

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB

Data Warehouse Fast Track Solution

Data Warehouse problem and a solution

The rapid growth of technology means that the amount of available data and the ability to collect that data increased to a level unthinkable as little as five years ago. As the volume and velocity of data increased, however, extracting meaningful insight in a timely manner became more complex. Therefore, opportunities are being missed and effort is being wasted. To compete, businesses in the 21st century are demanding the tools to derive true value from their data.

This Microsoft Data Warehouse Fast Track (DWFT) configuration for SQL Server 2016 improves time-to-value for data warehousing needs with a new scalable architecture. This solution in the Lenovo portfolio uses the high performance System x3850 X6 server combined with Intel P3700 NVMe PCIe flash adapters to solve SQL database warehouse needs up to 45 TB in size.

Microsoft Data Warehouse Fast Track program makes it easy to reduce costs, save time, and reduce risk with reliable, pretested hardware and best practices for data warehousing. This solution features the following highlights:

- Reduce time to value with pretested hardware configurations.
- Reduce hardware testing and reduce tuning immediately.
- Reduce total cost of ownership through better price and performance, rapid deployment, and advanced hardware.
- Optimize performance with pretested System x3850 X6 hardware configurations.

The configuration listed in this document has a Fast Track RowStore rating of I/O 10,981 MB/s and ColumnStore throughput of 3,015 Queries/Hr/TB.

Enterprise data warehouse with faster time-to-value

DWFT for SQL Server 2016 for System x offerings are methodically tested and tuned to save you months of configuration, setup, testing, and tuning. With these offerings from Lenovo, you can now complete the following tasks:

- Buy all the hardware that you need from only one vendor including servers, storage, and networking.
- Build, tune, and deploy with confidence by using established data warehouse best practices.
- Select from different levels of performance, scalability, and price to suit your business needs.
- Choose from 4 to up to 120 Intel Xeon processor cores.
- Run targeted query workloads that are patterned for large sequential data sets rather than small random transactions.
- Eliminate bottlenecks with optimized rapid data reads and query aggregations

HIGHLIGHTS

- First certified SQL 2016 DWFT configuration
- Advanced 45TB SQL solution from Lenovo
- Balanced and Optimized configuration
- Build, tune and deploy with confidence using established data warehouse best practices
- Eliminate bottlenecks with optimized rapid data reads and query aggregations
- Certified by Microsoft
- Reduced time to value



CONFIGURATION BRIEF

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB



Microsoft SQL Server 2016

Microsoft® SQL Server 2016 has made significant improvements in data warehousing technologies and performance, including column-store features as well as many other improvements. Column-store indices offer great advantages over traditional row stores for analytics and data warehousing queries. They are ideally suited for the star schemas, and tables with billions of rows which are commonly seen. Among their advantages for analytics are:

- **Up to 10X compression in data size** - Data warehouses are very large by nature, and the compression offered by column store index technologies offers both space and cost savings, but also significantly increased performance, due to the dramatically reduced IO requirements given by the compression, coupled by the ability to only scan the specific columns required by each query. Compression also reduces the amount of memory required to hold a given number of rows from the source data warehouse.
- **Additional Indices** - SQL Server 2016 adds the capability to add additional (B-Tree) indices to column store-based tables, which enables efficient single-row lookup.

In addition to these architectural features, query processing in column-store indices is further optimized in the following ways:

- **Operator Pushdown** - Pushdown refers to moving both filter and aggregation query operations closer to the data, so that many of the filters and calculations can be done in the scan operators, dramatically reducing the volume of data which needs to be handled further on in query processing.
- **Batch Mode Processing** - SQL Server 2016 includes enhancements in batch-mode processing which processes many rows at a time rather than serially doing calculations on each individual row. These batch operations are further optimized by leveraging Single Instruction Multiple Data (SIMD) vector processing CPU instructions in the Intel® architectures.

Configuration tested and certified

This configuration features the following main components:

- Server: Lenovo System x3850 X6
- Processors: Four Intel Xeon E7-8890 v4 24-core 2.2 GHz
- Memory: Ninety-six 16GB DDR4 DIMMs for a total 1,536 GB memory
- Storage: Eight P3700 2.0TB NVMe Enterprise Flash Adapters for data and tempdb
- OS Storage: Two 300 GB SAS HDDs
- Logging: Four 960 GB SSDs for log

CONFIGURATION BRIEF

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB

Powered by System x3850 X6 and Intel P3700 2.0TB NVMe Enterprise Performance Flash Adapters

The X6 server features Intel Xeon processors. With more cores and more memory, the X6 family is designed to be fast from the ground up. Every subsystem is tuned to maximize performance.

- Harness greatly increased processing power, with Xeon processor E7-8800 v4
- Get three times the memory capacity of previous platform with 96 DIMM sockets and 6TB memory in the x3850 X6 with 64GB DIMMs.
- New storage technologies, such as the P3700 2.0TB NVMe Enterprise Performance Flash Adapter for Lenovo System servers, which closely align the performance of storage with the power of the processors.

DWFT for SQL Server 2016 with System x3850 X6 features P3700 2.0TB NVMe Enterprise Performance Flash Adapters designed to improve productivity through data consolidation, availability, performance and scalability. These solid-state devices simplify DWFT storage configuration and maintenance versus the use of a SAN, which has more parts to maintain and manage.

Lenovo X6 servers continue to lead the way as the shift toward mission-critical scalable databases, business analytics, virtualization, enterprise applications and cloud- computing applications accelerates.

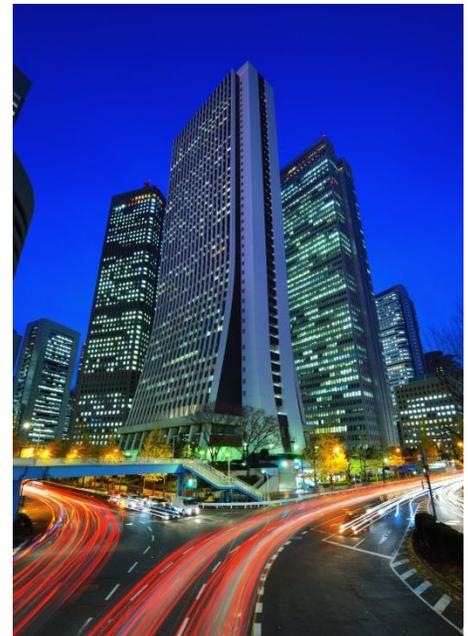
- Fast: x3850 X6 features outstanding compute performance and sophisticated memory technologies, providing lower latencies, faster response times and smarter data management in real time.
- Agile: Modular design for multiple generations of processors, providing a unique “pay-as-you-grow” capability and “fit-for-purpose” design for the lowest total cost of ownership.
- Resilient: X6 platforms maximize application uptime, provide fault tolerance and promote easy integration in virtual environments.

Best practices for Data Warehouse Fast Track

For a balanced and optimized Data Warehouse configuration:

- Configure UEFI settings to set Memory mode to Independent.
- Configure UEFI settings to set Operating mode to Maximum performance.
- Configure high availability for the OS with 2-disk Raid-1.
- Configure high availability for the log drive with 2-disk Raid-1 or Raid-10 with more disks based on performance needs.
- Data files and tempdb can be on Raid 0 drives. Spread data and tempdb files evenly across all data drives for even performance.
- Configure more than one tempdb files; at least one file per data drive.
- Enable lock pages in memory option using Windows Group policy tool to prevent paging of data.
- If the server is dedicated to data warehousing,
 - Set processor affinity for SQL Server to use all the processors in the system.
 - Set SQL Server Maximum Server Memory to 90% of the total memory available on the server.
 - Add –E and optionally –T834 to SQL Server Startup parameters.

Lenovo System x3850 X6 Mission Critical Rack Server with and Intel P3700 NVMe PCIe adapters



CONFIGURATION BRIEF

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB

Reference Architecture Certification

DWFT Certification #2016-001	Lenovo 45TB with System x3850 X6 DWFT Reference Architecture		Report Date: 6/6/2016		
DWFT Rev. 5.4					
System Provider	System Name	Processor Type		Memory	
	Lenovo System x3850 X6	Intel Xeon E7-8890 v4 2.2 GHz (4/96/192)		1536 GB	
Operating System			SQL Server Edition		
Windows Server 2012 R2			SQL Server 2016 Enterprise Edition		
Storage Provider	Storage Information				
	2x 300 GB SAS HDDs for OS (RAID 1) 8x Intel P3700 2.0 TB NVMe Enterprise Flash adapters for data and tempdb 4x 960 GB SSDs for log (RAID 10)				
Primary Metrics					
Rated User Data Capacity ¹	Row Store Relative Throughput ²	Column Store Relative Throughput ³	Maximum User Data Capacity ¹		
(TB)			(TB)		
45	373	464	51		
Row Store					
Relative Throughput ²	Measured Throughput	Measured Scan Rate Physical	Measured Scan Rate Logical	Measured I/O Throughput	Measured CPU (Avg.)
	(Queries/Hr/TB)	(MB/Sec)	(MB/Sec)	(MB/Sec)	(%)
373	383	10,042	11,920	10,981	71
Column Store					
Relative Throughput ²	Measured Throughput	Measured Scan Rate Physical	Measured Scan Rate Logical	Measured I/O Throughput	Measured CPU (Avg.)
	(Queries/Hr/TB)	(MB/Sec)	(MB/Sec)	(MB/Sec)	(%)
464	3,015	1,880	N/A	N/A	76
<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>					

CONFIGURATION BRIEF

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB

Bill of materials

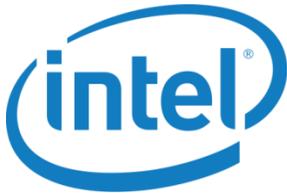
Feature code	Description	Quantity
6241-AC1	45TBDWFT: Lenovo System x3850 X6	1
A4A2	X6 Half-length I/O Book	2
A4A6	4x 2.5" HDD Riser	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	4
A3YZ	ServeRAID M5210 SAS/SATA Controller	1
A40R	Intel I350-T4 ML2 Quad Port GbE Adapter	1
AT7M	Intel P3700 2.0TB NVMe Enterprise Performance Flash Adapter	8
A4A1	X6 Storage Book	1
A4C2	HDD Filler ASM GEN 3	2
9206	No Preload Specify	1
A4VH	Lightpath LCD Op Panel	1
A4AA	Rail Kit	1
A4QX	Rating label for 1400W PS	4
A4C1	2U bracket for low profile-internal-storage adapter	1
A4A4	Mid-plane for 4U Chassis	1
A4BP	Short SAS cable to planar	2
ASFD	Labels GBM	1
A2HP	Configuration ID 01	1
ASFB	x3850/x3950 X6 I/O Planar III	1
A4AB	System 4U Packaging - WW	1
6134	Unique SBB for AC1/MC1 models	1
5977	Select Storage devices - no configured RAID required	1
A5B7	16GB TruDDR4 Memory (2Rx4, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM	96
AT89	300GB 10K 12Gbps SAS 2.5" G3HS HDD	2
AT8U	960GB Enterprise Entry SATA G3HS 2.5" SSD	4
AS91	Addl X6 DDR4 Compute Book Intel Xeon Processor E7-8890 v4 24C 2.2GHz	3
AS8B	X6 DDR4 Compute Book Intel Xeon Processor E7-8890 v4 24C 2.2GHz	1
ASMH	x3850 X6 4U Chassis	1
ASFA	System Documentation and Software-US English refresh	1
A54D	1400W HE Redundant Power Supply for altitudes >5000 meters	4



CONFIGURATION BRIEF

Lenovo Database Configuration for Microsoft SQL Server 2016 – 45TB

Mission Critical support with the Lenovo x3850 X6 server and Intel P3700 NVMe PCIe flash adapters



Why P3700 NVMe PCIe solid state storage from Intel

The Intel® P3700 NVMe PCIe Solid-State Drive brings extreme data throughput directly to Intel® Xeon® processors with up to six times faster data transfer speed than 6 Gbps SAS/SATA SSDs. The performance of a single Intel P3700 Series, can replace the performance of seven SATA SSDs aggregated through a host bus adapter.

Why Lenovo System servers for Microsoft SQL DWFT

Lenovo offers a wide range of servers and options. The Lenovo reference configurations for DWFT for SQL Server bring together the right mix of technology and software. The configurations integrate the latest powerful Lenovo System rack and enterprise servers, robust Lenovo Storage options, and the data warehouse capabilities of SQL Server 2016 Enterprise Edition.

Why Lenovo

Lenovo is a leading provider of x86 servers for the data center. Featuring rack, tower, blade, dense and converged systems, the Lenovo server portfolio provides excellent performance, reliability and security. Lenovo also offers a full range of networking, storage, software, solutions, and comprehensive services supporting business needs throughout the IT lifecycle. With options for planning, deployment, and support, Lenovo offers expertise and services needed to deliver better service-level agreements and generate greater end-user satisfaction.

For More Information

To learn more about the Lenovo Database Configuration for Microsoft SQL Server 2016 – 45 TB solution, contact your Lenovo Business Partner or visit:

<http://shop.lenovo.com/us/en/systems/solutions/database/>



© 2016 Lenovo. All rights reserved.

Availability: Offers, prices, specifications and availability may change without notice. Lenovo is not responsible for photographic or typographical errors. **Warranty:** For a copy of applicable warranties, write to: Lenovo Warranty Information, 1009 Think Place, Morrisville, NC, 27560, Lenovo makes no representation or warranty regarding third-party products or services.

Trademarks: Lenovo, the Lenovo logo, System x, ThinkServer are trademarks or registered trademarks of Lenovo. Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, the Intel logo, Xeon and Xeon Inside are registered trademarks of Intel Corporation in the U.S. and other countries. Other company, product, and service names may be trademarks or service marks of others. **CRN: DBSSQLX4562**

06/2016