



WHITE PAPER

HPE ProLiant DL380 Gen9 and HPE PCIe LE Workload Accelerator 55TB Data Warehouse Fast Track Reference Architecture

Based on the 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	3
About the HPE ProLiant DL380 Gen9	4
Fusion ioMemory SX350 PCIe Application Accelerator	4
New Data Warehouse Features in Microsoft SQL Server 2014.	4
About the Data Warehouse Fast Track Reference Architecture	5
Reference Architecture	5
General Settings.....	5
Operating System Settings	5
Windows Configuration – Power Settings.....	5
Fusion ioMemory SX350 Settings	6
Power Override.....	6
Storage Configuration.....	7
UEFI Configuration.....	7
SQL Server Settings	7
Database Configuration.....	7
TempDB Configuration	7
Memory Allocation	7
Local Security Policy	7
SQL Server 2014 Configuration Parameters.....	8
Measured Performance	8
SQL Server Data Warehouse Fast Track Certification	9
Summary	10
Bill of Materials	10
HPE ProLiant DL380 Gen9 – No High Availability Option	10

Executive Summary

This guide details the server, storage and software configurations for the HPE ProLiant DL380 Gen9 with Fusion ioMemory™ SX350 PCIe Application Accelerator devices.

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

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.

Fusion ioMemory products by SanDisk® provide the following significant benefits for data warehouses deployed on SQL Server 2014:

- I/O performance and resiliency
- Lower capital cost
- Lower operational costs
- Simplified management
- Predictable maintenance

Fusion ioMemory devices use flash like memory instead of like disk. This gives applications native access to data, delivering the lightning-fast response times that businesses depend on today.

Industries such as banking, social media, retail, transportation, healthcare, security, entertainment, and research and development, all benefit greatly from this solution. Service levels are increased, while complex and expensive storage sprawl is significantly reduced.

In addition, the Fusion ioMemory platform reduces energy consumption and total cost of ownership. The ioMemory platform integrates hardware and software to overcome the limitations of legacy architectures and specialized hardware.

Fusion ioMemory solutions significantly increase datacenter efficiency, while delivering enterprise-grade performance, reliability, availability, and manageability.

IT managers, database architects, and CTOs looking to explore and deploy data warehouses and BI applications now have added leverage – they can take advantage of SanDisk Fusion ioMemory engineering, integration and optimization to quickly build and deploy their next data warehouses.

About the HPE ProLiant DL380 Gen9



The HPE ProLiant DL380 Gen9 Server delivers the latest performance and expandability in the HPE 2P rack portfolio.

Reliability, serviceability and near continuous availability, backed by a comprehensive warranty, making it ideal for any server environment.

Designed to reduce costs and complexity, leveraging Intel's latest E5-2600 v3 processors with up to 70% performance gain, plus the latest HPE DDR4 SmartMemory supporting 1.5 TB and up to 14% performance increase. Additional support is available for 12Gb SAS and 40Gb NIC, with a broad range of graphics and compute options.

Fusion ioMemory™ SX350 PCIe Application Accelerator

The Fusion ioMemory SX350 PCIe application accelerator is the scalable capacity leader for PCIe flash solutions. The Fusion ioMemory SX350 PCIe series provides a cost-effective solution for read-intensive application workloads that include web hosting, data mining, seismic data processing, content caching, 3D animation, and CAD/CAM.

The Fusion ioMemory SX350 PCIe application accelerator is available in capacities from 1.25TB–6.4TB, with ultra-low 79µs/15µs read/write data access latency, superior reliability, and outstanding random read/write performance of up to 225K/385K IOPS. It also provides updated VSL™ (Virtual Storage Layer) software that delivers direct memory access, minimizes latency, and maximizes application throughput.

With this significant performance improvement, customers can reduce infrastructure and reduce power and cooling costs over a traditional hard disk drive infrastructure, for a lower total cost of ownership (TCO). With over 7,000 customers and over 250,000 units sold, this latest generation of PCIe application accelerators is designed to provide customers with the peace of mind that these products will perform in the field as intended.

New Data Warehouse Features in Microsoft® 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 solutions deliver an order-of-magnitude of 7x better query performance when using CCI. The CCI solution 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 – of new data while the query workload is running. This reduces the data latency from the time data is born to when it is available for querying.

About the Data Warehouse Fast Track Reference Architecture

The SQL Server Data Warehouse Fast Track reference architecture provides a scalable framework based 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 balances performance between the disk subsystem, CPU, and memory.

More information on SQL Server DWFT can be found here:

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

Reference Architecture

The following shows the configuration details for the Fusion ioMemory™ SX350 6400 and HPE ProLiant DL380 Gen9 Fast Track Data Warehouse architecture.

Server	HPE ProLiant DL380 Gen9
Operating System	Microsoft Windows Server 2012 R2 Standard Edition
CPU	Intel Xeon E5-2697 v3 @ 2.6 GHz
PCI-E Slots	Up to six PCIe 3.0 slots (two CPU and second riser card)
Drives	2 x 1.2TB SAS (OS)
RAM	256GB

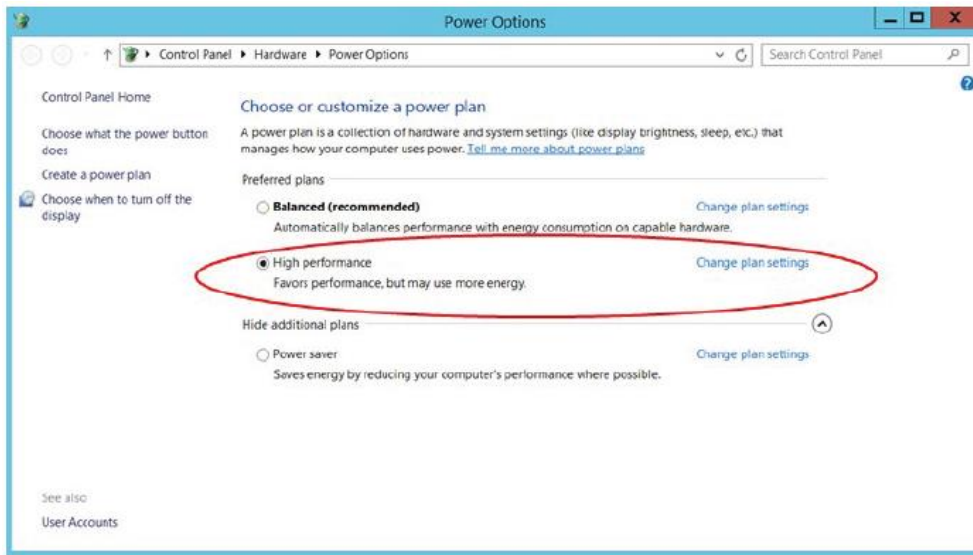
General Settings

Operating System Settings

The operating system used for this Fast Track Data Warehouse test was Microsoft Windows Server 2012 R2 Standard Edition. Standard installation steps were used to install the operating system with default values, followed by service packs and update patches.

Windows Configuration – Power Settings

The High Performance plan was chosen to reduce CPU throttling.



Fusion ioMemory™ SX350 Settings

Each device was formatted to its default capacity with a block size of 512B. The raw disk was then used as a mount point and an NTFS file system was applied using the default space and geometry, with an allocation unit size of 4096 KB.

Power Override

Enabling the power override setting on the Fusion ioMemory SX350 product line allows the device to draw up to 55 watts of power under heavy workloads and is required to achieve the performance results below. All cards must be configured at the same time and a server reboot is required for the setting to be active and persist.

Example:

```
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE SN:MW
```

where < SN > is the serial number of the card obtained from `fio-status`, and < MW > is the power in milliwatts. The following example configures the device with the given serial number to 55W:

```
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE 1410G0092:55000
```

Multiple cards must be configured with the same command:

```
fio-config -p FIO_EXTERNAL_POWER_OVERRIDE 1407G0327:55000,1504G0154:55000
```

Storage Configuration

The table below describes the storage configuration used for various devices and slots.

Slot	Device	Capacity	Mount Point	Allocation	Notes
2	Fusion ioMemory SX350	6.4TB	C:\DB\DATA	Data Files \ TempDB \ Logs	RAID 5
3	Fusion ioMemory SX350	6.4TB	C:\DB\DATA	Data Files \ TempDB \ Logs	RAID 5
5	Fusion ioMemory SX350	6.4TB	C:\DB\DATA	Data Files \ TempDB \ Logs	RAID 5
6	Fusion ioMemory SX350	6.4TB	C:\DB\DATA	Data Files \ TempDB \ Logs	RAID 5
N/A	2 x 10K SAS	1TB	C:\	OS	Mirrored

UEFI Configuration

- Hyper-threading was enabled.
- Operating mode was changed to "Maximum Performance"
- Fan Offset was set to "Increased Cooling".

SQL Server Settings

Database Configuration

A 1 TB data warehouse schema was created for benchmarking using the Fast Track toolkit. The schema used a master filegroup with 4 additional filegroups that represented the four partitions.

TempDB Configuration

Eight 10GB tempdb files were stored across four Fusion ioMemory volumes. The tempdb transaction log file was stored on the volume designated for log files.

Memory Allocation

SQL Server was allocated 118GB of the available server memory.

Local Security Policy

The SQL Server maintenance account was granted the following privileges:

- Enable Lock Pages in Memory
- Perform Volume Maintenance Tasks

SQL Server 2014 Configuration Parameters



Parameter	Setting	Description
Memory Allocation	118GB	This is the Fast Track-required value for a 2-socket, 1 TB database. Memory is deliberately constrained to enforce I/O pressure on the subsystem.
Max Degree of Parallelism	28	When SQL Server runs on a computer with more than one microprocessor or CPU, it detects the best degree of parallelism (the number of processors employed to run a single statement).
Resource Governor memory Allocation	12%	Default is 25%. This is reduced to 12% for both Row Store and Column Store to reduce the maximum memory consumed per query.
Fast Track Required Start-up Parameters	-T1117	-T1117 should be added to the startup options. This trace flag ensures even growth of all files in a file group in case <code>autogrow</code> is enabled. The standard FTDW recommendation for database file growth is to pre-allocate rather than use <code>autogrow</code> (with the exception of <code>tempdb</code>).
Optional Trace Flags	-T1118	Helps alleviate allocation bit map contention in <code>tempdb</code> by switching allocations to full extents (8 physically contiguous pages (64 KB)).

Measured Performance

During Fast Track Database Validation, Microsoft’s Reference Point tool drives multiple concurrent query workloads designed to identify bottlenecks. The tool establishes the key performance metrics in the table below.

Scan Rate Type	Scan rate	Description
Rated User Capacity	55TB	Represents the optimal Fast Track-certified data capacity of the configuration. Allows for 5:1 compression, recommended free space and applies some memory- and throughput-based limits.
Row Store Relative Throughput	186	Percentage throughput of this configuration in comparison to the FTDW reference configuration.
Column Store Relative Throughput	264	Percentage throughput of this configuration in comparison to the FTDW reference configuration.
Maximum User Data Capacity	65TB	Calculated, based on total disk capacity. Allows 5:1 compression. Factors recommended free space. Ignores the throughput limits that are applied to the Rated User Capacity.
Measured Throughput (Q/Hr/TB)	231	Number of combined benchmark queries completed during the measurement interval. Normalized to the 1TB database.

SQL Server Data Warehouse Fast Track Certification

DWFT Certification #2014-043	HP ProLiant DL380 (Gen9) Fusion ioMemory SX350 DWFT Reference Architecture			Report Date: 3/1/2016	
DWFT Rev. 5.4					
System Provider	System Name	Processor Type	Memory		
 Hewlett Packard Enterprise	HP ProLiant DL380 (Gen9)	Intel E5-2697 v3	768 GB		
Operating System		SQL Server Edition			
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition			
Storage Provider	Storage Information				
 a Western Digital brand	4 x Fusion ioMemory SX350-6400 (RAID5) 18,800 GB allocated to Data and TempDB 600 GB allocated to LOG				
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	186	264	65		
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.) (%)
186	231	4,903	6,049	5,476	83
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.) (%)
264	1,717	1,531	N/A	N/A	89
<p>The reference configuration is a 2 socket system rated for 25TB using the DWFT V4 methodology</p> <p>¹ Assumes 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>⁴ Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

Summary

Together, Hewlett Packard Enterprise and SanDisk dedicated hundreds of hours of testing to engineer the SQL Server DWFT solution to provide the most optimal reliability and performance. These series of tests pushed the HPE ProLiant DL380 Gen9 to peak performance without hardware failure of the hardware. The reliability and performance experienced during testing is what can be expected in production environments.

The same configuration meets the need of both Row Store and Column Store configurations, delivering high physical read throughput in the Row Store configuration at 4.9GB/s, and high query rates in the Column Store configuration at 1,531 Q/Hr/TB.

The Fusion ioMemory PCIe device simplifies storage configuration by reducing the importance of sequential I/O, as evangelized in previous Data Warehouse Fast Track Reference Architectures.

The HPE ProLiant DL380 Gen9 with Fusion io-Memory technology delivers superb performance, with the ability to host large data warehouses or a consolidation of data warehouses.

Fusion ioMemory products provide the following significant benefits for data warehouses deployed on SQL Server 2014:

- I/O performance and resiliency
- Lower capital cost
- Lower operational costs
- Simplified management
- Predictable maintenance

Bill of Materials

HPE ProLiant DL380 Gen9 – No High Availability Option

Qty	SKU	Description
1	719064-B21	HPE DL380 Gen9 8SFF CTO Server
1	719054-L21	HPE DL380 Gen9 E5-2697v3 FIO Kit
1	719054-B21	HPE DL380 Gen9 E5-2697v3 Kit
16	726719-B21	HPE 16GB 2Rx4 PC4-2133P-R Kit
2	718162-B21	P 1.2TB 6G SAS 10K rpm SFF (2.5-inch) SC Dual Port Enterprise 3yr Warranty Hard Drive
1	719073-B21	HPE DL380 Gen9 Secondary Riser
2	720479-B21	HPE 800W FS Plat Ht Plg Pwr Supply Kit
4	831739-B21	HPE 6.4TB Read Intensive PCIe Workload Accelerator

Specifications are subject to change. ©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. Optimus, Optimus Ascend, Guardian Technology, FlashGuard, DataGuard and EverGuard are trademarks 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).

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