

3rd Generation Intel Xeon Scalable processor features and configurations

Product features and 4S configuration

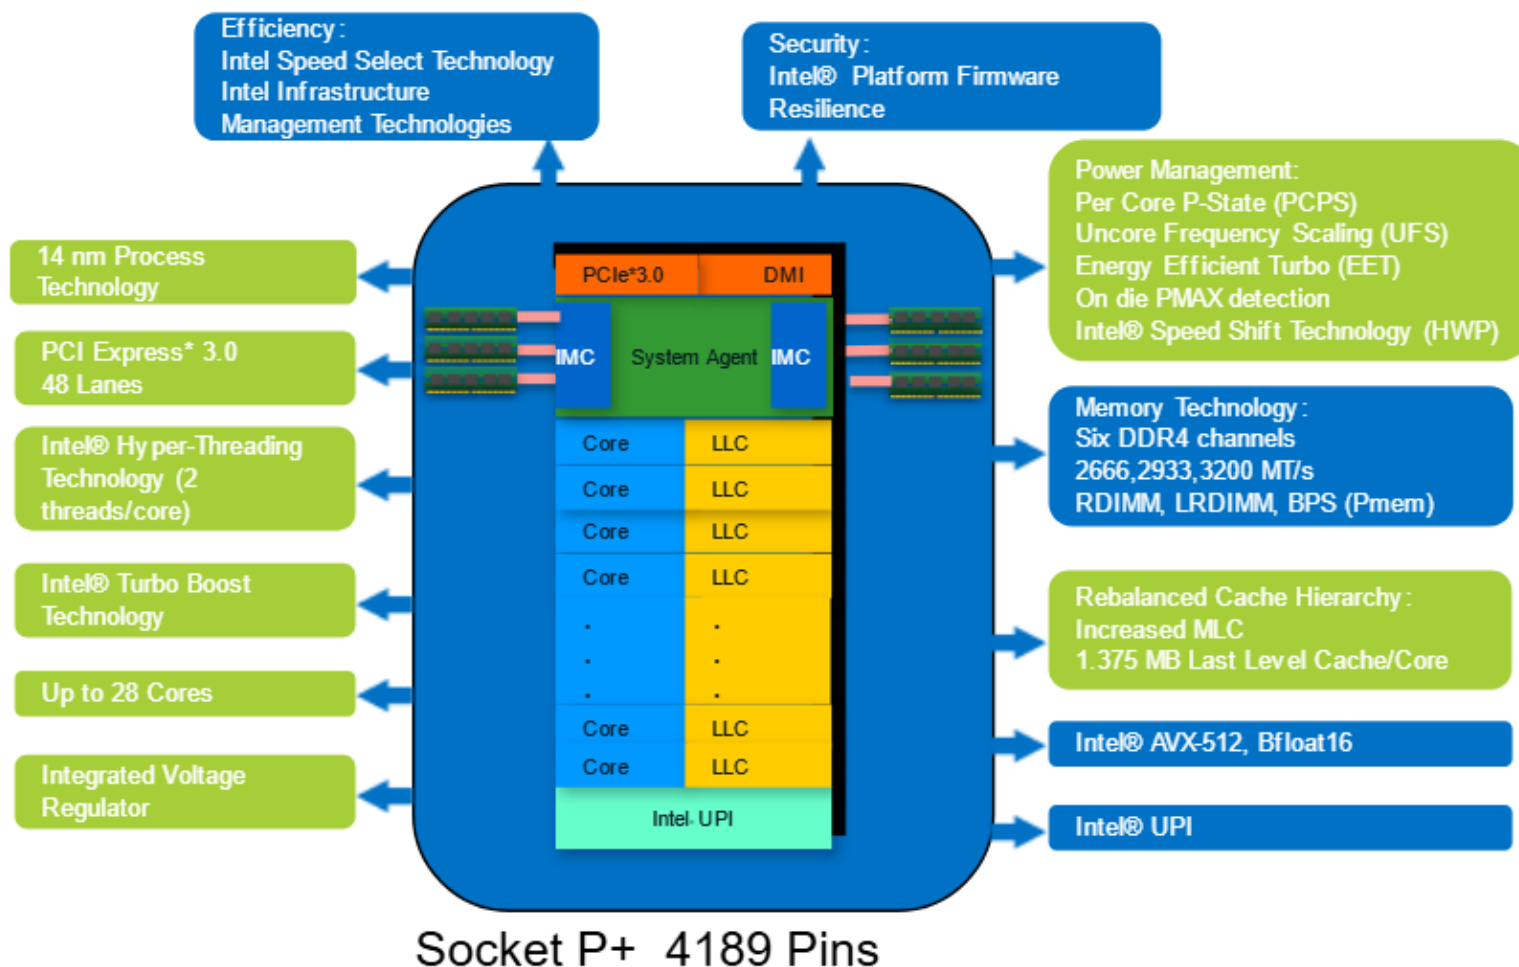
Lenovo

ThinkSystem V2 Series common features

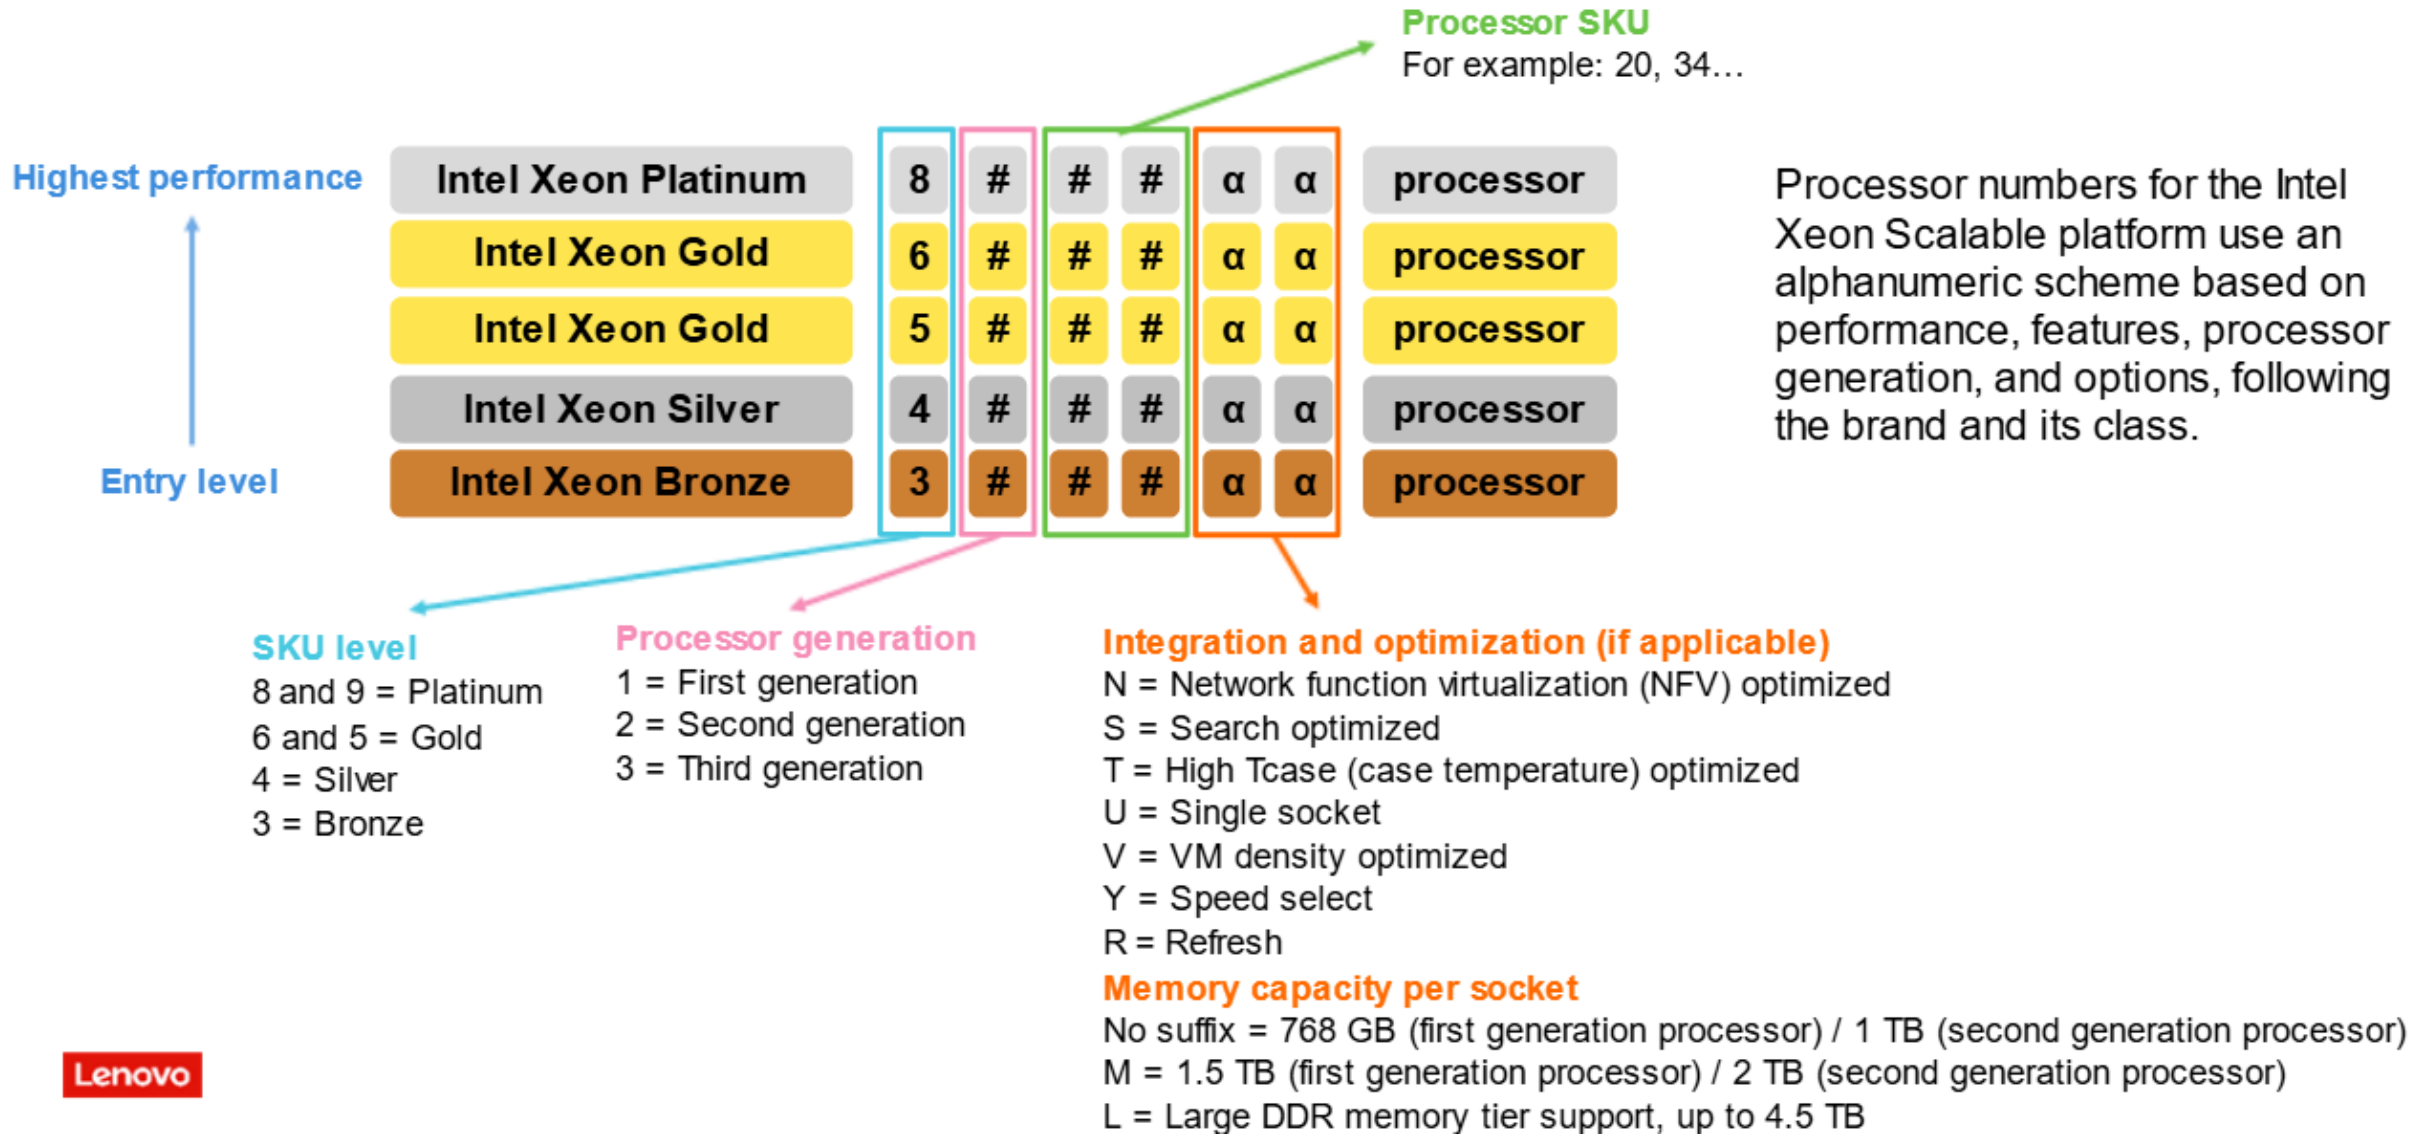
- Mechanical “brick” design – CPU installation and removal with the heat sink
- Heat sinks - 1U Entry and 1U Performance type for high TDP power models
- 2U/4U system rack mechanicals
- RAID adapters
 - Supports 14 varieties of RAID adapters – for more information, refer to course [ES51780C Servicing the ThinkSystem storage controllers](#)
- Storage backplanes with 4th generation mechanical trays for drives
- Three M.2 enablement kits
 - ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit
 - ThinkSystem M.2 NVMe 2-Bay RAID Enablement Kit
 - ThinkSystem M.2 SATA/NVMe 2-Bay Enablement Kit (non-RAID, currently for ThinkSystem SR860 V2 only)
- Two rear 7 mm HDD enablement kits
 - Rear 7 mm SATA SSD RAID Enablement Kit and Rear 7 mm NVMe SSD RAID Enablement Kit
- Power Supplies – common form factor version 4
- New XClarity Controller chipset (Emulex Pilot4 BMC, owned by ASpeed)
- OCP NIC 3.0
 - NIC Mezz Card from two-port 1 GbE to four-port 10/25 GbE

Cedar Island server processor

Note: Blue blocks contain information about new Cedar Island platform server features, and green blocks are for existing features.



Intel Xeon Scalable processor naming rules



3rd Generation Intel Xeon Scalable processor production SKU example

Key Specifications:

- Memory tiers: Base 1.12 TB, Large 4.5 TB
- All shelves will be 6 UPI and 10.4 GT/S
- Platinum SKUs support up to 8S, Gold SKUs up to 4S
- Memory speeds in MT/s :
 - Platinum : 3200 1DPC, 2933 2DPC
 - Gold 6: 2933 2DPC
 - Gold 5: 2666 2DPC

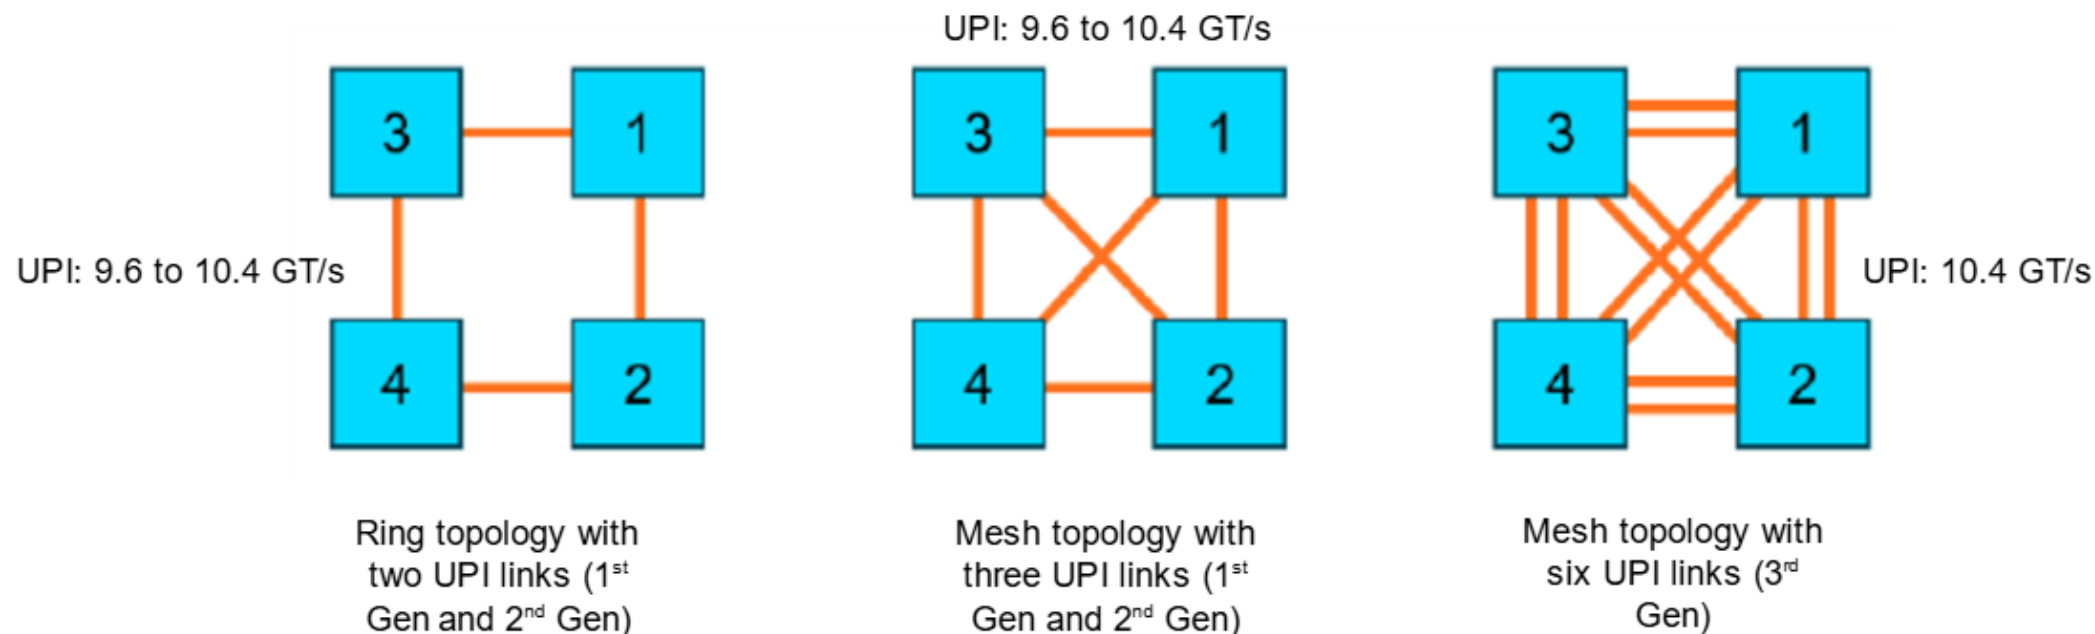
													Intel® Speed Select Technology	
Spec Code	MM #	Processor Number	Sku Shelf	Functional Cores	TDP	DDR4 Frequency	Intel® Optane Persistent Memory Frequency	Memory Max Capacity	RAS	4S/8S Capable	Base Configuration	Max Turbo Frequency	Intel® SST - CP (Core Power)	Intel® SST - TF (Turbo Frequency)
SRJXQ	99A33D	8380H	Intel(R) Xeon(R) Platinum	28 cores	250W	3200	2667	1.12TB	Advanced	8S	28c/250W/2.9GHz	4.3GHz	Disabled	Disabled
SRJXR	99A33N	8380HL	Intel(R) Xeon(R) Platinum	28 cores	250W	3200	2667	4.5TB	Advanced	8S	28c/250W/2.9GHz	4.3GHz	Disabled	Disabled
SRJXS	99A33P	8376H	Intel(R) Xeon(R) Platinum	28 cores	205W	3200	2667	1.12TB	Advanced	8S	28c/205W/2.6GHz	4.3GHz	Disabled	Disabled
SRJXT	99A33R	8376HL	Intel(R) Xeon(R) Platinum	28 cores	205W	3200	2667	4.5TB	Advanced	8S	28c/205W/2.6GHz	4.3GHz	Disabled	Disabled
SRK5Y	99A677	8354H	Intel(R) Xeon(R) Platinum	18 cores	205W	3200	2667	1.12TB	Advanced	8S	18c/205W/3.1GHz	4.3GHz	Disabled	Disabled
SRJXX	99A34A	6348H	Intel(R) Xeon(R) Gold	24 cores	165W	2933	2667	1.12TB	Advanced	4S	24c/165W/2.3GHz	4.2GHz	Disabled	Disabled
SRJXY	99A34C	6328H	Intel(R) Xeon(R) Gold	16 cores	165W	2933	2667	1.12TB	Advanced	4S	16c/165W/2.8GHz	4.3GHz	Enabled	Enabled
SRJXZ	99A34D	6328HL	Intel(R) Xeon(R) Gold	16 cores	165W	2933	2667	4.5TB	Advanced	4S	16c/165W/2.8GHz	4.3GHz	Enabled	Enabled
SRJY1	99A34G	5320H	Intel(R) Xeon(R) Gold	20 cores	150W	2667	2667	1.12TB	Advanced	4S	20c/150W/2.4GHz	4.2GHz	Enabled	Enabled
SRJY2	99A34H	8353H	Intel(R) Xeon(R) Platinum	18 cores	150W	3200	2667	1.12TB	Advanced	8S	18c/150W/2.5GHz	3.8GHz	Disabled	Disabled
SRJY3	99A34K	5318H	Intel(R) Xeon(R) Gold	18 cores	150W	2667	2667	1.12TB	Advanced	4S	18c/150W/2.5GHz	3.8GHz	Disabled	Disabled

Comparing Platinum and Gold SKUs

Platinum SKUs	Gold SKUs
<div>28C 250W 2.9G 8380H/HL</div> <div>28C 205W 2.6G 8376H/HL</div> <div>18C 205W 3.1G 8354H</div> <div>18C 150W 2.5G 8353H</div>	<div>24C 165W 2.3G 6348H</div> <div>16C 165W 2.8G 6328H/HL</div> <div>20C 150W 2.4G 5320H</div> <div>18C 150W 2.5G 5318H</div>
83xx (Platinum)	63xx, 53xx (Gold)
<ul style="list-style-type: none">▪ 4S-6UPI, up to 8S-6UPI▪ 6-ch DDR4 @ 3200 1DPC▪ Intel® Optane™ DC persistent memory 200 Series▪ 16Gb based DDR4 DIMM▪ 6 UPI links @ 10.4GT/s▪ Intel® Turbo Boost▪ Intel® Hyper-Threading▪ Intel® AVX-512 (2x 512-bit FMAs)▪ VNNI, bfloat16▪ 48 lanes PCIe Gen3▪ Node Controller Not Support▪ Advanced RAS	<ul style="list-style-type: none">▪ 4S-6UPI capability▪ 6-ch DDR4 @ 2933 1DPC for 63xxR, 2666 1DPC for 53xxR▪ Intel® Optane™ DC persistent memory 200 Series▪ 16Gb based DDR4 DIMM▪ 6 UPI links @ 10.4GT/s▪ Intel® Turbo Boost▪ Intel® Hyper-Threading▪ Intel® AVX-512 (2x 512-bit FMA on 63xx, 1x 512-bit FMA on 53xx)▪ VNNI, bfloat16▪ 48 lanes PCIe Gen3▪ Advanced RAS

3rd Generation Intel Xeon Scalable processor UPI topology

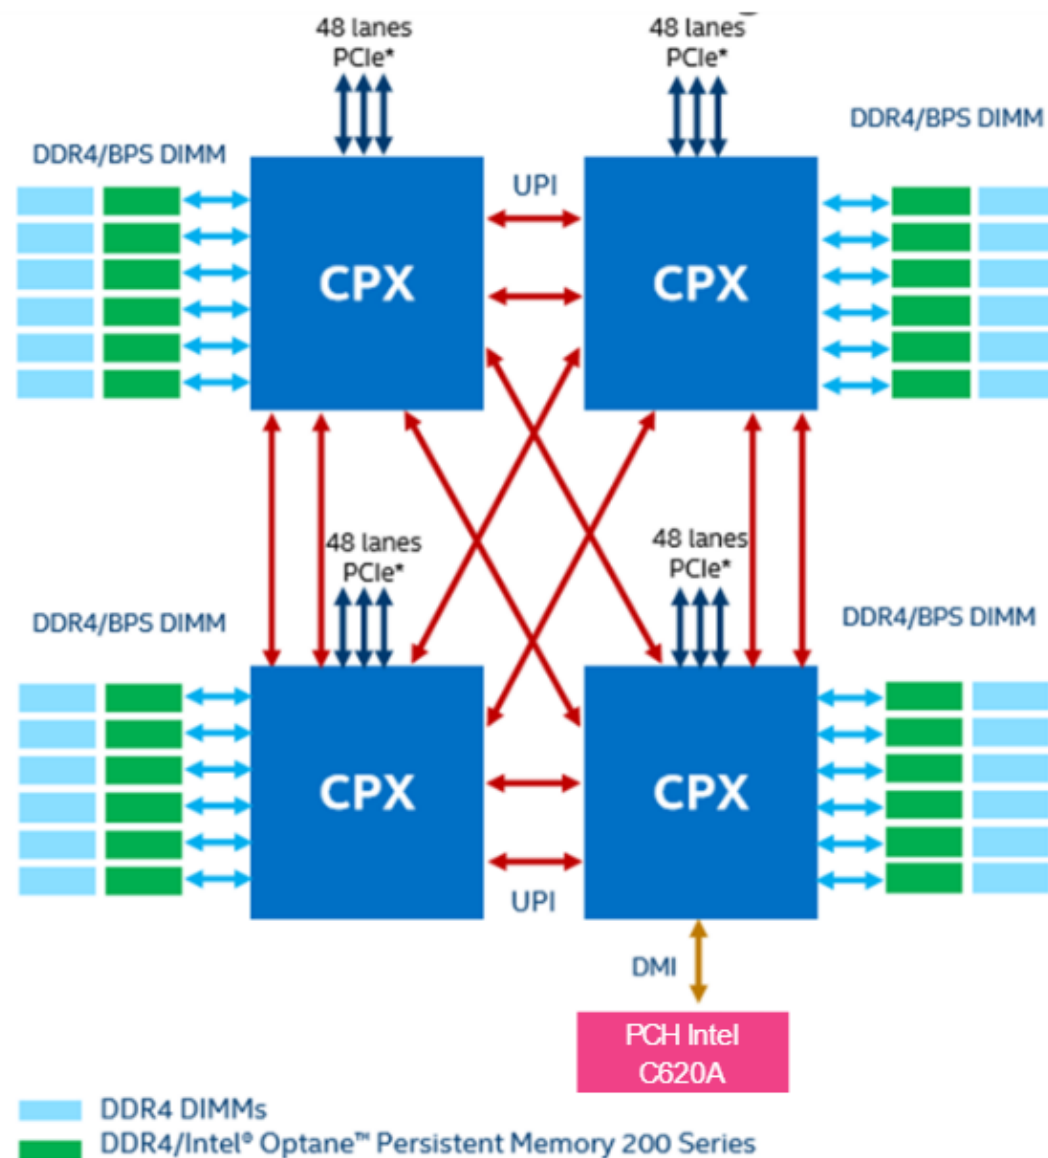
The four processors are connected using a mesh topology. A mesh topology allows each processor to be connected to each of the other processors, which improves the performance of processor-to-processor communications. The ThinkSystem V2 implements a mesh topology using six UPI links.



ThinkSystem V2 four-socket (4S) implementation diagram

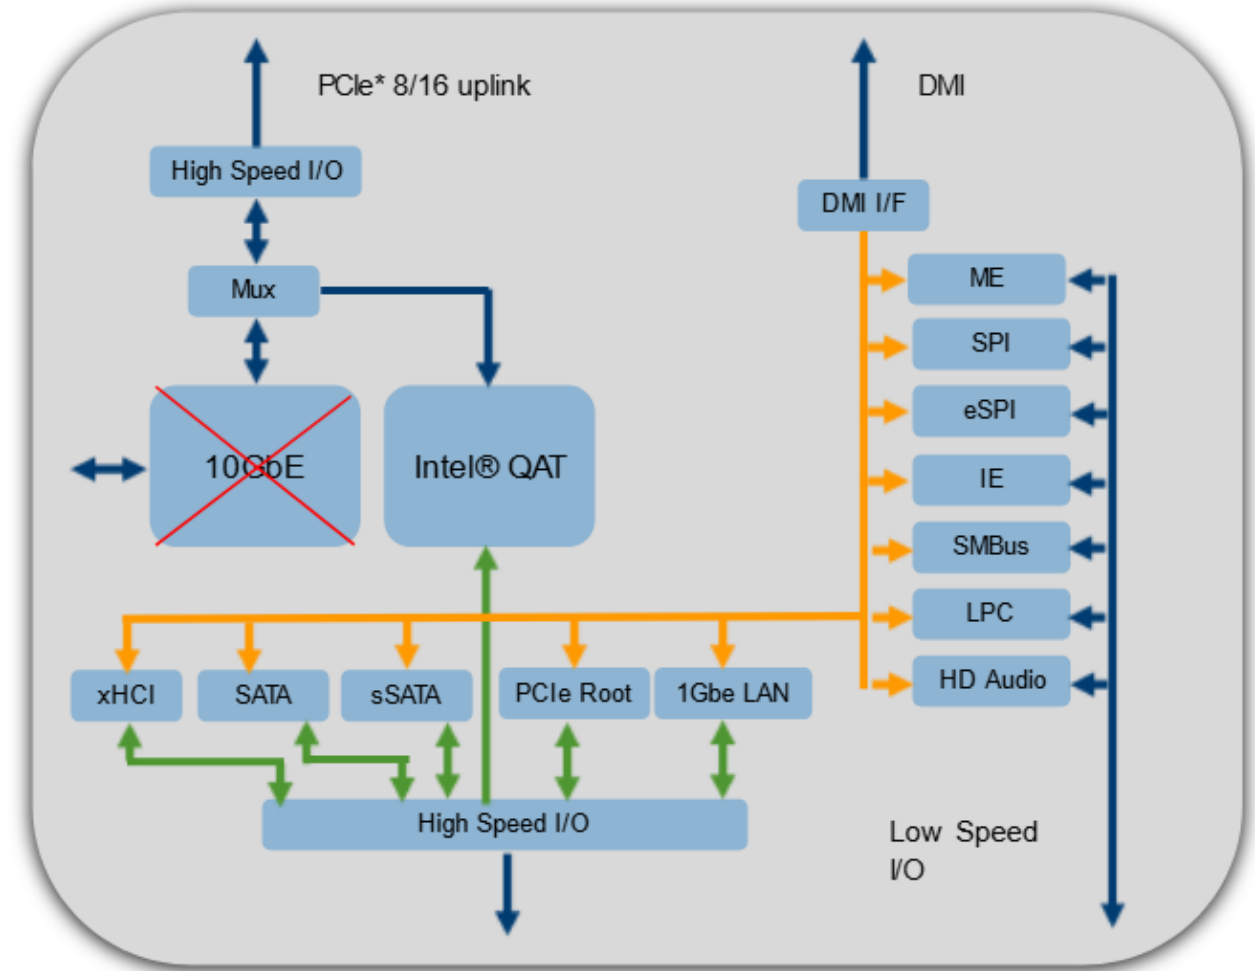
The figure on the right shows the four-socket configuration, which has the following features:

- Memory:
 - Six channels DDR4 per CPU at 3200 Hz, 3200 MHz is 1DPC only
 - Twelve DIMMs per socket
 - New Intel Optane Persistent Memory (PMem) 200 Series Module
 - PMems are valid on 4S only, only App Direct Mode is supported
- Six UPI links, up to 10.4 GT/s
- PCIe 3.0
 - 48 lanes per CPU
 - Bifurcation: x16, x8, x4
- PCH
 - Intel C620A Series Chipset (LBG-R)




Intel C620A southbridge

The ThinkSystem V2 uses the Intel C620A Series chipset as the Platform Control Hub (PCH) and links to the processor through the Direct Media Interface (DMI). Unlike the C620 Series that was used in the ThinkSystem Purley platform, the C620A does not have the integrated Intel X722 Ethernet Quad 10 GB function.



DC Persistent Memory second generation

The ThinkSystem V2 also supports Intel's second-generation 200 Series Optane DC Persistent Memory (PMem). Optane PMems will still be available in 128 GB, 256 GB, and 512 GB modules, and they will also still run at the same DDR4-2666 memory speed. In the ThinkSystem V2, PMems only support the four-socket configuration, and only APP Direct Mode is supported. The following figures show the PMem 200 Series module at the top, and the standard memory DIMM at the bottom.



ThinkSystem V2 RAS features improvement

RAS refers to server reliability, availability, and serviceability. The ThinkSystem V2 RAS features are based on the existing ThinkSystem RAS features but include the following Lenovo-designed improvements:

- Dual-UPI Link Error handler
- PPR (Post Package Repair) enhancement
 - The ability to repair defective memory cells by remapping an address of a defective portion of the memory cells to another portion of the memory cells
- Added Intel Advanced Memory Test Support
- Added Vendor (Hynix, Samsung) Memory Screen Code Support
- Page Retire algorithm enhancement
 - Memory RAS feature
- Auto Crash Dump

Note: For more information about ThinkSystem RAS features, refer to the following Lenovo Press article: <https://lenovopress.com/lp0777.pdf>

Intel Speed Select Technology – 2nd Gen and 3rd Gen comparison

Intel Speed Select Technology (Intel SST) provides a collection of features that give more granular control over CPU performance. With Intel SST, one server can be configured for power and performance for a variety of diverse workload requirements. The following table shows a comparison of 2nd Generation and 3rd Generation Intel Xeon Scalable processors.

	2nd Gen Intel Xeon Scalable processors	3rd Gen Intel Xeon Scalable processors
Intel Speed Select Technology – Performance Profile 1.0 The ability to set a guaranteed base frequency for a specific number of cores and assign this performance profile to a specific application/workload to guarantee performance requirements.	Platinum 8260Y Gold 6240Y Silver 4214	N/A
Intel Speed Select Technology – Base Frequency Allows the configuration of a guaranteed higher base frequency, for a specific number of cores, to support those workloads and applications that are not optimized for turbo frequencies.	Gold 6252N Gold 6230N Gold 5128N	N/A
Intel Speed Select Technology – Core Power Enables flexibility for workloads that benefit from higher base frequencies on a subset of processor cores, and lower base frequencies on the remaining cores, all while maintaining max turbo frequencies across all cores.	N/A	Gold 6238HL Gold 6328H Gold 5320H
Intel Speed Select Technology – Turbo Frequency Enables flexibility for workloads that benefit from higher turbo frequencies on a subset of processor cores, and lower turbo frequencies on the remaining cores, all while maintaining base frequencies across all cores.	N/A	Gold 6238HL Gold 6328H Gold 5320H

Memory installation rules

Memory modules must be installed in a specific order based on the memory configuration that you implement and the number of processors and memory modules installed in the server. Each server has its own installation rules, and servicers should always check the memory installation rules for individual ThinkSystem servers.

The procedure:

1. Go to the [ThinkSystem Documentation Information Center](#) Web site.
2. Select the ThinkSystem portfolio in the left column. In this example, **ThinkSystem rack servers** has been selected.
3. Select a machine type. In this example, **ThinkSystem SR850 V2 Types 7D31, 7D32** has been selected.
4. Select **Server hardware setup** → **Install server hardware options in the server** → **Install a memory module** → **Memory module installation rules and order**.



Note: When you see an magnifier icon  on a figure. Click the figure to enlarge it.

Intel Virtual RAID on CPU

Intel Virtual RAID on CPU (VROC) provides an enterprise RAID solution on platforms that support Intel Volume Management Device (VMD). Intel VMD provides support for RAID on PCIe NVMe SSDs.

Intel VROC, combined with Intel RSTe 5.0 and VMD, allows bootable RAID on PCIe NVMe SSDs directly attached to the CPU PCIe lanes.

Two types of VROCs are supported on the Lenovo ThinkSystem V2:

- Onboard SATA ports with software RAID support (Intel VROC SATA RAID, formerly known as Intel RSTe)
- Onboard NVMe ports with software RAID support (Intel VROC NVMe RAID)
 - Standard: Intel VROC (VMD NVMe RAID), supports Intel SSDs only
 - Premium (requires an FoD key): Intel VROC (VMD NVMe RAID), which enables RAID support for non-Intel NVMe SSDs

For more information about VROC, refer to courses [ES51780C Servicing the ThinkSystem storage controllers](#) and [ES51757B Introducing ThinkSystem tools](#).

VROC settings in UEFI

To use Intel VROC, first enable **Intel VMD technology** in UEFI. You will then be able to see **Intel VROC controller** and **Intel Virtual RAID on CPU** under **System Settings → Storage**.

