



WHITE PAPER

HPE ProLiant DL580 Gen10 and Ultrastar[®] SS300 SSD 195TB Microsoft[®] SQL Server[®] Data Warehouse Fast Track Reference Architecture

BASED ON MICROSOFT SQL SERVER 2017 DATA WAREHOUSE FAST TRACK REFERENCE ARCHITECTURE

Western Digital.

Table of Contents

Executive Summary	3
Performance Highlights	3
About the Data Warehouse Fast Track Reference Architecture	3
About the HPE DL580 Gen10.....	3
Scalable Performance in an Expandable 4U Form Factor	4
Boost IOPS Performance with Ultrastar® SSDs.....	4
Data Warehouse Features in SQL Server 2017	5
Storage Configuration.....	5
Physical Storage Layout	5
Standalone Storage Spaces Configuration	6
Database Configuration	6
Tempdb Configuration	7
SQL Server 2017 Configuration Parameters	7
Windows Server 2016 and BIOS Configuration	7
SQL Server Data Warehouse Fast Track Reference Architecture Results.....	8
Comparing the 120TB vs. 195TB Data Warehouse Fast Track Ras	9
Primary Metrics	9
Row store Comparison.....	9
Columnstore Comparison	9
Summary.....	10
Appendix A: Bill of Materials	10
HP Enterprise ProLiant DL580 Gen10: Non-High-Availability Option	10
Appendix B: SQL Server Read Analysis – Extended Events.....	11
Top Read Block Sizes.....	11

Executive Summary

The Microsoft SQL Server® Data Warehouse Fast Track (DWFT) reference architecture is designed to eliminate the complexity of adequately sizing hardware, which helps reduce unnecessary scale-out of storage and servers. The sizing techniques used in SQL Server Data Warehouse Fast Track will correctly 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 (Cloud Providers, BI Architects, DBAs, Report Developers, and IT Directors) involved in purchasing or system architecture decision making who are looking for guidance when designing enterprise Business Intelligence applications.

Reference Architecture Performance Highlights

Peak Throughput	Measure I/O Throughput	Average CPU Utilization%	Average Read Latency	Queries/Hr/TB	Max User Data Capacity
42GB/s	20GB/s	95%	500 µsec	5850	219TB

About the Data Warehouse Fast Track Reference Architecture

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

Based on the HPE ProLiant DL580 Gen10 and Ultrastar® SS300 SSD, this optimized configuration is for data warehouse (scan I/O) workloads and is rated by Microsoft for up to 195TB of compressed data.

For more information on the SQL Server Data Warehouse Fast Track program click here:

<https://www.microsoft.com/en-us/sql-server/data-warehousing>

About the HPE ProLiant DL580 Gen10

The HPE ProLiant DL580 Gen10 Server is a highly expandable, 4P server with high performance, scalability, and availability in a 4U chassis. Supporting the Intel® Xeon® Scalable processors with up to a 28%¹ performance gain, the HPE ProLiant DL580 Gen10 Server delivers higher processing power than previous generations. The processor improvements provide up to 6 TB of 2666 MT/s memory with up to 66% greater memory bandwidth² and up to 16 PCIe 3.0 slots, plus the simplicity of automated management with HPE OneView and HPE iLO 5.



For additional details, go to www.hpe.com/servers/dl580-gen10.

¹ Intel measurements. Up to 28% performance increase of Intel Xeon Platinum versus previous generation E7-8800 v4 average performance based on generational gains on HPE servers comparing 4-socket Intel Xeon Platinum 8180 to E7-8890 v4 family processors. Any difference in system hardware or software design or configuration may affect actual performance. May 2017.

² Percentage compares Gen10 vs Gen9: Gen10 = 12 Channels x 2666 data rate x 8 bytes = 256 GB/sec. Gen9 = 8 channels x 2400 x 8 bytes = 154 GB/sec. 256/154 = 1.66 or Gen10 is 66% greater bandwidth, July 2017.

Features

Scalable Performance in an Expandable 4U Form Factor

- The completely redesigned HPE ProLiant DL580 Gen10 Server provides 4P computing in an expandable 4U form factor. It also supports up to four Intel Xeon Platinum and Gold processors, which provide up to 28% more processor performance and 17%³ more cores than the previous generation.
- Up to 48 DIMM slots are provided, which support up to 6 TB for 2666 MT/s HPE DDR4 SmartMemory. HPE DDR4 SmartMemory improves workload performance and power efficiency while reducing data loss and downtime with enhanced error handling.
- A choice of HPE FlexibleLOM Adapters offers a range of networking speeds (1GbE to 25GbE) and fabrics so your systems can adapt and grow to changing business needs.
- HPE Innovation with Intelligent System Tuning enhances workload performance using customized profiles to tune internal resources. This improves workloads such as high-frequency trading with jitter smoothing and provides added performance with a core-boosting SKU.

Boost IOPS Performance with Ultrastar® SSDs

Offered in a 2.5-inch small form factor, the Ultrastar® SS300 solid-state drive (SSD) delivers ultra-high performance to power through the most challenging workloads.

Designed with a 12Gb/s SAS interface for seamless integration into enterprise environments, the Ultrastar SS300 delivers high sequential throughput—up to 2100MB/s read at 14W power. Available with capacities from 7.68TB⁴ to 400GB, the SS300 delivers up to 400,000 read, and 200,000 write IOPS. This provides rapid access to “hot” enterprise data for higher productivity and operational efficiency. The Ultrastar SS300 family offers significant value regarding IOPS per watt, reducing total cost of ownership (TCO) through low power consumption, efficient cooling and reduced space requirements.

For more information on the Ultrastar SS300, visit <http://www.hgst.com/products/solid-state-solutions/ultrastar-ss300>



³ Up to 17% performance increase of Intel Xeon Platinum vs. previous generation comparing 4-socket Intel Xeon Platinum 8180 (28 cores) to E7-8890 v4 (24 cores). Calculation 28 cores/24 cores= 1.167 = 17%. August 2017.

⁴ One MB is equal to one million bytes, one GB is equal to one billion bytes, one TB equals 1,000GB and one PB equals 1,000TB when referring to HDD/SSD capacity. Accessible capacity will vary from the stated capacity due to formatting and partitioning of the HDD/SSD drive, the computer's operating system, and other factors.

Data Warehouse Features in SQL Server 2017


Microsoft added Clustered Columnstore Indexes (CCI) in SQL Server 2012, which are designed to decrease query response times and deliver more profound levels of data compression. CCI eliminates the need to build summary tables, thus further reducing ETL run times.

- Our solution delivers 10x 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 referenced in the query.
- CCI is updatable, allowing concurrent insert – both bulk import and trickle insert – of new data while query workload is running. Updateability reduces the data latency from the time data is born to when it is available for querying.

Storage Configuration

Physical Storage Layout

Forty-eight Ultrastar SS300 SSDs connect to six Smart Array controllers, which feed into three processors via Gen3 x8 PCIe slots. Each Smart Array controller hosts eight Ultrastar SS300 SSDs. Under synthetic load, peak performance reached 41GB/s. Each controller’s peak performance reached 6.83GB/s.




Standalone Storage Spaces Configuration

Standalone Storage Spaces was used to simplify the management of data protection. The HPE ProLiant DL580 Gen10 was configured with six Smart Array P408i-p SR SAS controllers. Pass-thru mode was enabled on the SAS controllers to allow Windows Server 2016 to manage data resiliency.

One storage pool with simple resiliency hosts five volumes. All volumes were configured with mirroring.

Physical Devices	48	Total Capacity
Storage Pool(s)	1	153TB (raw)
Data Volumes	4	63TB (useable)
Log Volume	1	3.5TB



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_LOG	1 (transaction log on mount volume)

Tempdb Configuration

160 x 12.5GB tempdb files were evenly divided across the 4 data volumes. During testing, we noticed a performance increase when running Columnstore tests when using a large number of tempdb files*. The map contention was alleviated as the number of tempdb files increased. The tempdb transaction log file was stored on the volume designated for log files.

*Your results may vary with Tempdb configuration. Always configure Tempdb with a minimum of 8 data files. Increase the number of data files until the contention is alleviated.

SQL Server 2017 Configuration Parameters

Parameter	Setting	Description
Memory Allocation	236GB	This is the Fast Track-required value for a 4-socket system with a 2 TB database. Memory is deliberately constrained to enforce I/O pressure on the storage subsystem.
Max Degree of Parallelism, Row store	40	When SQL Server runs on a computer with more than one CPU, it detects the best degree of parallelism (the number of processors used in the execution of a parallel plan).
Max Degree of Parallelism, Columnstore	160	
Memory Grant %, Row store	6%	The default is 25%. This is reduced to 6% for Row store to reduce the maximum memory consumed per query.
Memory Grant %, Columnstore	25%	The default is 25%. This value was not changed for the official test runs.
Startup Parameters, Row store only	834	When this flag is set, SQL Server uses Windows large-page memory allocations for the buffer pool. This trace flag can improve throughput rates for many data warehouse workloads. This value is disabled for Columnstore runs.

Windows Server 2016 and BIOS Configuration

Parameter	Setting	Description
Power Setting	High Performance	The default Power Setting is "balanced". High Performance disables power throttling in Windows.
Hyper-threading	Enabled	Hyper-threading improves parallelism by virtualizing physical cores. Enabling hyper-threading is a best practice for SQL Server workloads.
System Profile	I/O Throughput	The I/O Throughput profile setting configures the BIOS for a database workload. All P and C-states were disabled.
Lock Pages in Memory	Enabled for SQL Service Account	This prevents the SQL Server service account from paging to virtual memory on disk.

SQL Server Data Warehouse Fast Track Reference Architecture Results

DWFT Certification #2017-012	HPE ProLiant DL580 Gen10 HGST Ultrastar SS300 SSD DWFT Reference Architecture		Report Date: 3/19/2018
DWFT Rev. 5.4			
System Provider	System Name	Processor Type	Memory
	HPE DL580 Gen10	Intel Platinum 8180 2.5 GHz (4/112/224)	4096 GB
Operating System		SQL Server Edition	
Windows Server 2016		SQL Server 2017 Enterprise Edition	
Storage Provider	Storage Information		
	48 x 3.2TB SSD for log, data and tempdb 2 x 340 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)
195	689	900	219

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.) (%)
689	711	18,597	22,010	20,304	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/O Throughput (MB/Sec)	Measured CPU (Avg.) (%)
900	5,850	3,691	N/A	N/A	94

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.

Comparing the 120TB vs. 195TB Data Warehouse Fast Track RAs

Our previous SQL Server Data Warehouse Fast Track with the HPE ProLiant DL580 Gen9 scored a 120TB User Data Capacity rating. The 120TB solution leveraged our legacy Fusion-io/SanDisk ioMemory PCIe SSD. Mentioned earlier, this solution leverages our Ultrastar SS300 SAS SSD. To demonstrate the performance leap between the legacy Fusion-io/SanDisk and the new Ultrastar SS300 SSD we compared the certifications side-by-side.

Primary Metrics

Parameter	120TB	195TB	% Difference
Rated User Data Capacity (TB)	120	195	28%
Row store Relative Throughput	468	689	33%
Columnstore Relative Throughput	569	900	36%
Maximum User Data Capacity (TB)	178	219	19%

Row store Comparison

Parameter	120TB	195TB	% Difference
Measured Throughput (Queries/Hr/TB)	468	711	34%
Measured Scan Rate, Physical	11,356	18,597	39%
Measured Scan Rate, Logical	13,826	22,010	37%
Measure I/O Throughput	12,591	20,304	38%
Measured CPU %	80	96	16%

Columnstore Comparison

Parameter	120TB	195TB	% Difference
Measured Throughput (Queries/Hr/TB)	3,696	5,850	39%
Measured Scan Rate, Physical	2,611	3,691	29%
Measured Scan Rate, Logical	N/A	N/A	N/A
Measure I/O Throughput	N/A	N/A	N/A
Measured CPU %	93	94	1%

Summary

As with all our SQL Server Data Warehouse Fast Track solutions, this solution went through hundreds of hours of testing and engineering to provide the most optimal and reliable configuration. The Standalone Storage Spaces simplifies the storage configuration while the Ultrastar SS300 SSD provides excellent performance.

The HPE ProLiant DL580 Gen10 with the Ultrastar SS300 SSD delivers industry-leading performance with the ability to host vast data warehouses or consolidate data warehouses. It features sub-500 microsecond latency, 25GB/s average throughput, 42GB/s of peak database read throughput, and an SQL Server Data Warehouse Fast Track record-breaking rated user data capacity of 195TB.

Appendix A: Bill of Materials

HP Enterprise ProLiant DL580 Gen10: Non-High-Availability Configuration

Item#	Qty	Product #	Product Description
0100	1	869854-B21	HPE ProLiant DL580 Gen10 8SFF Configure-to-order Server
0101	1	878143-L21	HPE DL580 Gen10 Intel Xeon-Gold 6148 (2.4GHz/20-core/145W) FIO Processor Kit
0102	3	878143-B21	HPE DL580 Gen10 Intel Xeon-Gold 6148 (2.4GHz/20-core/145W) Processor Kit
0103	64	815101-B21	HPE 64GB (1x64GB) Quad Rank x4 DDR4-2666 CAS-19-19-19 Load Reduced Smart Memory Kit
0104	5	878366-B21	HPE DL580 Gen10 8SFF HDD Bay Kit
0105	1	872340-B21	HPE DL580 Gen10 9-slot 6 x8/3 x16 Secondary Riser Kit
0106	1	878214-B21	HPE DL580 Gen10 7-slot 4 x8/3 x16 Primary Riser Kit
0107	6	830824-B21	HPE Smart Array P408i-p SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS PCIe Plug-in Controller
0108	2	875241-B21	HPE 96W Smart Storage Battery (up to 20 /145mm Cable) Kit
0109	4	830272-B21	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
0110	1	872222-B21	HPE DL5x0 Gen10 CPU Mezzanine Board Kit
0200	1	700751-B21	HPE FlexFabric 10Gb 2-port 534FLR-SFP+ Adapter
0300	1	E6U59ABE	HPE iLO Advanced including 1yr 24x7 Technical Support and Updates E-LTU
0400	48	872386-B21	HPE 3.2TB 12G SAS MU SFF SC DS SSD

Appendix B: SQL Server Read Analysis – Extended Events

The summary data below was captured using Extended Events during a row store tests on the 165TB rated SQL Server Data Warehouse Fast Track Reference Architecture. SQL Server primarily reads in two block sizes: 8KB (the size of a data page) and 64KB (the size of an extent, or 8*8K data pages). During testing, we observed I/O requests from SQL Server fetching I/Os as large as 4MB.

8KB and 64KB block sizes account for 78% of the read I/O from SQL Server. More read I/Os were issued for block sizes smaller than 512KB, most at 8KB and 64KB. However, the accumulative amount of data read from the database were from block sizes greater than 512KB (SQL Server read more data from block sizes higher than 512KB).

Block Size	Unique Block Size Count	Read I/Os	Read I/O %	Total Read Gigabytes	Total Read Gigabytes %
< 512K	53	293854137	90%	22955	45%
>= 512K	176	32237242	10%	28063	55%

Top Read Block Sizes

Read Block Size (Bytes)	Reads	% of Reads	Total GB Read
65536	231556889	71%	14133
8192	24174862	7%	184
524288	13502172	4%	6593
344064	13072897	4%	4189
131072	7570073	2%	924
983040	3991866	1%	3655
720896	3265702	1%	2193
73728	2638968	1%	181
458752	2446798	1%	1045
196608	2382649	1%	436
786432	2110010	1%	1545
57344	1077050	0%	58
16384	1020929	0%	16
262144	1013864	0%	248
270336	997190	0%	251
24576	982225	0%	22
32768	973772	0%	30
4194304	930228	0%	3634

©2018 Western Digital Corporation or its affiliates. All rights reserved. Produced 3/18.

Western Digital, the Western Digital logo and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Microsoft®, Windows® and SQL Server® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. All other marks are the property of their respective owners.

DPL20182EN 20180328