

WHITE PAPER

Lenovo® System x3650 M5 with HGST Ultrastar® SAS SSDs 9TB Data Warehouse Fast Track Reference Architecture on SQL Server 2016 SP1 Standard Edition

Based on the SQL Server® 2016 Data Warehouse Fast Track Reference
Architecture (DWFTRA)

Table of Contents

Executive Summary	3
Audience	3
About the Lenovo System x3650 M5	4
HGST Ultrastar® SSD1600MM	4
New Data Warehouse Features in Microsoft® SQL Server® 2016 SP1	5
About the Data Warehouse Fast Track Reference Architecture	5
Reference Architecture	5
Server UEFI/BIOS Configuration	6
General Settings	6
Operating System Settings	6
Windows Configuration – Power Settings.....	6
Storage Configuration.....	6
Storage Pool: Physical Disks	7
Storage Pool: Virtual Disk and Volume.....	7
SQL Server Settings.....	7
Database Configuration.....	7
tempdb Configuration	7
Memory Allocation.....	7
Local Security Policy	7
SQL Server 2016 Configuration Parameters.....	8
Measured Performance: Reference Point	8
SQL Server Data Warehouse Fast Track Certification	9
Summary	10
Bill of Materials	11
Lenovo System x3650 M5	11

Executive Summary

This white paper details the server, storage, and software configuration for a Microsoft® SQL Server®-based Data Warehouse using the Lenovo® System x3650 M5 with certain HGST-brand Ultrastar® 12G SAS SSDs, which Microsoft has certified for a Fast Track Rated User Capacity of 9TB.

The goal with this configuration was to provide a hardware and software configuration that maximized savings while simultaneously providing the performance to pass Fast Track certification. This was achieved by installing a single CPU in the two-socket x3650 M5 system, using the more affordable SQL Server 2016 SP1 Standard Edition (SE) and Windows Server® 2016 Standard Edition software, and scaling the storage configuration to balance performance, capacity, and cost.

The goal of a Data Warehouse Fast Track Reference Architecture is to achieve an efficient resource balance between SQL Server data processing capability and component hardware throughput, and to provide customers standard and proven system architectures optimized for a range of enterprise data warehousing needs.

SQL Server 2016 SP1 SE exposes features which were formerly available only in Enterprise Edition (EE), such as columnstore in-memory technology, data compression, and table partitioning. These features are leveraged in this reference architecture and qualify use of SQL Server SE for Fast Track-certification.

This configuration uses Standard Edition. Compared to an Enterprise Edition configuration, *licensing savings exceed the cost of the hardware.*

The system characterized in this document delivers data warehouse performance comparable to a much more expensive system: 80% row store relative throughput, and 122% columnstore relative throughput on this SQL Server Standard Edition configuration, compared to the 2-socket reference configuration running SQL Server 2014 Enterprise Edition.

Although this configuration used only four HGST Ultrastar 1.6TB SSDs, it has a Rated User Capacity of 9TB. Adding another SSD resulted in a test rating of 12TB, which can also be obtained by using a Storage Spaces resiliency setting of “simple” instead of “parity.” Thus, a 33% greater Rated Capacity can be achieved by adding a single 1.6TB device or by providing availability beyond the storage layer (not uncommon for data warehouse stores).

This system scales easily to handle larger, more demanding data warehouse workloads by substituting a processor with more cores, populating the second socket, adding system memory, adding storage, or upgrading to SQL Server 2016 Enterprise Edition. Lenovo System x-based systems with validated Fast Track Rated User Capacities of up to 145TB have been published.

Audience

This document is for individuals (Business Intelligence database, and system architects; DBAs; report developers; IT directors) involved in decision making who are looking for guidance when designing enterprise business intelligence and analytics applications.

About the Lenovo System x3650 M5



The Lenovo System x3650 M5 Server delivers the latest performance and expandability in the LENOVO 2U rack portfolio. Reliability, serviceability, and near continuous availability, backed by a comprehensive warranty, make it ideal for any server environment.

Lenovo System servers, such as the System x3650 M5 server, feature the latest Intel® Xeon® E5-2600 v4 series (Broadwell-EP) processors. With more cores and more memory, the new M5 family is fast. The greatly increased processing power is provided by the latest E5-2600 v4 processors. Incorporating flash from HGST improves productivity through data consolidation, availability, performance, and scalability. These solid-state devices simplify DWFT storage configuration and maintenance versus the use of SANs, which demand more management overhead and are significantly more expensive. Lenovo System x3650 M5 servers include the following features:

- Two times the memory capacity of previous generation processors, with 24 DIMM sockets in the x3650 M5
- Support for 64 GB TruDDR4 Memory LRDIMMs, up to 1.5 TB of memory in the x3650 M5
- New storage technologies, such as enterprise flash from HGST

A comprehensive, easy-to-understand review of options and features is available in this authoritative reference: [Lenovo System x3650 M5 \(Machine Type 8871\) Product Guide](https://lenovopress.com/lp0068-lenovo-system-x3650-m5-machine-type-8871)
<https://lenovopress.com/lp0068-lenovo-system-x3650-m5-machine-type-8871>

HGST Ultrastar® SSD1600MM SAS SSDs

To meet today's ever-increasing needs for reliability, endurance, and performance in demanding enterprise environments, the second generation Ultrastar SSD1600MM 12Gb/s SAS SSD is a solid choice. It is the result of decades of proven enterprise storage expertise in SAS design, reliability, firmware, customer qualification, and system integration. The Ultrastar SSD1600MM SSD delivers 10 drive writes per day (DW/D) for five years, and it offers a broad range of encryption options.

The Ultrastar SSD1600MM doubles interface speed from the prior generation, providing high performance for mission-critical application workloads for hyperscale, cloud, and virtualized data centers. It offers a feature-rich, robust design, including full data path protection (T10-DIF support), temperature-based power control, SED

instant secure erase, an MTBF of 2.5 million hours, and power-fail protection. This SSD is backward-compatible with 6Gb/s SAS and offers a single firmware binary platform for seamless integration and reduced qualification times. For more information, refer to the SSD1600MM datasheet:

http://www.hgst.com/sites/default/files/resources/US_SSD1600MM_ds.pdf

New Data Warehouse Features in Microsoft® SQL Server® 2016 SP1

With the [release](#) of SQL Server 2016 SP1, Microsoft announced the availability of a common programming surface across editions. Features formerly available only in Enterprise Edition were exposed in all editions, including columnstore in-memory indexes, data compression, and table partitioning, are now available in Standard Edition. The system characterized in this document leverages these capabilities, proving out the performance by passing Fast Track certification in a scalable architecture. For more information on columnstore indexes in SQL Server 2016 SP1, see:

Developers choice: Columnstore index in Standard and Express Edition with SQL Server 2016 (SP1)

https://blogs.msdn.microsoft.com/sql_server_team/columnstore-index-standard-and-express-editions-with-sql-server-2016-sp1

The article documents that columnstore queries in Standard Edition are limited to two threads. This limitation—a reasonable compromise for substantially reduced licensing costs—did not thwart this system from achieving Fast Track certification.

About the Data Warehouse Fast Track Reference Architecture

The SQL Server Data Warehouse Fast Track reference architecture (DWFTRA) provides a scalable framework to achieve maximum performance from SMP-based servers. 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 storage subsystem, CPU, and memory.

More information on SQL Server DWFTRA can be found here:

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

Reference Architecture

The following table shows the configuration details for the Lenovo System x3650 M5 with Ultrastar SSDs Data Warehouse Fast Track Reference Architecture.

Server	Lenovo System x3650 M5
Operating System	Microsoft Windows Server 2016 Standard Edition
CPU	Intel® Xeon® E5-2667 v4 @3.2 GHz (1S/8C/16T)
PCI-E Slots	Up to nine PCIe 3.0 slots Slots 4, 5, and 9 are the fixed slots on the system board, and the remaining slots depend on the riser cards installed.
OS Drives	2x 1.2TB 10K 12Gbs SAS 2.5in HDD (mirrored)
Data Drives	4x 1.6TB Ultrastar SSD1600MM 12Gbs SAS SSD (Storage Spaces)
RAM	128GB

Server UEFI/BIOS Configuration

- TurboBoost: Power Optimized
- Operating Mode: Maximum Performance
- C-Status: Disable
- Power/Performance Bias: Platform Controlled
- Platform Controlled Type: Maximum Performance
- Hyper-Threading: Enable

General Settings

Operating System Settings

The operating system used for this Data Warehouse Fast Track test was Microsoft Windows Server 2016 Standard Edition. Typical installation steps were used to install the operating system with default values, followed by Windows Updates.

Windows Configuration – Power Settings

The High Performance plan was chosen to eliminate CPU throttling.

Storage Configuration

A mirrored pair of 10K SAS HDDs hosted the operating system.

Four 1.6TB Ultrastar SSD1600MM SSDs were configured using Storage Spaces. This technology enables Windows Server to virtualize storage by grouping industry-standard disks into storage pools and then creating virtual disks (Storage Spaces) from the available capacity in those storage pools.

Two virtual disks were carved from the physical media. SQL Server user database data and tempdb data were hosted on a virtual disk configured with simple resiliency. User database log and tempdb log files were hosted on a virtual disk configured with parity.

Note that the onboard RAID controller was disabled. The non-RAID I/O controller described in the [Bill of Materials](#) was required to provide disk consistency and facilitate storage spaces configuration.

The performance provided by Ultrastar SSDs minimizes investment costs by eliminating the need for external storage and related peripherals. SQL Server wait statistics confirmed that storage was not a performance bottleneck at any time.

Storage Pool: Physical Disks

Slot	BUS	Physical Disk	Raw Capacity	Storage Pool
3	SAS	Ultrastar SSD1600MM	1.6TB	SQLdata SQLlog
4				
5				
6				

Storage Pool: Virtual Disk and Volume

Name	Layout	Provisioning	Capacity	Storage Pool
User Data	Simple	Fixed	3.6TB	SQLdata
Tempdb Data				
User Log	Parity	Fixed	200GB	SQLlog
Tempdb Log				

SQL Server Settings

Database Configuration

A 1TB data warehouse schema was created for benchmarking using the Fast Track toolkit. The TPC™-H-like schema was comprised of eight user tables distributed across ten filegroups.

tempdb Configuration

Eight 512MB data files and a 32MB log file were configured for the tempdb. These files provided sufficient capacity to avoid growth events.

Memory Allocation

Of the 128GB of available server memory, 118GB was allocated to SQL Server. This standardized configuration is defined by the Fast Track test criteria to drive disk activity, especially during the row store query runs. *Actual production configurations may benefit from additional memory.*

Local Security Policy

The SQL Server service account was granted the following privileges:

- Enable lock pages in memory
- Perform volume maintenance tasks

SQL Server 2016 Configuration Parameters

Parameter	Setting	Description
Memory Allocation	118GB	This value is required by Fast Track protocols. Memory is deliberately constrained to drive I/O pressure on the storage subsystem.
Max Degree of Parallelism	16	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). For this configuration, it is equivalent to the logical processors available from a single socket.
Startup Parameter	-T834	Tests on this configuration leveraging this trace flag yielded statistically equivalent results compared to the default (no trace flag). When set, SQL Server uses Windows large-page memory allocations for the buffer pool. This trace flag can improve throughput rates for data warehousing workloads.

Measured Performance: Reference Point



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, for a single-socket configuration.

As noted previously, the configuration documented here has a Rated User Capacity of 9TB. However, adding an additional HGST Ultrastar 1.6TB SSD resulted in a test rating of 12TB. Likewise, this result was obtained by using a Storage Spaces resiliency setting of “simple” instead of “parity.”

The base reference configuration from which relative throughput was calculated is a 2-socket system. It is rated for 25TB and used the DWFT v4 methodology with SQL Server Enterprise Edition. The results documented here for a single processor on SQL Server Standard Edition compare favorably or exceed the performance of the reference implementation using a significantly more expensive Enterprise Edition.

Scan Rate Type	Scan rate	Description
Rated User Capacity	9TB	Represents the optimal Fast Track-certified data capacity of the configuration. It allows for 5:1 compression with 10% recommended free space. 25% capacity is reserved for tempdb, while some memory and throughput based limits are also applied.
Row Store Relative Throughput	80	Percentage throughput of this configuration in comparison to the FTDW reference configuration. This result almost matched the reference configuration, using only a single socket. The reference architecture is a 25TB dual-socket configuration.
Columnstore Relative Throughput	122	Percentage throughput of this configuration compared to the FTDW reference configuration.
Maximum User Data Capacity	11TB	Calculated, based on total disk capacity. Allows 5:1 compression. This factors in the recommended free space but ignores the throughput limits applied to the Rated User Capacity.
Row Store Measured Throughput (Q/Hr/TB)	107	Row store benchmark queries completed during the one- hour measurement interval, normalized to the 1TB database.
Columnstore Measured Throughput (Q/Hr/TB)	794	Columnstore combined benchmark queries completed during the one-hour measurement interval, normalized to the 1TB database.

Data Warehouse Fast Track Certification

DWFT Certification #2016-018	Lenovo System x3650 M5 single-socket 8-core with HGST Ultrastar SAS SSDs DWFTRA for Microsoft SQL Server 2016 SP1 Standard Edition DWFT Reference Architecture		Report Date: 5/26/2016		
DWFT Rev. 5.4					
System Provider	System Name	Processor Type	Memory		
	Lenovo System x3650 M5	Intel Xeon E5-2667 v4 3.2GHz (1/8/16)	128GB 4x32GB DDR4		
Operating System		SQL Server Edition			
Windows Server 2016 Standard Edition		SQL Server 2016 SP1 Standard Edition			
Storage Provider	Storage Information				
 a Western Digital brand	4x 1.6TB HGST Ultrastar 12Gb/s SAS SSDs for data, tempdb, & logs Storage Spaces (Data: parity resiliency; Logs: mirrored resiliency) 2x 1TB HDD for OS (mirrored)				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
9	80	122	11		
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.) (%)
80	107	2,076	2,652	2,364	99
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.) (%)
122	794	868	N/A	N/A	71
<p>The reference configuration is a 2 socket system rated for 25TB using SQL Server 2014 and 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>* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

Summary

Together, Lenovo and Western Digital employees spent hundreds of hours testing to ensure the SQL Server DWFTRA solution provided optimal reliability and performance. These tests pushed the Lenovo System x3650 M5 to peak performance without failure. The reliability and performance experienced during testing is what can be expected in production environments.

Although this configuration used only four HGST Ultrastar 1.6TB SSDs, it has a Rated User Capacity of 9TB. Adding another SSD resulted in a test rating of 12TB, which can also be obtained by using a Storage Spaces resiliency setting of “simple” instead of “parity.” Thus, a 33% greater Rated Capacity can be achieved by adding a single 1.6TB device or by providing availability beyond the storage layer (not uncommon for data warehouse stores).

The same configuration satisfied both row store and columnstore configurations, delivering high physical read throughput in the row store configuration exceeding 2.0 GB/s, and high query rates in the columnstore configuration at 794 Q/Hr/TB.

This performance was achieved using Microsoft Windows Server 2016 Standard Edition (instead of Datacenter Edition), hosting SQL Server 2016 SP1 Standard Edition (instead of Enterprise Edition). Using the Standard Editions minimized licensing costs by 3.8x. This configuration increased savings while providing the performance to achieve Fast Track certification. With this configuration, *licensing savings exceed the cost of the hardware.*

The architecture scales easily, including adding a processor with more cores, populating the second socket, adding more memory, adding additional storage, or upgrading to SQL Server 2016 Enterprise Edition.

In summary, the system characterized here offers affordability, reliability, performance, scalability all in a configuration validated by Microsoft.

Bill of Materials

Lenovo System x3650 M5

The table below lists the components of the reference architecture system. It shows the recommended options for high performance networking needed for data loading, backup and restore, and redundancy (power supplies, etc.).

Quantity	Feature Code	Component	Description
1	8871-AC1	Server	9TB DWFT: Lenovo System x3650 M5
1	ATES	Processor	Intel Xeon Processor E5-2667 v4 8C 3.2GHz 25MB Cache 2400MHz 135W
4	ATCB	Memory	32GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM
4	AS7J	SSD Flash Storage (database)	1.6TB 12G SAS 2.5" MLC G3HS Enterprise SSD
1	N2215	I/O Controller (non-RAID)	SAS/SATA HBA
1	A3PN	NIC	Mellanox® ConnectX®-3 2x40GbE / FDR IB VPI Adapter
2	AT8C	Boot Devices	1.2TB 10K 12Gbps SAS 2.5" G3HS HDD
1	A5G6		x3650 M5 8x 2.5" HS HDD Assembly Kit (Single RAID)
1	5977		Select Storage devices; no IBM® configured RAID required
2	ATDY		System x3650 M5 24x 2.5" Base without Power Supply BDW
2	A5EW	Power Supply	System x 900W High Efficiency Platinum AC Power Supply
1	9206		No preload specified
1	6311		2.8 m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable
1	A5FW		Lenovo System Gen-II Universal Slides Kit
1	ATG5		Lenovo System x3650 M5 WW Packaging
1	A5R6		System x3650 M5 PCIe Riser 2 (2 x8 FH/FL + 1 x8 FH/HL Slots)
1	ATE4		System x3650 M5 Planar BDW
1	A59E		Lenovo System x3550 M5 Label GMB
1	A5FP		System x3650 M5 PCIe Riser 1 (2 x8 FH/FL + 1 x8 FH/HL Slots)
	01GX-366		3YR Tech Install Parts 24x7x4

Options and features may be reviewed in this authoritative reference:

Lenovo System x3650 M5 (Machine Type 8871) Product Guide

<https://lenovopress.com/lp0068-lenovo-system-x3650-m5-machine-type-8871>

Specifications are subject to change. ©2017 Western Digital Corporation or its affiliates. All rights reserved. Western Digital, SanDisk, the SanDisk logo, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Lenovo® is a trademark of Lenovo in the United States, other countries, or both. Microsoft, Windows, SQL Server and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other marks are the property of their respective owners. Mellanox and ConnectX are registered trademarks of Mellanox Technologies, Ltd. IBM is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide. 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.

DPL1701EN 20170628