

Smarter technology for all

Lenovo ThinkSystem SR650 & SR650 V2

STAC-M3™ Antuco Suite

Performance Leadership

Joe Jakubowski

Principal Engineer and Performance Architect

Infrastructure Solutions Group

November 2021

Lenovo

Being reliable means providing deep optimization expertise

Lenovo's tuning scenarios are designed for key workloads from database to virtualization to AI

More than 3X as many #1 results as any other server vendor

283 World records in workload performance

As of September 2, 2021

100

World records
set by servers equipped
3rd Generation Intel®
Xeon® Scalable
Processors

52

World records
set by servers equipped
2nd Generation Intel®
Xeon® Scalable
Processors

25

World records
set by servers using 1st
Generation Intel® Xeon®
Scalable Processors

75

World records
set by servers running
3rd Gen AMD EPYC™

11

World records
set by servers running
2nd Gen AMD EPYC™

20

World records
set by a server using the
Intel® Xeon® D and E
Processors

Lenovo

SR650 V2 system specs



Customer selectable preset UEFI operating modes (low-level optimization)

- Maximum Performance
- Efficiency – Favor Performance
- Efficiency – Favor Power
- Minimal Power
- Custom

2021 Lenovo. All rights reserved.

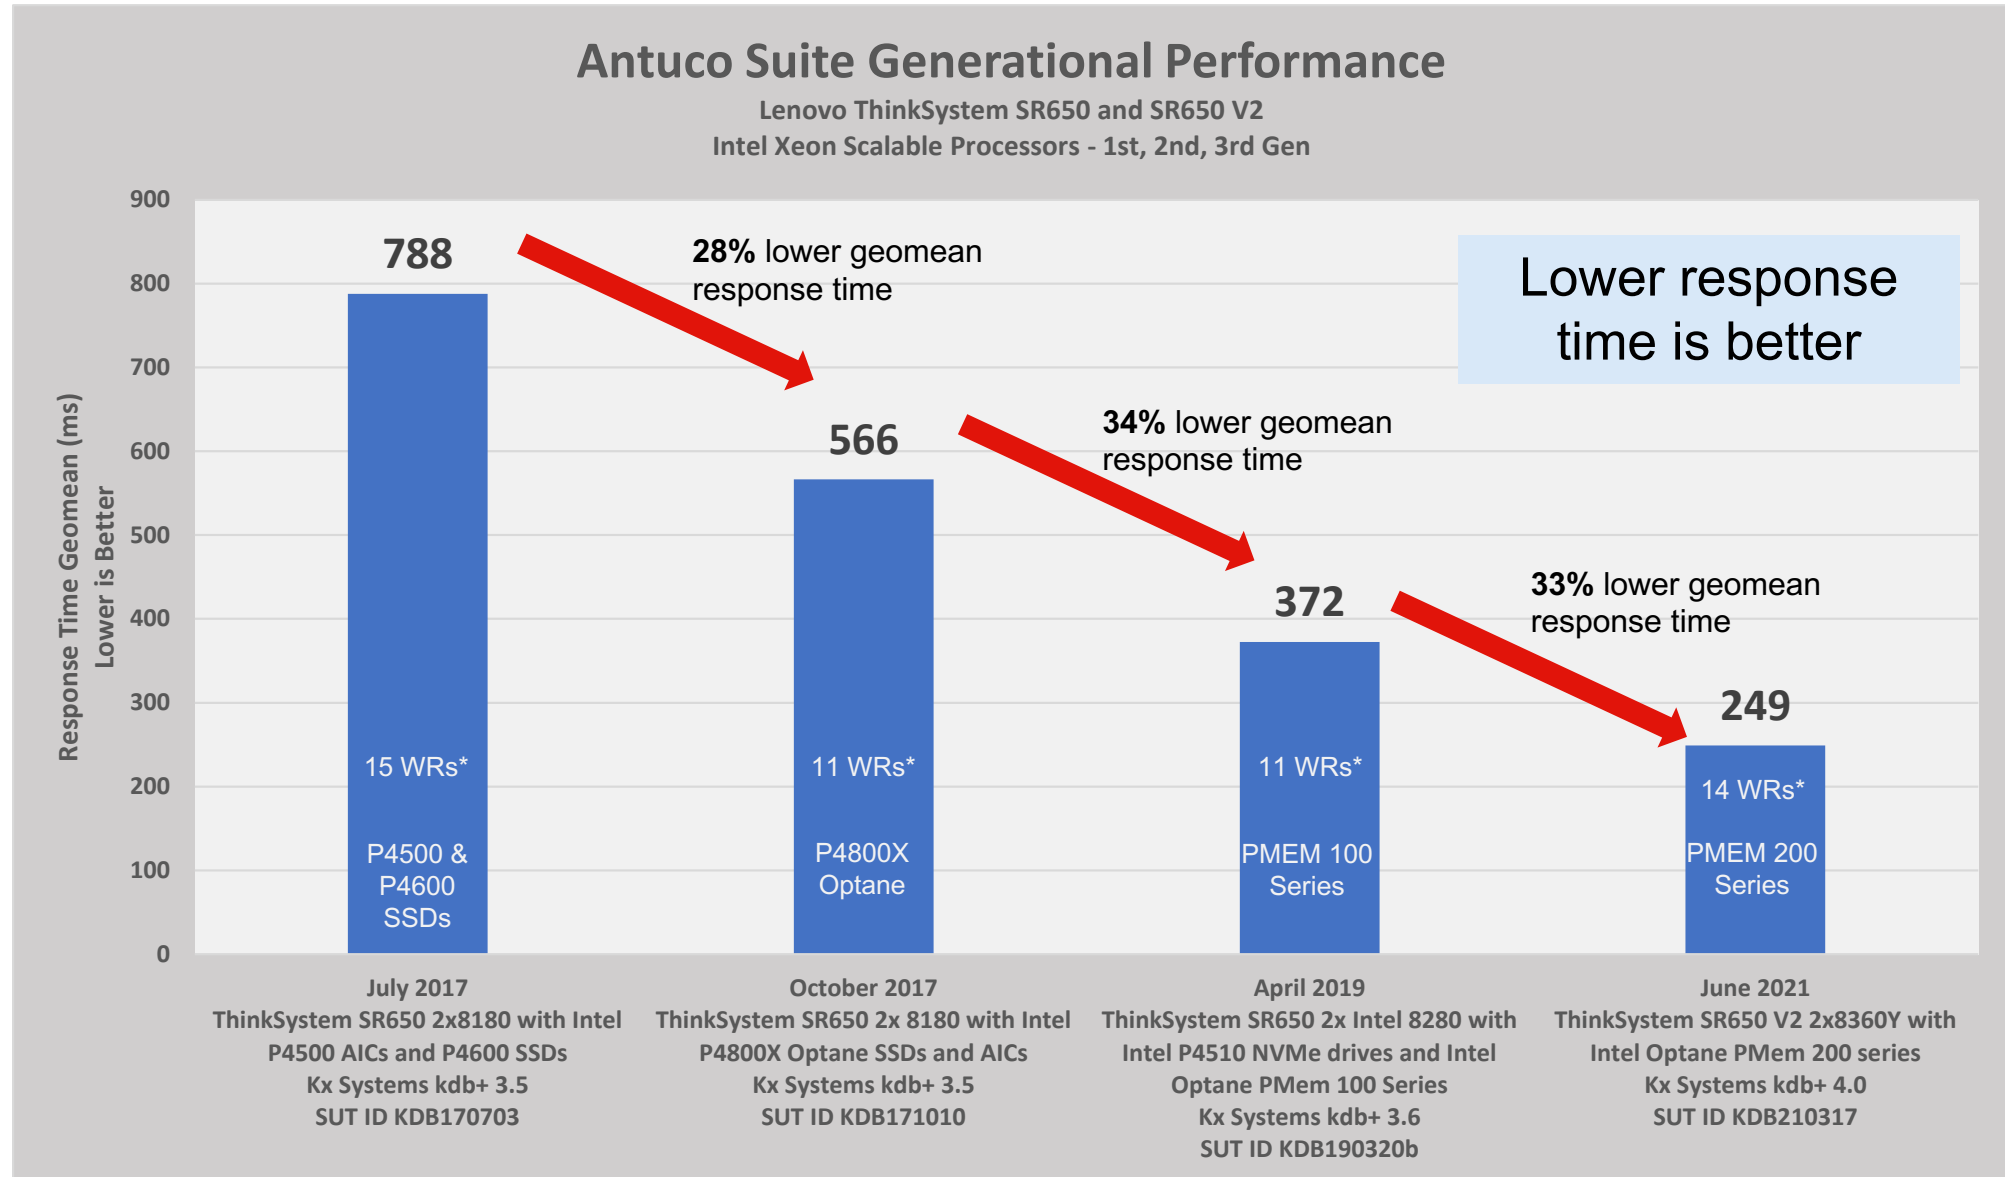
Lenovo

Feature	Specification
Form Factor	• 2U2S (445 x 86.5 x 754.7mm)
Compute	• Support for 3rd Gen Intel Xeon Scalable Processors, up to 270W per CPU
GPU	<ul style="list-style-type: none"> • 3x DW GPU up to 300W each • 8x SW GPU up to 75W each • 6x SW GPU up to 150W each
Memory	<ul style="list-style-type: none"> • Up-to 32 TruDDR4 DIMMs. 16 Intel Optane PMEM 200 Series • TruDDR4 RDIMM ECC 16GB/32GB/64GB/128GB/256GB • Intel Optane PMEM 200 Series 128GB,256GB/512GB • Supports 1 & 2DPC@3200MHz
Storage	<ul style="list-style-type: none"> • Up to 20x 3.5" or 40x 2.5" drive bays • Front: up to 12x3.5" / 24x2.5", hot-swap • Mid (internal): up to 4x3.5" / 8x2.5", Hot swap • Rear: up to 4x3.5" / 8x2.5"*, 2x7mm boot drive, hot-swap • 24 NVMe with 1:1 connection, 32 NVMe with a retimer card • Direct NVMe, VROC/VMD, and onboard SATA w/ SW RAID • Internal dual M.2 media supports simple swap • Common building block PCIe RAID and HBA adapters
Cooling	• Redundant hot swap, standard and performance fan (up to 6)
Networking	• LOM via OCP
Management and security	<ul style="list-style-type: none"> • Full xClarity Support • Redfish support • Optional Chassis Intrusion Switch
Front IO and LED	<ul style="list-style-type: none"> • Selectable LED (internal) or LCD (external) • Default external diagnostic handset connector • Optional VGA • 2 x USB 3.0 with XCC connection to a mobile device
Rear IO and LED	<ul style="list-style-type: none"> • 3 x USB 3.0 • Optional Serial Port using one PCIe slot • 1GbE Dedicated Management port • 8x PCIe slot, 1x OCP Mezz slot • Internal cabled raid slot
Power Supply	• Dual Redundant Platinum and Titanium AC & DC (multiple wattages)

STAC-M3 Antuco benchmark suite


Suite	Purpose	Dataset size*	Concurrent requests	Operations	Constraints related to memory and storage
Antuco	Using a limited dataset size for convenience, simulate performance that would be obtained with a larger real-world dataset residing mostly on non-volatile media. Study a broad range of read and write operations. 17 mean response-time benchmark categories	4.5 TB	1 to 100	Range of compute-bound and storage-bound analytics. A few operations involving writes.	<ul style="list-style-type: none"> • No pre-loading into memory • File system cache cleared at several points in test run
Shasta	Study performance across a broad range of operations for datasets that are relatively small in the real world. (While the dataset tested is the same size as in Antuco, there is no attempt to simulate the storage-access pattern of a larger dataset.)	4.5 TB	1 to 100	Same as Antuco except operations involving writes are optional.	<ul style="list-style-type: none"> • Pre-loading into memory is allowed (most recent data first) • Caches not cleared during test run
Kanaga	Study performance on large datasets with large numbers of concurrent requests.	33 TB to 897 TB	1 to 450	A few storage-intensive queries.	<ul style="list-style-type: none"> • Pre-loading into memory is allowed (most recent data first) • Caches not cleared during test run • Storing certain data into faster storage tiers is allowed

Antuco suite generational comparison



Summary

- **Lenovo ThinkSystem servers are synonymous with industry leading performance**
 - More than **3X current industry benchmark world records*** vs. our nearest competitor as of September 2, 2021
- **Lenovo ThinkSystem SR650 and SR650 V2 servers with Intel Optane** technology have **driven significant generational reductions** in Antuco benchmark suite geometric mean response time
 - More than **3X Antuco geomean response time reduction** since 2017
 - Combination of performance improvements with ThinkSystem hardware + low-level UEFI optimization, Intel Optane PMem 100/200 Series and the Kx Systems kdb+ database



To learn more about Lenovo ThinkSystem server performance and optimization

- Lenovo ThinkSystem - Data Centric Leadership Performance, Intel Product Performance
 - <https://edc.intel.com/content/www/us/en/benchmarks/data-centric-leadership-performance/lenovo/>
- Lenovo ThinkSystem server industry benchmark performance leadership white paper
 - <https://lenovopress.com/lp1145.pdf>
- Balanced Memory Configurations for 2-Socket Servers with 3rd-Gen Intel Xeon Scalable Processors
 - <https://lenovopress.com/lp1517.pdf>
- Tuning UEFI Settings for Performance and Energy Efficiency on Intel Xeon Scalable Processor-Based ThinkSystem Servers
 - <https://lenovopress.com/lp1477.pdf>



Lenovo

Moving at the Speed of our Clients

Best Performance

Leadership Workload Performance

Advanced Memory Technology

Secure IT Foundation

STAC® Leadership

Best Reliability

Better Utilize Data

Lower Overall IT Costs

Reduce Downtime

Scale to your Business

Best Business Value

Focus on Innovation

Leverage AI for Smarter Infra

Get to Market Faster

Show Return on Investment

Choice of CapEx or OpEx

Lenovo

2021 Lenovo Internal. All rights reserved.

Smarter
technology
for all

Lenovo

thanks.

BACKUP

Configuration

SUT REPORT AND CONFIGURATION INFORMATION				
Report Date	July 2017	October 2017	April 2019	June 2021
SUT ID	KDB170703	KDB171010	KDB190320b	KDB210317
Vendor	Lenovo	Lenovo	Lenovo	Lenovo
System	ThinkSystem SR650	ThinkSystem SR650	ThinkSystem SR650	ThinkSystem SR650 V2
nodes	1	1	1	1
Sockets	2	2	2	2
CPU	Intel Xeon Scalable Processor 8180	Intel Xeon Scalable Processor 8180	2nd Generation Intel Xeon Scalable Processor 8280L	3rd Generation Intel Xeon Scalable Processor 8360Y
Database Software	Kx Systems kdb+ 3.5 2017.06.19	Kx Systems kdb+ 3.5 2017.06.19	Kx Systems kdb+ 3.6 2018.11.09	Kx Systems kdb+ 4.0
Storage	4 x 1.6TB SSD P4600 U.2 2 x 4TB SSD P4500 AIC (HHHL form factor), over provisioned to 1.6 TB LUNs: 6 partitions, one in each SSD	4x Intel SSD DC P4800X U.2 750GB 2 x Intel SSD DC P4800X AIC 750GB (Optane) drives LUNs: 6 partitions, one in each SSD	2 x Intel SSD DC P4510 8TB NVMe (3DNAND) 12 x Intel® Optane® DC Persistent Memory Module (DCPMM) 512GB LUNs: 2 partitions, one in each NVMe drive	16x 256GiB Intel® Optane™ PMem 200 series

Data

Antuco Suite Benchmark Data				
Report Date	July 2017	October 2017	April 2019	June 2021
SUT ID	KDB170703	KDB171010	KDB190320b	KDB210317
	Mean Response Time (ms)	Mean Response Time (ms)	Mean Response Time (ms)	Mean Response Time (ms)
100T.STATS-UI.TIME	5492	6647	4026	1898
100T.VWAB-12D-NO.TIME	10653	4281	3426	1933
10T.MKTSNAP.TIME	233	31	87	14
10T.STATS-AGG.TIME	7492	10233	7691	4338
10T.STATS-UI.TIME	2636	3584	428	682
10T.THEOPL.TIME	56	22	13	16
10T.VOLCURV.TIME	4685	3112	3440	2761
1T.MOHIBID.TIME	66	30	24	16
1T.NBBO.TIME	19814	14925	17883	13747
1T.QTRHIBID.TIME	96	83	61	53
1T.STATS-UI.TIME	571	784	75	137
1T.VWAB-D.TIME	28	25	12	11
1T.WKHIBID.TIME	54	18	19	12
1T.WRITE.TIME	8864	7974	9333	3930
1T.YRHIBID-2.TIME	128	113	107	58
1T.YRHIBID.TIME	342	324	224	132
50T.STATS-UI.TIME	5844	7515	2874	1326
Geomean	788	566	372	249