

NVMe-Rich Configurations of the ThinkSystem SR650 Reference Information

The Lenovo ThinkSystem SR650 is a mainstream 2U 2-socket server with industry-leading reliability, management, and security features, and is designed to handle a wide range of workloads.

New to the SR650 is support for up to 24 NVMe solid-state drives. With this support, the SR650 is an excellent choice for workloads that need large amounts of low-latency high-bandwidth storage, including virtualized clustered SAN solutions, software-defined storage, and applications leveraging NVMe over Fabrics (NVMeOF).




Figure 1. ThinkSystem SR650 with 24 NVMe drives

Three new configurations are now available:

- 16 NVMe drives + 8 SAS/SATA drives
- 20 NVMe drives
- 24 NVMe drives

NVMe (Non-Volatile Memory Express) is a technology that overcomes SAS/SATA SSD performance limitations by optimizing hardware and software to take full advantage of flash technology. Intel Xeon processors efficiently transfer data in fewer clock cycles with the NVMe optimized software stack compared to the legacy AHCI stack, thereby reducing latency and overhead. NVMe SSDs connect directly to the processor via the PCIe bus, further reducing latency. NVMe drives are characterized by very high bandwidth and very low latency.

Ordering information

These configurations are available configure-to-order (CTO) in the Lenovo Data Center Solution Configurator (DCSC), <https://dcsc.lenovo.com>. The following table lists the feature codes related to the NVMe drive subsystem. The configurator will derive any additional components that are needed.

Field upgrades: The 20x NVMe and 24x NVMe drive configurations are also available as field upgrades as described in the [Field upgrades](#) section.

Table 1. Feature codes for CTO orders

Feature code	Description
PCIe Switch Adapters	
B22D	ThinkSystem 810-4P NVMe Switch Adapter (PCIe x8 adapter with four x4 drive connectors)
AUV2	ThinkSystem 1610-4P NVMe Switch Adapter (PCIe x16 adapter with four x4 drive connectors)
B4PA	ThinkSystem 1610-8P NVMe Switch Adapter (PCIe x16 adapter with four connectors to connect to eight drives)
NVMe Backplane	
B4PC	ThinkSystem SR650 2.5" NVMe 8-Bay Backplane
Riser Cards	
AUR3	ThinkSystem SR550/SR590/SR650 x16/x8 PCIe FH Riser 1 Kit (x16+x8 PCIe Riser for Riser 1, for 16 and 20-drive configurations)
B4PB	ThinkSystem SR650 x16/x8/x16 PCIe Riser1 (x16+x8+x16 PCIe Riser for Riser 1, for 24-drive configurations)
AURC	ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit (x16+x16 PCIe Riser for Riser 2, for all three configurations)

Note the following requirements for any of the three NVMe-rich configurations:

- Two processors
- No high-thermal processors:
 - 200 W or 205 W TDP are not supported
 - Gold 6126T, Gold 6144, Gold 6146, or Platinum 8160T processors are not supported
 - Gold 6230N, Gold 6240Y, and Gold 6244 processors are not supported
- No GPU adapters installed
- No PCIe flash adapters installed
- No PCIe adapters with more than 25 W TDP installed
- 1100 W or 1600 W power supplies installed.
- Ambient temperature of up to 30 °C (86 °F)
- If a fan fails and the ambient temperature is above 27 °C, system performance may be reduced.

Although not required, it is expected that these configurations will be fully populated with NVMe drives. Maximum performance is achieved when all NVMe drive bays are filled with drives.

To verify support and ensure that the right power supply is chosen for optimal performance, validate your server configuration using the latest version of the Lenovo Capacity Planner:

<http://datacentersupport.lenovo.com/us/en/solutions/Invo-lcp>

Supported NVMe drives

See the ThinkSystem SR650 product guide for the complete list of NVMe drives that are supported in the server: <https://lenovopress.com/lp0644#drives-for-internal-storage>

The NVMe drives listed in the following table are not supported in the three NVMe-rich configurations.

Table 2. NVMe drives that are not supported in the 16, 20, and 24x NVMe drive configurations

Part number	Feature code	Description
Unsupported NVMe drives		
7SD7A05770	B11L	ThinkSystem U.2 Intel P4600 6.4TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD
7N47A00984	AUV0	ThinkSystem U.2 PM963 1.92TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD
7N47A00985	AUUU	ThinkSystem U.2 PM963 3.84TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD

Configuration 1: 16x NVMe drives + 8x SAS/SATA

The 16x NVMe drive configuration has the following features:

- 16 NVMe 2.5-inch drive bays plus eight SAS/SATA 2.5-inch drive bays. All drives are hot-swap from the front of the server (provided the operating system supports hot-swap).
- The NVMe drives are connected to the processors either via NVMe Switch Adapters or via the onboard NVMe connectors on the system board of the server.
- The eight SAS/SATA drive bays are connected to a supported 8-port RAID adapter or SAS HBA.
- One PCIe x16 slot is available for high-speed networking such as a 100 GbE adapter, InfiniBand or OPA adapter. If you elect not to configure the eight SAS/SATA drive bays, then you can free up an additional x8 slot for a second networking adapter.
- The LOM (LAN on Motherboard) slot is also available for 1Gb or 10Gb Ethernet connections. Supported LOM adapters are the following:
 - ThinkSystem 1Gb 2-port RJ45 LOM
 - ThinkSystem 1Gb 4-port RJ45 LOM
 - ThinkSystem 10Gb 2-port Base-T LOM
 - ThinkSystem 10Gb 2-port SFP+ LOM
 - ThinkSystem 10Gb 4-port Base-T LOM
 - ThinkSystem 10Gb 4-port SFP+ LOM
- Additional support for one or two M.2 drives, if needed

The 16x NVMe drive configuration has the following performance characteristics:

- **Balanced NVMe configuration.** In this 16-NVMe drive configuration, each processor is connected to 8 drives. Such a balanced configuration ensures maximum performance by ensuring the processors are equally occupied handling I/O requests to and from the NVMe drives.
- **No oversubscription.** Lenovo NVMe drives connect using four PCIe lanes, and in this configuration, each drive is allocated 4 lanes from the processor. The 1:1 ratio means no oversubscription of the PCIe lanes from the processors and results in maximum NVMe drive bandwidth.

In the 16x NVMe drive configuration, the drive bays are configured as follows:

- Bays 0-15: NVMe drives
- Bays 16-23: SAS or SATA drives

The PCIe slots in the server are configured as follows:

- Slot 1: 1610-4P NVMe Switch Adapter
- Slot 2: Not present
- Slot 3: Supported RAID adapter for SAS/SATA drives
- Slot 4: 810-4P NVMe Switch Adapter
- Slot 5: Available x16 slot
- Slot 6: 1610-4P NVMe Switch Adapter
- Slot 7 (internal slot): 810-4P NVMe Switch Adapter

The front and rear views of the SR650 with 16x NVMe drives and 8x SAS/SATA drives is shown in the following figure.




Figure 2. SR650 front and rear views of the 16-NVMe drive configuration

The following figure shows a block diagram of how the PCIe lanes are routed from the processors to the NVMe drives.

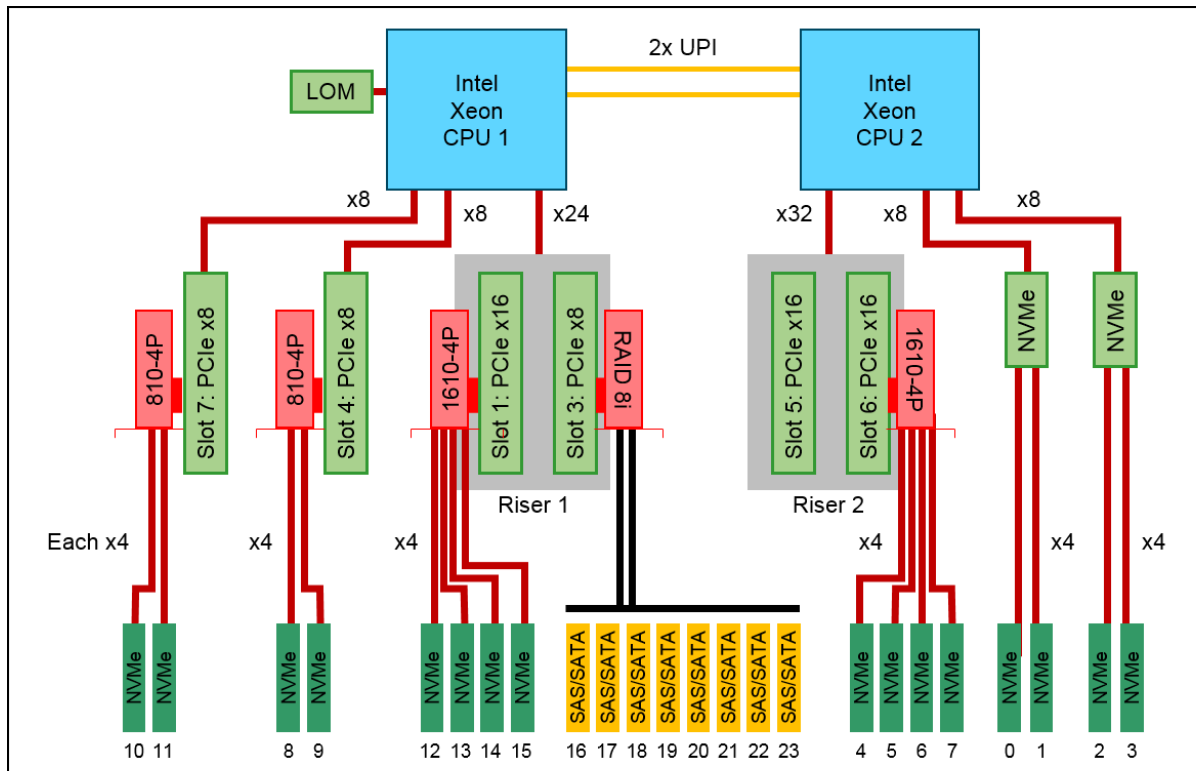


Figure 3. SR650 block diagram of the 16-NVMe drive configuration

The details of the connections are listed in the following table.

Table 3. Drive connections

Drive bay	Drive type	Drive lanes	Adapter	Slot	Host lanes	CPU
0	NVMe	PCIe x4	Onboard NVMe port	None	PCIe x8	2
1	NVMe	PCIe x4				2
2	NVMe	PCIe x4	Onboard NVMe port	None	PCIe x8	2
3	NVMe	PCIe x4				2
4	NVMe	PCIe x4	1610-4P	Slot 6 (Riser 2)	PCIe x16	2
5	NVMe	PCIe x4				2
6	NVMe	PCIe x4				2
7	NVMe	PCIe x4				2
8	NVMe	PCIe x4	810-4P	Slot 4 (vertical)	PCIe x8	1
9	NVMe	PCIe x4				1
10	NVMe	PCIe x4	810-4P	Slot 7 (internal)	PCIe x8	1
11	NVMe	PCIe x4				1
12	NVMe	PCIe x4	1610-4P	Slot 1 (Riser 1)	PCIe x16	1
13	NVMe	PCIe x4				1
14	NVMe	PCIe x4				1
15	NVMe	PCIe x4				1
16	SAS or SATA		RAID 8i	Slot 3 (Riser 1)	PCIe x8	1
17	SAS or SATA					1
18	SAS or SATA					1
19	SAS or SATA					1
20	SAS or SATA					1
21	SAS or SATA					1
22	SAS or SATA					1
23	SAS or SATA					1

Configuration 2: 20x NVMe drives

The 20x NVMe drive configuration has the following features:

- 20 NVMe 2.5-inch drive bays. All drives are hot-swap from the front of the server (provided the operating system supports hot-swap). The other 4 bays are unavailable and are covered by a 4-bay blank.
- The NVMe drives are connected to the processors either via NVMe Switch Adapters or via the onboard NVMe connectors on the system board of the server.
- One PCIe x8 slot is available for networking or other needs. The LOM (LAN on Motherboard) slot is also available for 1Gb or 10Gb Ethernet connections. Supported LOM adapters are the following:
 - ThinkSystem 1Gb 2-port RJ45 LOM
 - ThinkSystem 1Gb 4-port RJ45 LOM
 - ThinkSystem 10Gb 2-port Base-T LOM
 - ThinkSystem 10Gb 2-port SFP+ LOM
 - ThinkSystem 10Gb 4-port Base-T LOM
 - ThinkSystem 10Gb 4-port SFP+ LOM
- Additional support for one or two M.2 drives, if needed

The 20x NVMe drive configuration has the following performance characteristics:

- **No oversubscription.** Lenovo NVMe drives connect using four PCIe lanes, and in this configuration, each drive is allocated 4 lanes from the processor. The 1:1 ratio means no oversubscription of the PCIe lanes from the processors and results in maximum NVMe drive bandwidth.
- **Near-balanced NVMe configuration.** Unlike the 16-drive and 24-drive configurations, that 20-drive configuration has eight NVMe drives connected to processor 1, and 12 NVMe drives connected to processor 2. As a result, we recommend you to only choose this configuration if you need the additional capacity that four drives provide above the 16-drive configuration, and your workload can fully operate without an equal number of drives connected to each processor.

The PCIe slots in the server are configured as follows:

- Slot 1: 1610-4P NVMe Switch Adapter
- Slot 2: Not present
- Slot 3: Available x8 slot
- Slot 4: 810-4P NVMe Switch Adapter
- Slot 5: 1610-4P NVMe Switch Adapter
- Slot 6: 1610-4P NVMe Switch Adapter
- Slot 7 (internal slot): 810-4P NVMe Switch Adapter

The front and rear views of the SR650 with 20x NVMe drives is shown in the following figure.




Figure 4. SR650 front and rear views of the 20-NVMe drive configuration

The following figure shows a block diagram of how the PCIe lanes are routed from the processors to the NVMe drives.

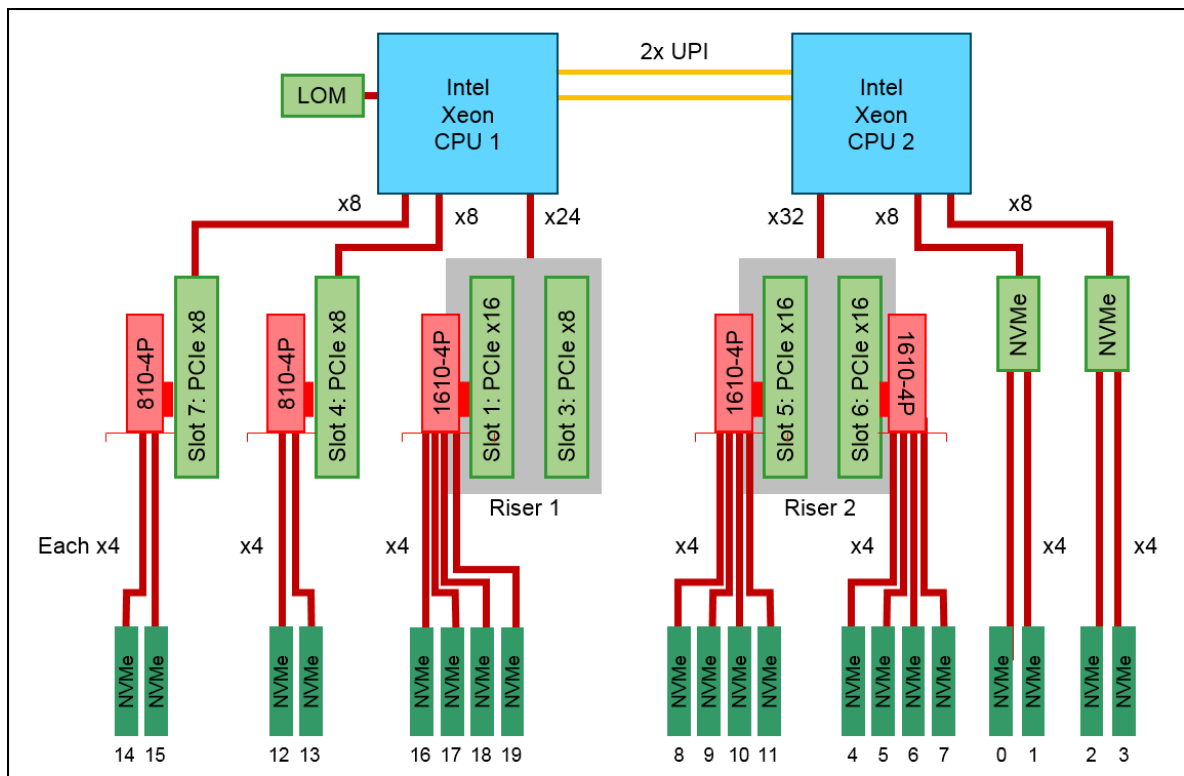


Figure 5. SR650 block diagram of the 20-NVMe drive configuration

The details of the connections are listed in the following table.

Table 4. Drive connections

Drive bay	Drive type	Drive lanes	Adapter	Slot	Host lanes	CPU
0	NVMe	PCIe x4	Onboard NVMe port	None	PCIe x8	2
1	NVMe	PCIe x4				2
2	NVMe	PCIe x4	Onboard NVMe port	None	PCIe x8	2
3	NVMe	PCIe x4				2
4	NVMe	PCIe x4	1610-4P	Slot 6 (Riser 2)	PCIe x16	2
5	NVMe	PCIe x4				2
6	NVMe	PCIe x4				2
7	NVMe	PCIe x4				2
8	NVMe	PCIe x4	1610-4P	Slot 5 (Riser 2)	PCIe x16	2
9	NVMe	PCIe x4				2
10	NVMe	PCIe x4				2
11	NVMe	PCIe x4				2
12	NVMe	PCIe x4	810-4P	Slot 4 (vertical)	PCIe x8	1
13	NVMe	PCIe x4				1
14	NVMe	PCIe x4	810-4P	Slot 7 (internal)	PCIe x8	1
15	NVMe	PCIe x4				1
16	NVMe	PCIe x4	1610-4P	Slot 1 (Riser 1)	PCIe x16	1
17	NVMe	PCIe x4				1
18	NVMe	PCIe x4				1
19	NVMe	PCIe x4				1
20	Blank bay - no connection					
21	Blank bay - no connection					
22	Blank bay - no connection					
23	Blank bay - no connection					

Configuration 3: 24x NVMe drives

The 24x NVMe drive configuration has the following features:

- 24 NVMe 2.5-inch drive bays. All drives are hot-swap from the front of the server (provided the operating system supports hot-swap).
- The NVMe drives are connected to the processors via NVMe Switch Adapters. The onboard NVMe connectors are routed to a riser card installed in Riser slot 1.
- Two x16 slots (one connected to each processor) are available for high-speed networking such as a 100 GbE adapter, InfiniBand or OPA adapter.
- The LOM (LAN on Motherboard) slot is also available for 1Gb or 10Gb Ethernet connections. Supported LOM adapters are the following:
 - ThinkSystem 1Gb 2-port RJ45 LOM
 - ThinkSystem 1Gb 4-port RJ45 LOM
 - ThinkSystem 10Gb 2-port Base-T LOM
 - ThinkSystem 10Gb 2-port SFP+ LOM
 - ThinkSystem 10Gb 4-port Base-T LOM
 - ThinkSystem 10Gb 4-port SFP+ LOM
- Additional support for one or two M.2 drives, if needed

The 24x NVMe drive configuration has the following performance characteristics:

- **Balanced NVMe configuration.** In this 24-NVMe drive configuration, each processor is connected to 12 drives. Such a balanced configuration provides maximum performance by ensuring the processors are equally occupied handling I/O requests to and from the NVMe drives.
- **2:1 oversubscription.** Lenovo NVMe drives connect using four PCIe lanes, and in this configuration each drive is allocated 2 lanes from the processor, resulting in a 2:1 oversubscription of the PCIe lanes. With 24 drives, there are simply not enough PCIe lanes in a two-socket server to support no oversubscription. As a result, the design objective is to minimize the oversubscription while still maintaining balance across all lanes.
- **Balanced open slots.** This configuration has two open PCIe x16 slots, one connected to each processor. These slots could be used for a pair of high-speed network cards and the result would be balanced configuration.

The PCIe slots in the server are configured as follows:

- Slot 1: 1610-8P NVMe Switch Adapter
- Slot 2: 810-4P NVMe Switch Adapter
- Slot 3: Available x16 slot
- Slot 4: 810-4P NVMe Switch Adapter
- Slot 5: Available x16 slot
- Slot 6: 810-4P NVMe Switch Adapter
- Slot 7 (internal slot): 810-4P NVMe Switch Adapter

The front and rear views of the SR650 with 24x NVMe drives is shown in the following figure.




Figure 6. SR650 front and rear views of the 24-NVMe drive configuration

The following figure shows a block diagram of how the PCIe lanes are routed from the processors to the NVMe drives.

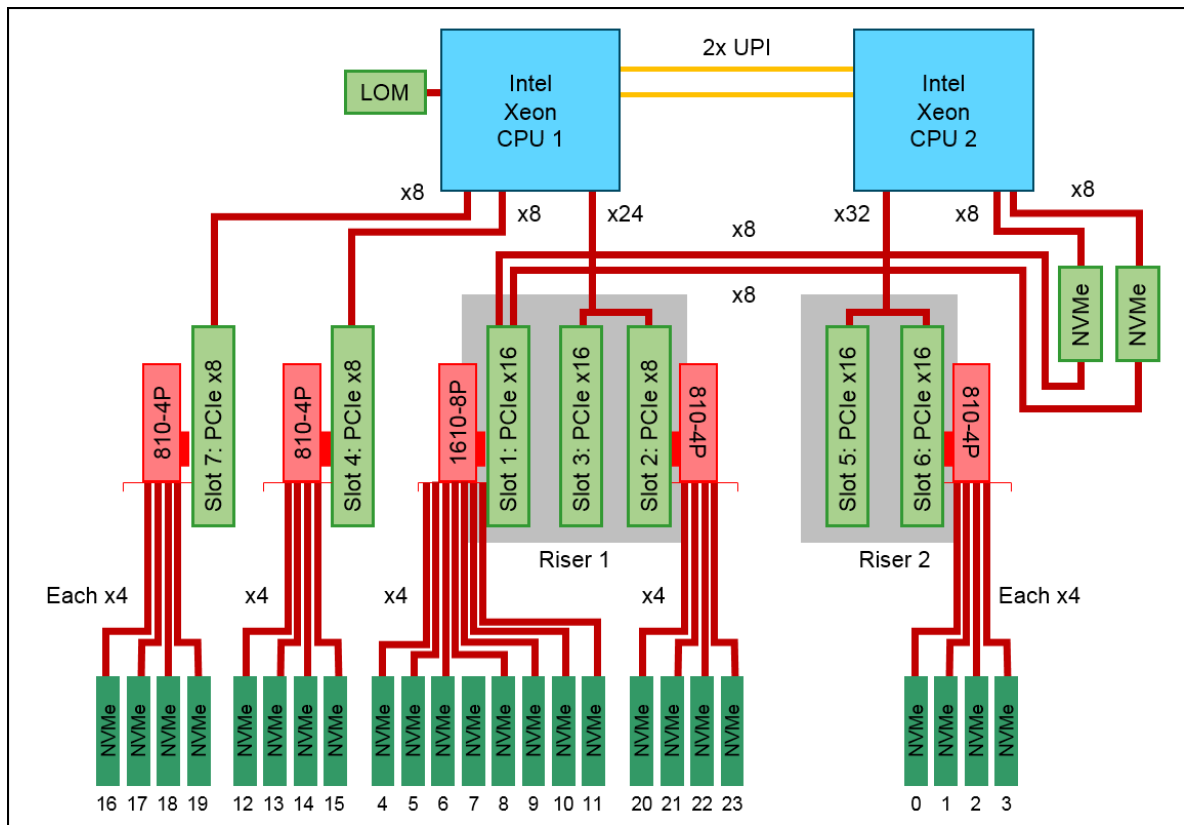


Figure 7. SR650 block diagram of the 24-NVMe drive configuration

The details of the connections are listed in the following table.

Table 5. Drive connections

Drive bay	Drive type	Drive lanes	Adapter	Slot	Host lanes	CPU
0	NVMe	PCIe x4	810-4P	Slot 6 (Riser 2)	PCIe x8	2
1	NVMe	PCIe x4				2
2	NVMe	PCIe x4				2
3	NVMe	PCIe x4				2
4	NVMe	PCIe x4	1610-8P	Slot 1 (Riser 1)	PCIe x16 (from onboard NVMe ports)	2
5	NVMe	PCIe x4				2
6	NVMe	PCIe x4				2
7	NVMe	PCIe x4				2
8	NVMe	PCIe x4				2
9	NVMe	PCIe x4				2
10	NVMe	PCIe x4				2
11	NVMe	PCIe x4				2
12	NVMe	PCIe x4	810-4P	Slot 4 (vertical)	PCIe x8	1
13	NVMe	PCIe x4				1
14	NVMe	PCIe x4				1
15	NVMe	PCIe x4				1
16	NVMe	PCIe x4	810-4P	Slot 7 (internal)	PCIe x8	1
17	NVMe	PCIe x4				1
18	NVMe	PCIe x4				1
19	NVMe	PCIe x4				1
20	NVMe	PCIe x4	810-4P	Slot 2 (Riser 1)	PCIe x8	1
21	NVMe	PCIe x4				1
22	NVMe	PCIe x4				1
23	NVMe	PCIe x4				1

Field upgrades

The following two field upgrade option kits are available to upgrade existing SAS/SATA or AnyBay drive configurations based on the 24x 2.5" chassis (feature code AUVV) to either the 20-drive or 24-drive NVMe configurations.

Table 6. Field upgrades

Part number	Feature code	Description
4XH7A09819	B64L	ThinkSystem SR650 U.2 20-Bays Upgrade Kit
4XH7A08810	B64K	ThinkSystem SR650 U.2 24-Bays Upgrade Kit

These kits include drive backplanes and required NVMe cables, power cables, drive bay fillers, and NVMe switch adapters.

No 16-drive upgrade kit: There is no upgrade kit for the 16x NVMe drive configuration.

The ThinkSystem SR650 U.2 20-Bays Upgrade Kit includes the following components:

- Two 810-4P NVMe Switch Adapters
- Three 1610-4P NVMe Switch Adapters
- One x16/x8 PCIe Riser for Riser 1
- One x16/x16 PCIe Riser for Riser 2
- Three 8-bay NVMe drive backplanes
- One 4-bay drive bay filler
- NVMe and power cables
- Brackets and screws
- Drive bay labels for the front bezel

The ThinkSystem SR650 U.2 24-Bays Upgrade Kit includes the following components:

- Four 810-4P NVMe Switch Adapters
- One 1610-8P NVMe Switch Adapter
- One x16/x8/x16 PCIe Riser for Riser 1
- One x16/x16 PCIe Riser for Riser 2
- Three 8-bay NVMe drive backplanes
- NVMe and power cables
- Brackets and screws
- Drive bay labels for the front bezel

Further information

For more information, see these resources:

- ThinkSystem SR650 product guide
<https://lenovopress.com/lp0644-lenovo-thinksystem-sr650-server>
- Product Guides for ThinkSystem NVMe drives:
<https://lenovopress.com/servers/options/drives#term=nvme&rt=product-guide>
- Paper, *Implementing NVMe Drives on Lenovo Servers*
<https://lenovopress.com/lp0508-implementing-nvme-drives-on-lenovo-servers>
- Paper, *Comparing the Effect of PCIe Host Connections on NVMe Drive Performance*
<https://lenovopress.com/lp0865-comparing-the-effect-of-pcie-host-connections-on-nvme-drive-performance>
- Data Center Solution Configurator (DCSC) configurator
<https://dcsc.lenovo.com/>

Related product families

Product families related to this document are the following:

- [ThinkSystem SR650 Server](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2020. All rights reserved.

This document, LP0904, was created or updated on April 2, 2019.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<http://lenovopress.com/LP0904>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <http://lenovopress.com/LP0904>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

AnyBay®

Lenovo®

ThinkSystem

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.