

Lenovo Database Configuration

for Microsoft SQL Server DWFT – 90TB

Reduce time to value with pretested hardware configurations

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 2017 improves time-to-value for data warehousing needs with a new scalable architecture. This solution in the Lenovo portfolio uses the high performance Lenovo ThinkSystem SR950 server combined with Lenovo NVMe Enterprise Mainstream Flash Adapters to solve SQL database warehouse needs up to 90 TB in size.

Enterprise data warehouse with faster time-to-value

DWFT for SQL Server 2017 for Lenovo solution 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
- Pre-optimized system tuned and tested with Microsoft certification and deploy with confidence for your demanding data warehouse performance needs
- Select from different levels of performance, scalability, and price to suit your business needs
- Run targeted query workloads patterned for large sequential data sets rather than small random transactions
- Eliminate bottlenecks with optimized rapid data reads and query aggregations

Highlights

- Reduce time to value with pretested hardware configurations
- Microsoft Certified solution with detailed performance data
- Reduce TCO through better performance, rapid deployment and advanced hardware
- Optimize performance with pretested ThinkSystem SR950 hardware configurations
- Consolidate storage and match IT investment-to-information-value with Lenovo NVMe Enterprise Mainstream Flash Adapters



MICROSOFT SQL SERVER 2017

Microsoft® SQL Server 2017 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 – Microsoft SQL Server 2017 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 – Microsoft SQL Server 2017 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.

This configuration features the following main components:

- Server: Lenovo ThinkSystem SR950
- Processor: Intel Xeon Platinum 8180 28C 2.5GHz
- Memory: 1536 GB of TRUDDR4 memory
- Storage: Eight 3.84TB Lenovo NVMe Enterprise Mainstream Flash Adapters for data and tempdb
- OS Storage: Two 480 GB SATA SSDs for the operating system (RAID 1)
- Logging: Six 800 GB SAS SSDs for log (RAID 10)
- Software:
 - Microsoft Windows Server 2016 Standard Edition
 - Microsoft SQL Server 2017 Enterprise Edition

This data warehouse solution with Microsoft SQL Server 2017 Enterprise Edition features the Lenovo NVMe Enterprise Mainstream Flash Adapters. These flash adapters improve productivity through data consolidation, availability performance, and scalability. Solid-state adapters simplify DWFT storage configuration and maintenance versus the use of SAN which has more parts to maintain and manage.



Lenovo PX04PMC NVMe Enterprise Mainstream flash adapters are engineered for greater performance and endurance in a cost-effective design

Best practices for Data Warehouse Fast Track

For a balanced and optimized Data Warehouse configuration:

- 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.





Solution Benefits

- **Advanced 90 TB SQL DWFT solution from Lenovo**
- **Balanced and optimized configuration**
- **Features High performance Lenovo NVMe storage**
- **Certified by Microsoft**
- **Reduced time to value**



| Bill of materials | | |
|--------------------------|---|-----------------|
| Feature code | Description | Quantity |
| 7X12-CTO1WW | ThinkSystem SR950 - 3yr Warranty | 1 |
| AUNY | SR950 2S to 4S Base | 1 |
| AX6G | Intel Xeon Platinum 8180 28C 205W 2.5GHz Processor | 4 |
| AUND | ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM | 48 |
| AUN6 | ThinkSystem 2.5" SAS/SATA 2x2 Bay Backplane Kit | 2 |
| AUNG | ThinkSystem RAID 530-8i PCIe 12Gb Adapter | 1 |
| AUMD | ThinkSystem 2.5" PM1635a 800GB Mainstream SAS 12Gb Hot Swap SSD | 6 |
| AUM9 | ThinkSystem 2.5" PM863a 480GB Entry SATA 6Gb Hot Swap SSD | 2 |
| AUN2 | ThinkSystem SR950 (4) x8 PCIe Riser | 2 |
| AUN4 | ThinkSystem SR950 (2) x8 PCIe Riser | 1 |
| AUKG | ThinkSystem 1Gb 2-port RJ45 LOM | 1 |
| AUVY | ThinkSystem HHHL PX04PMC 3.84TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter | 8 |
| AUMY | SR950 Lift Handles | 1 |
| AVNV | High Efficiency 1100W Power Supply for SR950 | 4 |
| 6311 | 2.8m, 10A/100-250V, C13 to C14 Jumper Cord | 4 |
| AUMT | SR950 Main I/O Planar | 1 |
| AUN0 | ThinkSystem SR950 2-CPU, 24 DIMM Compute System Board | 2 |
| AUMW | Front operator panel with LCD display | 1 |
| A4AA | Rail Kit | 1 |
| 5977 | Select Storage devices - no configured RAID required | 1 |
| AUK7 | Feature Win Server 2016 Configuration (TPM 2.0 and Secure Boot) | 1 |
| B0ML | Feature Enable TPM on MB | 1 |
| AX3Y | 12 Hard Drive Maximum System | 1 |
| AULF | Front CPU Bulkhead Filler | 1 |
| AULE | Two Slot PCIe Riser Bulkhead Filler | 1 |
| AVPV | Drive ID Label for SAS/SATA/SSD | 1 |
| AVJ2 | ThinkSystem 4R CPU HS Clip | 4 |
| AXCR | CPU Heatsink for Highest Wattage CPUs | 4 |
| AUPH | Middle SAS Cable to Controller | 1 |
| 9206 | No Generic Preload Specify | 1 |
| AULN | SR950 Lenovo Bezel Label | 1 |
| AUMR | System Documentation | 1 |
| AVKJ | ThinkSystem 2x2 Quad Bay Gen4 2.5" HDD Filler | 1 |
| AUL9 | Short SAS Cable to Controller | 1 |
| AUVE | PHY Card Installation Bracket and Screw | 1 |
| AUN7 | SR950 WW Packaging | 1 |
| AULK | SR950 Labels | 1 |
| 5AS7-A02047 | Hardware Installation Server (Business Hours) | 1 |
| 5PS7-A01828 | Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData | 1 |

Configuration Certification

| | | | | | |
|--|---|---|---|--|-----------------------------------|
| DWFT Certification #2017-003 | Lenovo ThinkSystem SR950 - 90TB DWFT Reference Architecture | | | Report Date: 10/08/2017 | |
| DWFT Rev. 5.4 | | | | | |
| System Provider | System Name | Processor Type | | Memory | |
|  | Lenovo ThinkSystem SR950 | Intel Xeon Platinum 8180 2.5GHz (4/112/224) | | 1536 GB | |
| Operating System | | | SQL Server Edition | | |
| Windows Server 2016 Standard Edition | | | SQL Server 2017 Enterprise Edition | | |
| Storage Provider | Storage Information | | | | |
|  | 8x 3.84 TB NVMe SSD for data and tempdb 2x 480 GB SSD for OS (RAID 1) 6x 800 GB SSD for log (RAID 10) | | | | |
| Primary Metrics | | | | | |
| Rated User Data Capacity ¹ (TB) | Row Store Relative Throughput ² | Column Store Relative Throughput ³ | Maximum User Data Capacity ¹ (TB) | | |
| 90 | 501 | 720 | 102 | | |
| 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.) (%) |
| 501 | 611 | 12,358 | 17,156 | 14,757 | 68 |
| 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.) (%) |
| 720 | 4,679 | 2,421 | N/A | N/A | 63 |
| 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. | | | | | |

High performance
and high capacity
database solution
with the Lenovo
ThinkSystem
SR950 server and
Lenovo NVMe
Enterprise Flash
Adapters



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 DWFT – 90TB, contact your Lenovo Business Partner or visit:

www.lenovo.com/systems/solutions

For more information on the Lenovo NVMe Mainstream Flash Adapters, visit:

<https://lenovopress.com/lp0627.pdf>

© 2017 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, ThinkSystem, 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: DBSMS02XX73

10/2017