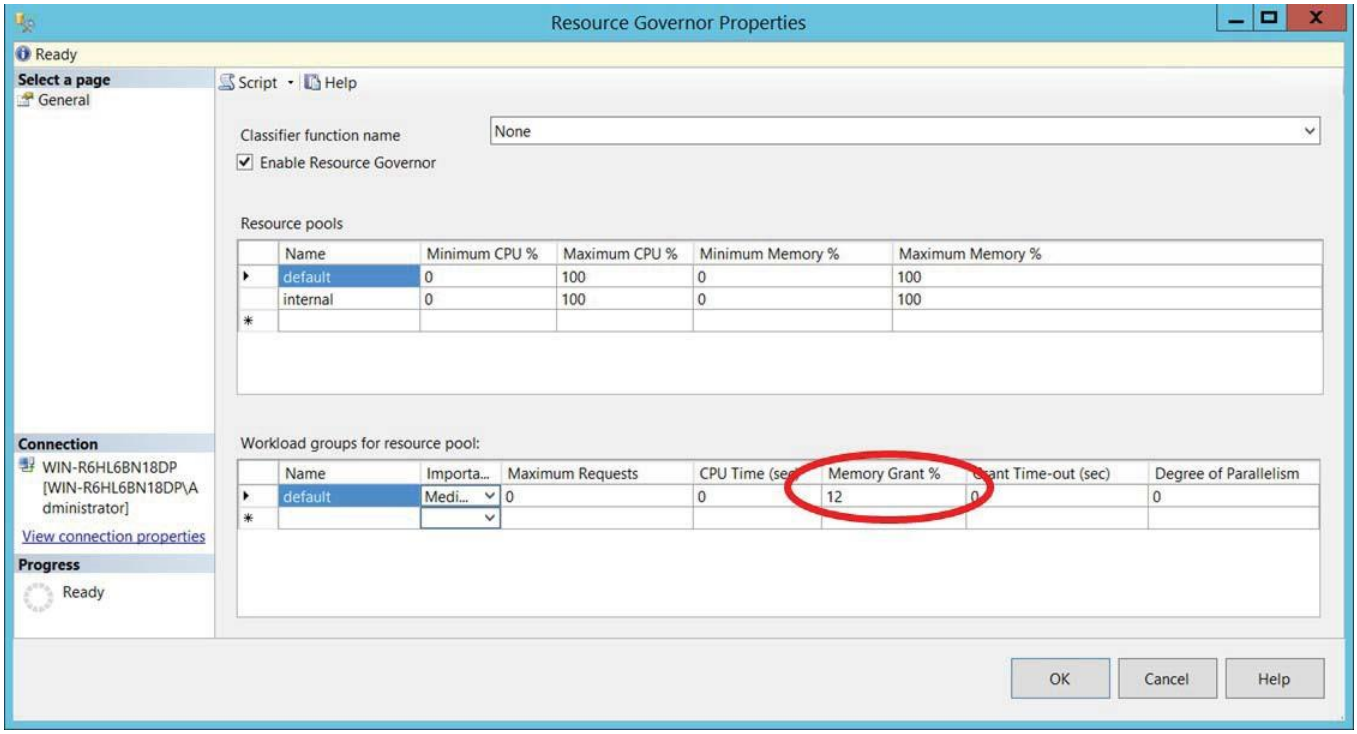

Max Degree of Parallelism (MDOP)

MD0P was set to 24 for row store and 48 for column store. These values provided the best scan rates for the respective runs without pegging the processors @ 100% utilization.

Example:

```
--for row store runs
EXEC sp_configure 'max degree of parallelism', 24
GO
RECONFIGURE WITH OVERRIDE
GO
--for column store runs
EXEC sp_configure 'max degree of parallelism', 48
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/0 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.

Server Configuration

Windows Server 20l2 R2 Configuration

Power Settings

The High Performance plan was chosen to reduce CPU throttling.

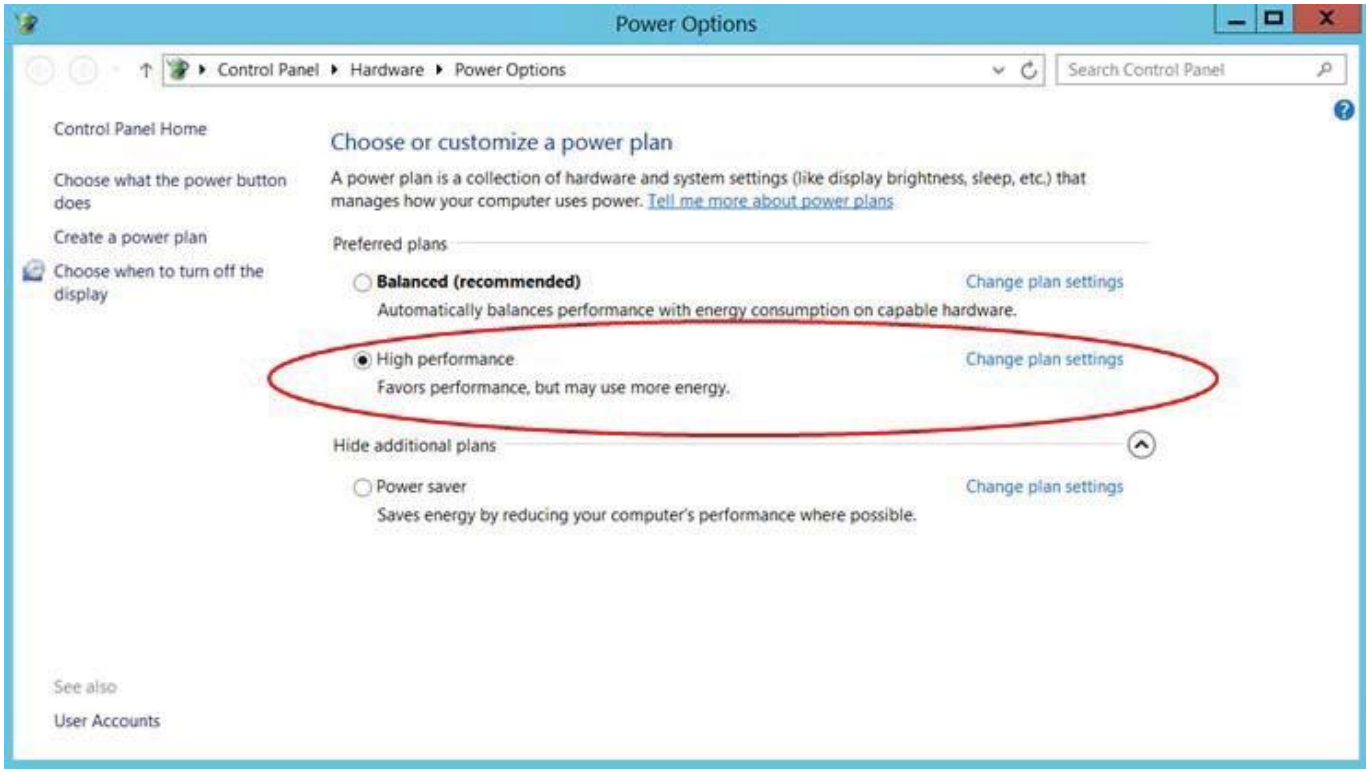

BIOS Configuration

- Hyper-threading was enabled.
- System Profile was set to High Performance mode.
- Fan 0ffset was set to "Increased Cooling".

SQl Server Data Warehouse Fast Track Reference Architecture Results

28TB Certification

DWFT Certification #2014-003 DWFT Rev. 5.4	HP ProLiant DL380 Gen8 with HP PCIe LE Workload Accelerator 28TB reference architecture for Microsoft SQL Server 2014 Data Warehouse Fast Track		Report Date 9/3/2014
System Provider	**System Name**	**Processor Type**	**Memory**
hp invent	HP ProLiant DL380 Gen8	Intel Xeon E5-2697v2 2.7 GHz (12/24/48)	768GB
Operating System		**SQl Server Edition**	
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition	
Storage Provider	**Storage Information**		
hp invent	4 x 2.6TB HP PCIe LE Workload Accelerators for data and tempdb 2 x 300GB HDDs for 0S (RAID 1) 8 x 15K 300GB HDDs for log (RAID 10)		

Primary Metrics			
Rated User Data Capacity[1] (TB)	Row Store Relative Throughput[2]	Column Store Relative Throughput[3]	Maximum User Data Capacity (TB)
28	170	218	32

Row Store					
Relative Throughput[2]	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/0 Throughput (MB/Sec)	Measured CPU (Avg.) (%)
170	202	4,508	5,500	5,004	96

Column Store					
Relative Throughput[2]	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/0 Throughput (MB/Sec)	Measured CPU (Avg.) (%)
218	1,414	1,389	N/A	N/A	98

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.

45TB Certification

DWFT Certification #2014-003 DWFT Rev. 5.4	HP ProLiant DL380 Gen8 with HP PCIe LE Workload Accelerator 45TB reference architecture for Microsoft SQL Server 2014 Data Warehouse Fast Track		Report Date 9/3/2014
System Provider	**System Name**	**Processor Type**	**Memory**
hp invent	HP ProLiant DL380 Gen8	Intel Xeon E5-2697v2 2.7 GHz (12/24/48)	768 GB
Operating System		**SQl Server Edition**	
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition	
Storage Provider	**Storage Information**		
hp invent	4 x 5.2TB HP PCIe LE Workload Accelerators for data and tempdb 2 x 300GB HDDs for 0S (RAID 1) 8 x 15K 300GB HDDs for log (RAID 10)		

Primary Metrics			
Rated User Data Capacity[1] (TB)	Row Store Relative Throughput[2]	Column Store Relative Throughput[3]	Maximum User Data Capacity (TB)
45	161	227	72

Row Store					
Relative Throughput[2]	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/0 Throughput (MB/Sec)	Measured CPU (Avg.) (%)
161	198	4,220	5,240	4,730	97

Column Store					
Relative Throughput[2]	Measured Throughput (Queries/Hr/TB)	Measured Scan Rate Physical (MB/Sec)	Measured Scan Rate Logical (MB/Sec)	Measured I/0 Throughput (MB/Sec)	Measured CPU (Avg.) (%)
227	1,476	1,443	N/A	N/A	99

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

These solutions went through hundreds of hours of testing and engineering to provide the most optimal and reliable configuration for the HP and SanDisk SQL Server DWFT Reference Architecture. The HP ProLiant DL380 Gen8 and HP PCIe LE Workload Accelerators deliver 5000+MB/s of consistent database performance while providing the best reliability in the industry. The HP PCIe LE Workload Accelerator simplifies storage configuration by reducing the importance of sequential I/0, as evangelized in previous DWFT Reference Architectures.

With a rated user capacity of 28TB and 45TB, these 2U configurations deliver the best mix of performance and data capacity, which reduces the need to scale your data warehouse. In fact, these solutions allow for massive database consolidation projects, allowing your organization to save on licensing cost.

The results highlight how HP and SanDisk can deliver enterprise-level solutions that serve as the foundation for data warehouse or database consolidation projects.

Bill of Materials

SKU	Description	Quantity
	28TB Configuration	
653200-B21	HP ProLiant DL380p Gen8 8 SFF Configure-to-order Server	1
715224-L21	HP ProLiant DL380p Gen8 Intel® Xeon® E5-2697v2 (2.7GHz/12-core/30MB/130W) FI0 Processor Kit	2
708643-B21	HP 32GB (1x32GB) Quad Rank x4 PC3-14900L (DDR3-1866) Load Reduced CAS-13 Memory Kit	24
775670-B21	HP 2.6TB HH/HL PCIe Light Endurance (LE) Workload Accelerator	4
652611-B21	HP 300GB 6G SAS 15K rpm SFF (2.5-inch) SC Enterprise 3yr Warranty Hard Drive	8
684208-B21	HP Ethernet 1Gb 4-port 331FLR FI0 Adapter	1
656363-B21	HP 750W Common Slot Platinum Plus Hot Plug Power Supply Kit	2
663476-B21	HP 2U FI0 Friction Rail Kit	1
755996-B21	Windows Server 2012 R2 Standard Edition 2P Pre-Installed on ProLiant Servers	
	45TB Configuration	
653200-B21	HP ProLiant DL380p Gen8 8 SFF Configure-to-order Server	1
715224-L21	HP ProLiant DL380p Gen8 Intel® Xeon® E5-2697v2 (2.7GHz/12-core/30MB/130W) FI0 Processor Kit	2
708643-B21	HP 32GB (1x32GB) Quad Rank x4 PC3-14900L (DDR3-1866) Load Reduced CAS-13 Memory Kit	24
775672-B21	HP 5.2TB FH/HL PCIe Light Endurance (LE) Workload Accelerator	4
652611-B21	HP 300GB 6G SAS 15K rpm SFF (2.5-inch) SC Enterprise 3yr Warranty Hard Drive	8
684208-B21	HP Ethernet 1Gb 4-port 331FLR FI0 Adapter	1
656363-B21	HP 750W Common Slot Platinum Plus Hot Plug Power Supply Kit	2
663476-B21	HP 2U FI0 Friction Rail Kit	1
755996-B21	Windows Server 2012 R2 Standard Edition 2P Pre-Installed on ProLiant Servers	

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