



WHITE PAPER

Lenovo® System x3650 M5 with Optimus Ascend™ SAS SSDs from SanDisk® – 7TB Data Warehouse Fast Track Reference Architecture

Based on the 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 Lenovo® System x3650 M5	3
Optimus® Product Family of SAS SSDs	4
Guardian Technology™ Platform	4
New Data Warehouse Features in Microsoft® SQL Server® 2014.	5
About the Data Warehouse Fast Track Reference Architecture	5
Reference Architecture	5
General Settings.....	6
Operating System Settings.....	6
Windows Configuration – Power Settings	6
Optimus Ascend Configuration.....	6
ServeRAID M5210 Controller.....	6
Storage Configuration.....	7
UEFI Configuration.....	7
SQL Server Settings	7
SQL Server 2014 Configuration Parameters	7
Measured Performance	8
SQL Server Data Warehouse Fast Track Certification	9
Summary	10
Bill of Materials – North America Region	11
Lenovo System x3650 M5	11
Bill of Materials – EMEA Region	12
Lenovo System x3650 M5	12

Executive Summary

This guide details the server, storage, and software configurations for the Lenovo System x3650 M5 with Optimus Ascend SSDs from SanDisk (6Gb/s SAS, 400GB), which Microsoft has validated for a Fast Track-rated user data capacity of 7TB.

By installing a single CPU in this two-socket system, and with 16 small form factor (SFF) drive bays, this 7TB architecture offers scalable capacity and performance, while minimizing software licensing costs.

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.

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 the 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.

About the Lenovo® System x3650 M5



With the powerful, versatile new 2U two-socket Lenovo System x3650 M5 rack server, you can run even more workloads, 24 x 7, and gain faster business insights. Integrated with the Intel® Xeon® processor E5-2600 v3 product family and industry-leading two-socket storage capacity, the x3650 M5 fast-forwards your business. You can select from an impressive array of storage configurations (up to 28 small form factor [SFF] drive bays) that optimize diverse workloads from cloud to big data.

Combining balanced performance and flexibility, the x3650 M5 is a great choice for small and medium businesses up to the large enterprise. It can provide outstanding uptime to keep business-critical applications and cloud deployments running safely. Ease of use and comprehensive systems management tools help make deployment easier. Outstanding reliability, availability, and serviceability (RAS) and high-efficiency design improve your business environment and help save operational costs.

<https://lenovopress.com/tips1193-lenovo-system-x3650-m5>

Optimus® Product Family of SAS SSDs

Available in capacities of 100GB – 2TB, Optimus SAS SSDs from SanDisk provide

- Random read performance up to 95K IOPS
- Random write performance up to 40K IOPS
- Sequential read data transfer up to 550MB/s
- Sequential write data transfer up to 540MB/s

The Optimus SAS SSD portfolio has been designed to offer a range of performance and capacity solutions for a variety of applications that include mixed-use, read-intensive, and write-intensive workloads. Optimus SAS SSDs range from one to 45 full random drive writes per day (DWPD). The Optimus SAS SSD has a five-year warranty. For more information on Optimus SAS SSDs, refer to the product datasheet:

https://www.sandisk.com/content/dam/sandisk-main/en_us/assets/resources/enterprise/data-sheets/optimus-sas-ssd-product-family-datasheet.pdf

Guardian Technology™ Platform

The Optimus SAS SSD product family is powered by SanDisk's proprietary Guardian Technology Platform, a comprehensive suite of enterprise features comprised of FlashGuard™, DataGuard™ and EverGuard™ technologies.

- FlashGuard Technology combines Aggregated Flash Management, which treats all flash elements as a single system, and Advanced Signal Processing. The Advanced Signal Processing feature dynamically adjusts flash parameters throughout the life of the SSD, to reliably extract significantly more life from commercial-grade MLC flash. This makes it suitable for mixed-use, read-intensive, and write-intensive application workloads.
- DataGuard Technology features full data path protection, safeguarding user data from corruption along all data paths in the SSD.
- EverGuard Technology protects against loss of user data in the event of unexpected power interruptions. It uses a third-generation backup power circuit design and high-reliability discrete capacitors. Optimus SAS SSDs are designed for I/O intensive, mission-critical, 24/7 applications.

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 Optimus SAS SSD engineering, integration and optimization to quickly build and deploy their next data warehouses.

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. The SQL Server Data WarehouseFast Track eliminates the complexity of sizing servers with data warehouses by providing data consumption rates that properly balance 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 Optimus Ascend 7 TB Fast Track Data Warehouse.

Server	Lenovo System x3650 M5
Operating System	Microsoft Windows Server 2012 R2 Standard Edition
CPU	Intel Xeon E5-2667v3 3.2GHz (1/8/16) (single socket)
Primary Data Drives	10 x 400GB Optimus Ascend SAS SSD presented via Storage Spaces, with an OS-configured RAID 5.
RAID Controller	ServeRAID M5210 SAS/SATA Controller
OS Drives	2 x 1TB 10K SAS Drives (Mirrored)
RAM	128GB

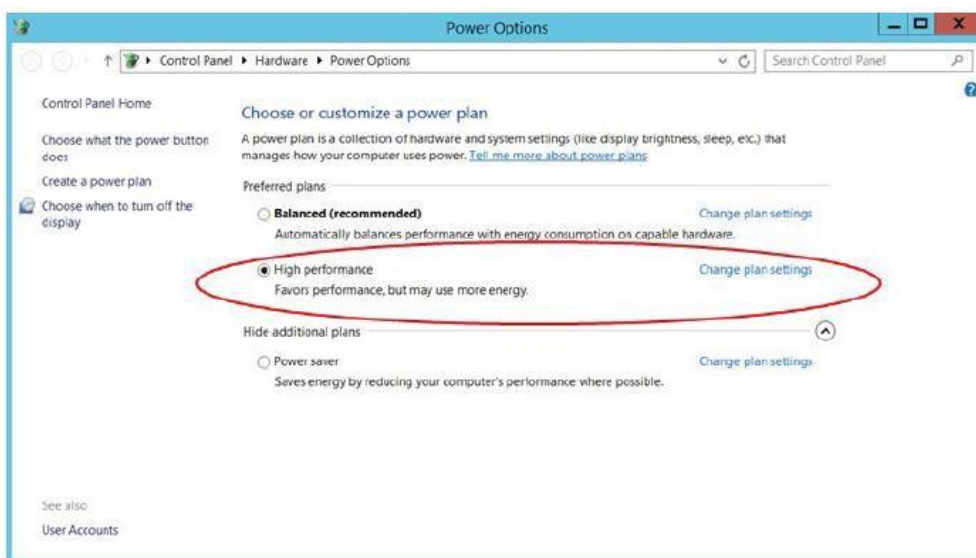
General Settings

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 OS with default values, followed by service packs and update patches.

Windows Configuration – Power Settings

The High Performance plan was chosen to reduce CPU throttling, as shown below.



Optimus Ascend Configuration

Each Optimus Ascend SAS SSD was attached to the ServeRAID M5210 Controller and configured as a single Drive Group (RAID 5) set, with a single virtual drive presented to the Windows OS.

ServeRAID M5210 Controller

The ServeRAID M5210 SAS/SATA controller is a PCIe 3.0, half-size, half-height RAID controller based on the LSISAS3108 PCI Express-SAS/SATA I/O Processor chip.

The controller controls eight internal SAS/SATA ports through two SFF-8643 4i internal mini-SAS HD connectors. The controller integrates eight high-performance SAS/SATA PHY chips and a PCI Express bus master DMA core. Each of the eight PHYs is capable of 12.0Gb/s SAS link rates and 6.0Gb/s SATA III link rates.

Storage Configuration

The following configuration was chosen.

Slot	Device	Capacity	Mount Point	Allocation	Notes
N/A	Optimus Ascend	2.2TB	C:\mounthead\	Data Files\TempDB\LOG	RAID5
N/A	10K SAS HDD	1.0TB	C:\	OS	Mirrored

UEFI Configuration

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

SQL Server Settings

- **Database Configuration:** A 1TB data warehouse schema was created for benchmarking using the Fast Track toolkit. The schema used a single primary filegroup (8 files).
- **TempDB Configuration:** In total, eight x 10GB tempdb files were created.
- **Memory Allocation:** SQL Server was allocated 118GB of the available servermemory.
- **Local Security Policy:** The SQL Server maintenance account was granted the privileges of “Enable Lock Pages in Memory” and “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, 1TB database. Memory is deliberately constrained to enforce I/O pressure on the subsystem.
Max Degree of Parallelism	8	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 autogrow is enabled. The standard FTDW recommendation for database file growth is to pre-allocate rather than use autogrow (with the exception oftempdb).
Optional Trace Flags	None	No additional Trace Flags were used.

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	7TB	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	58	Percentage throughput of this configuration in comparison to the FTDW reference configuration. The reference configuration value is 100, but this is achieved with both sockets populated.
Column Store Relative Throughput	87	Percentage throughput of this configuration in comparison to the FTDW reference configuration. The reference configuration value is 100, but this is achieved with both sockets populated.
Maximum User Data Capacity	9TB	Calculated, based on total disk capacity. Allows 5:1 compression. Factors recommended free space and ignores the throughput limits that are applied to the Rated User Capacity.
Measured Throughput (Q/Hr/TB)	64	Number of combined benchmark queries completed during the measurement interval. Normalized to the 1TB database.

SQL Server Data Warehouse Fast Track Certification

DWFT Certification #2014-047	Lenovo x3650 M5 (Single Socket) with Optimus Ascend from SanDisk 7TB Reference Architecture for Microsoft SQL Server 2014		Report Date: 3/21/2016		
DWFT Rev. 5.4	DWFT Reference Architecture				
System Provider	System Name	Processor Type	Memory		
	Lenovo x3650 M5	E5-2667 v3 3.2GHz (1/8/16)	128GB		
Operating System		SQL Server Edition			
Windows Server 2012 R2		SQL Server 2014 Enterprise Edition			
Storage Provider	Storage Information				
	10 x 400GB Optimus Ascend from SanDisk SSD allocated to Storage Spaces 2820GB allocated to Data & TempDB. 300GB (RAID 5) 2 x 1TB 10K SAS Drives (RAID 1) OS				
Primary Metrics					
Rated User Data Capacity ¹ (TB)	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹ (TB)		
7	58	87	9		
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.) (%)
58	64	1,551	1,874	1,713	76
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.) (%)
87	564	491	N/A	N/A	76
<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>* Reported metrics are based on the qualification configuration which specifies database size and SQL Server memory.</p>					

Summary

Together, Lenovo and SanDisk dedicated hundreds of hours of testing to engineer this SQL Server DWFT solution to provide the most optimal reliability and performance. These series of tests pushed the Lenovo System x3650 M5 to peak performance without 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 1.5GB/s, and high query rates in the Column Store configuration at 491 Q/Hr/TB.

The Lenovo System x3650 M5 with Optimus Ascend SSD technology from SanDisk delivers superb performance, with the ability to host larger data warehouses or consolidate multiple data warehouses.

Bill of Materials – North America Region

Lenovo System x3650 M5

Qty	Feature Code	Description
1	8871KEU	X3650 M5 MLK, Xeon 8C E5-2620 v4 85W 2.1GHz/2133MHz/20MB, 1x16GB
7	46W0829	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM
1	00FK658	System x3650 M5 Rear 2x2.5in HDD kit
8	00AJ212	400GB SAS 2.5in MLC G3HS Enterprise SSD
2	00WG685	300GB 10K 12Gbps SAS 2.5in G3HS HDD
1	46C9110	ServeRAID M5210 SAS/SATA Controller
1	44T1370	Broadcom NetXtreme 2x10GbE Base T Adapter
1	00FK936	System x 900W High Efficiency Platinum AC Power Supply
1	90Y3901	Integrated Management Module Advanced Upgrade
1	00YD002	System x3650 M5 LCD Light Path Upgrade Kit
1	39Y7937	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable

Bill of Materials – EMEA Region

Lenovo System x3650 M5

Qty	SKU	Description
1	5462AC1	8coreTopSeller : Lenovo System x3650 M5
1	A1ML	Integrated Management Module Advanced Upgrade
2	A5EU	System x 750W High Efficiency Platinum AC Power Supply
1	A483	Populate and boot from rear drives
1	A5GH	System x3650 M5 Rear 2x 2.5" HDD Kit (Independent RAID)
1	A483	Populate and Boot From Rear Drives
1	A5EY	System Documentation and Software-US English
1	A5FV	System x Enterprise Slides Kit
1	A5EA	System x3650 M5 Planar
8	A5B7	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM
1	A5FH	System x3650 M5 Agency Label GBM
1	A5FM	System x3650 M5 System Level Code
1	9206	No Preload Specify
6	A4C2	HDD Filler ASM Gen3
1	A4VH	Lightpath LCD Op Panel
1	A5FZ	System x3650 M5 Riser Filler
2	A5V6	System x3650 M5 System Fan Filler
1	ASQA	System x3650 M5 Rear 2x 2.5" HDD Label (Independent RAID-Riser1)
2	A2HP	Configuration ID 01

1	A5FT	System x3650 M5 Power Paddle Card
1	A5G1	System x3650 M5 EIA Plate
1	A5H0	2U Bracket for Broadcom NetXtreme 2x10GbE BaseT Adapter
1	A5V4	System x3650 M5 Right EIA
1	A5FC	System x3650 M5 WW Packaging
1	A5G4	System x3650 M5 ODD Filler
1	A5DL	Intel Xeon Processor E5-2667 v3 8C 3.2GHz 20MB Cache 2133MHz 135W
1	5977	Select Storage devices - no configured RAID required
1	A5GF	x3650 M5 16x 2.5" HS HDD Assembly Kit (Single RAID)
1	A3YY	N2215 SAS/SATA HBA
1	A45W	ServeRAID M1215 SAS/SATA Controller
1	A5FD	System x3650 M5 2.5" Base without Power Supply
1	A5FY	System x3650 M5 2.5" ODD/LCD Light Path Bay
2	A4TL	300GB 10K 6Gbps SAS 2.5" G3HS HDD
10	A4UB	400GB SAS 2.5" MLC G3HS Enterprise SSD
1	A5G3	System x3650 M5 2.5" ODD Bezel with LCD Light Path
1	A5GZ	Broadcom NetXtreme 2x10GbE BaseT Adapter
1	5374CM1	5374CM1 : Configuration Instruction
1	A5M2	ServeRAID M1215 SAS/SATA Controller Upgrade Placement
1	A2HP	Configuration ID 01
1	A2JX	Controller 01
1	5374CM1	5374CM1: Configuration Instruction

1	A2HP	Configuration ID 02
1	A46U	N2215 SAS/SATA HBA Replacement
1	A2JY	Controller 02

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). 5096EN 20160615

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