Processor

Skylake processor features and configurations

Skylake server processor

- Internal architecture moves from ring to mesh topology.
- Six memory channels at 2666 MHz.
  - Operating speed depends on the processor model and UEFI operating mode. For example: 31xx only supports 2133 memory frequency.
- 4x L2 cache, lower last level cache (LLC) per core.
- Use Intel UltraPath Interconnect (UPI) as the new protocol to replace QPI 1.1.

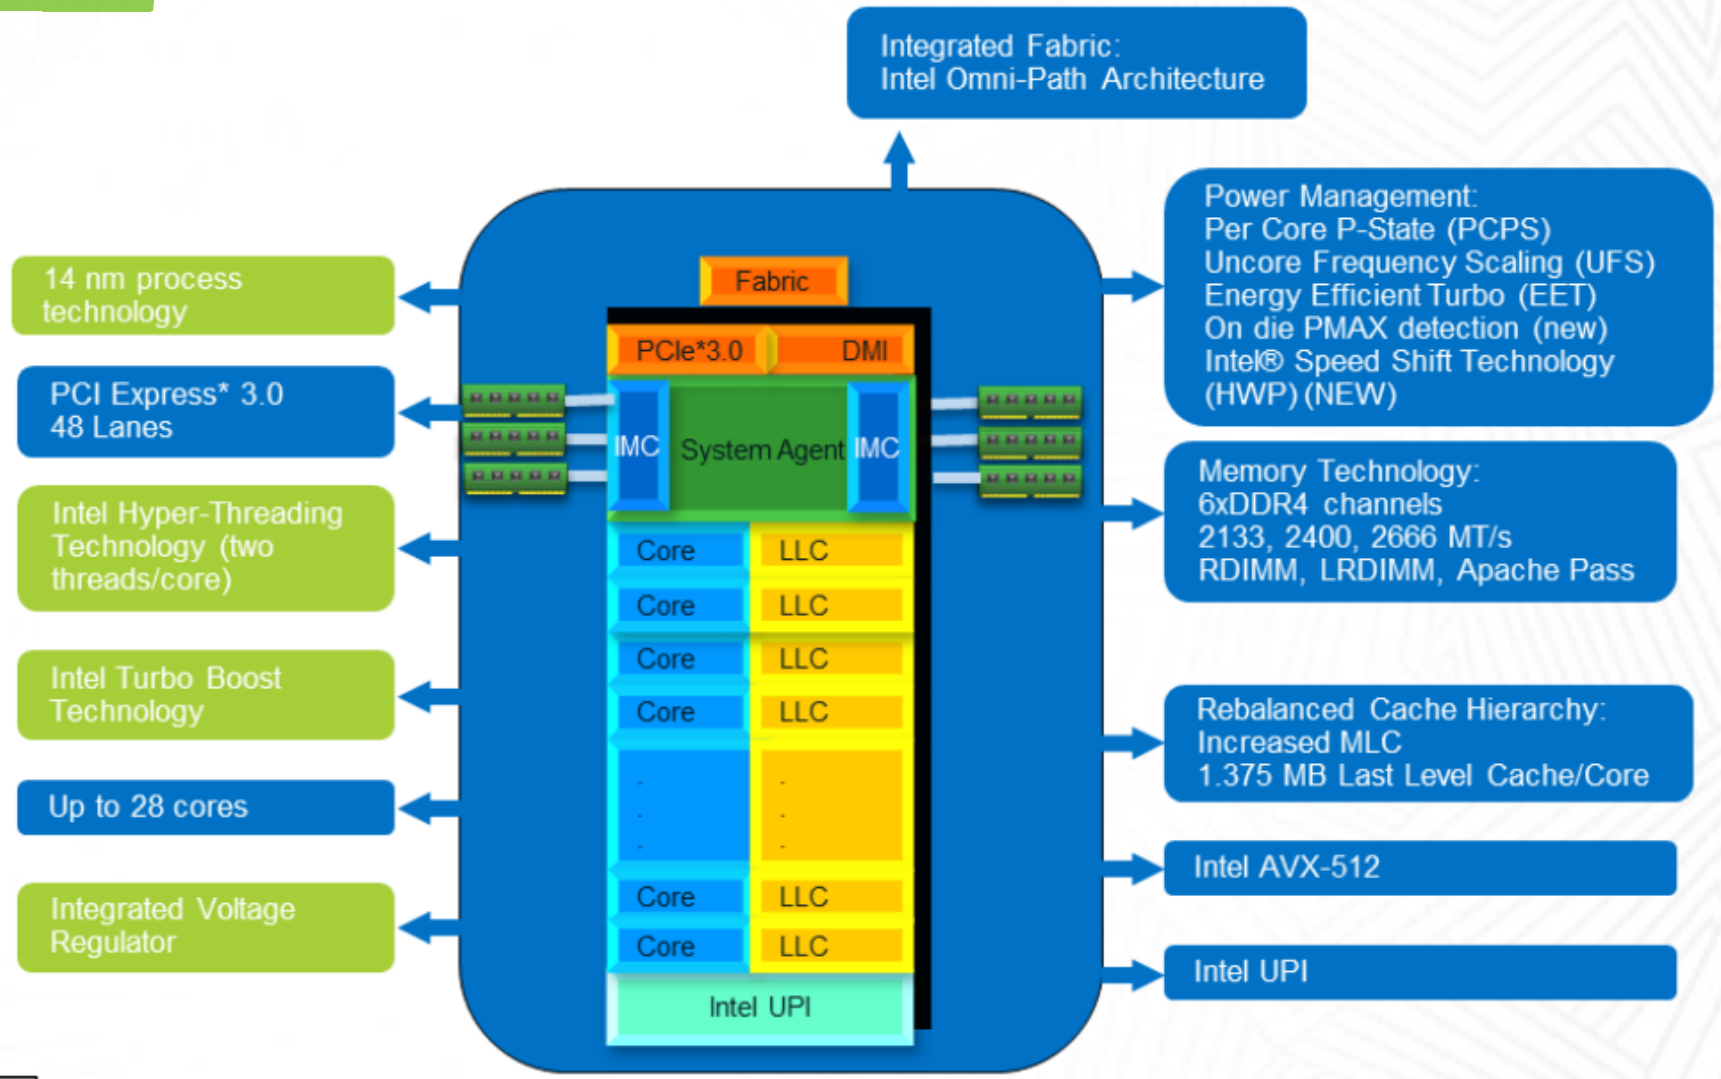

Information:

IMC represents integrated memory controller.

Processor product numbering

The image explains processor numbering rules.

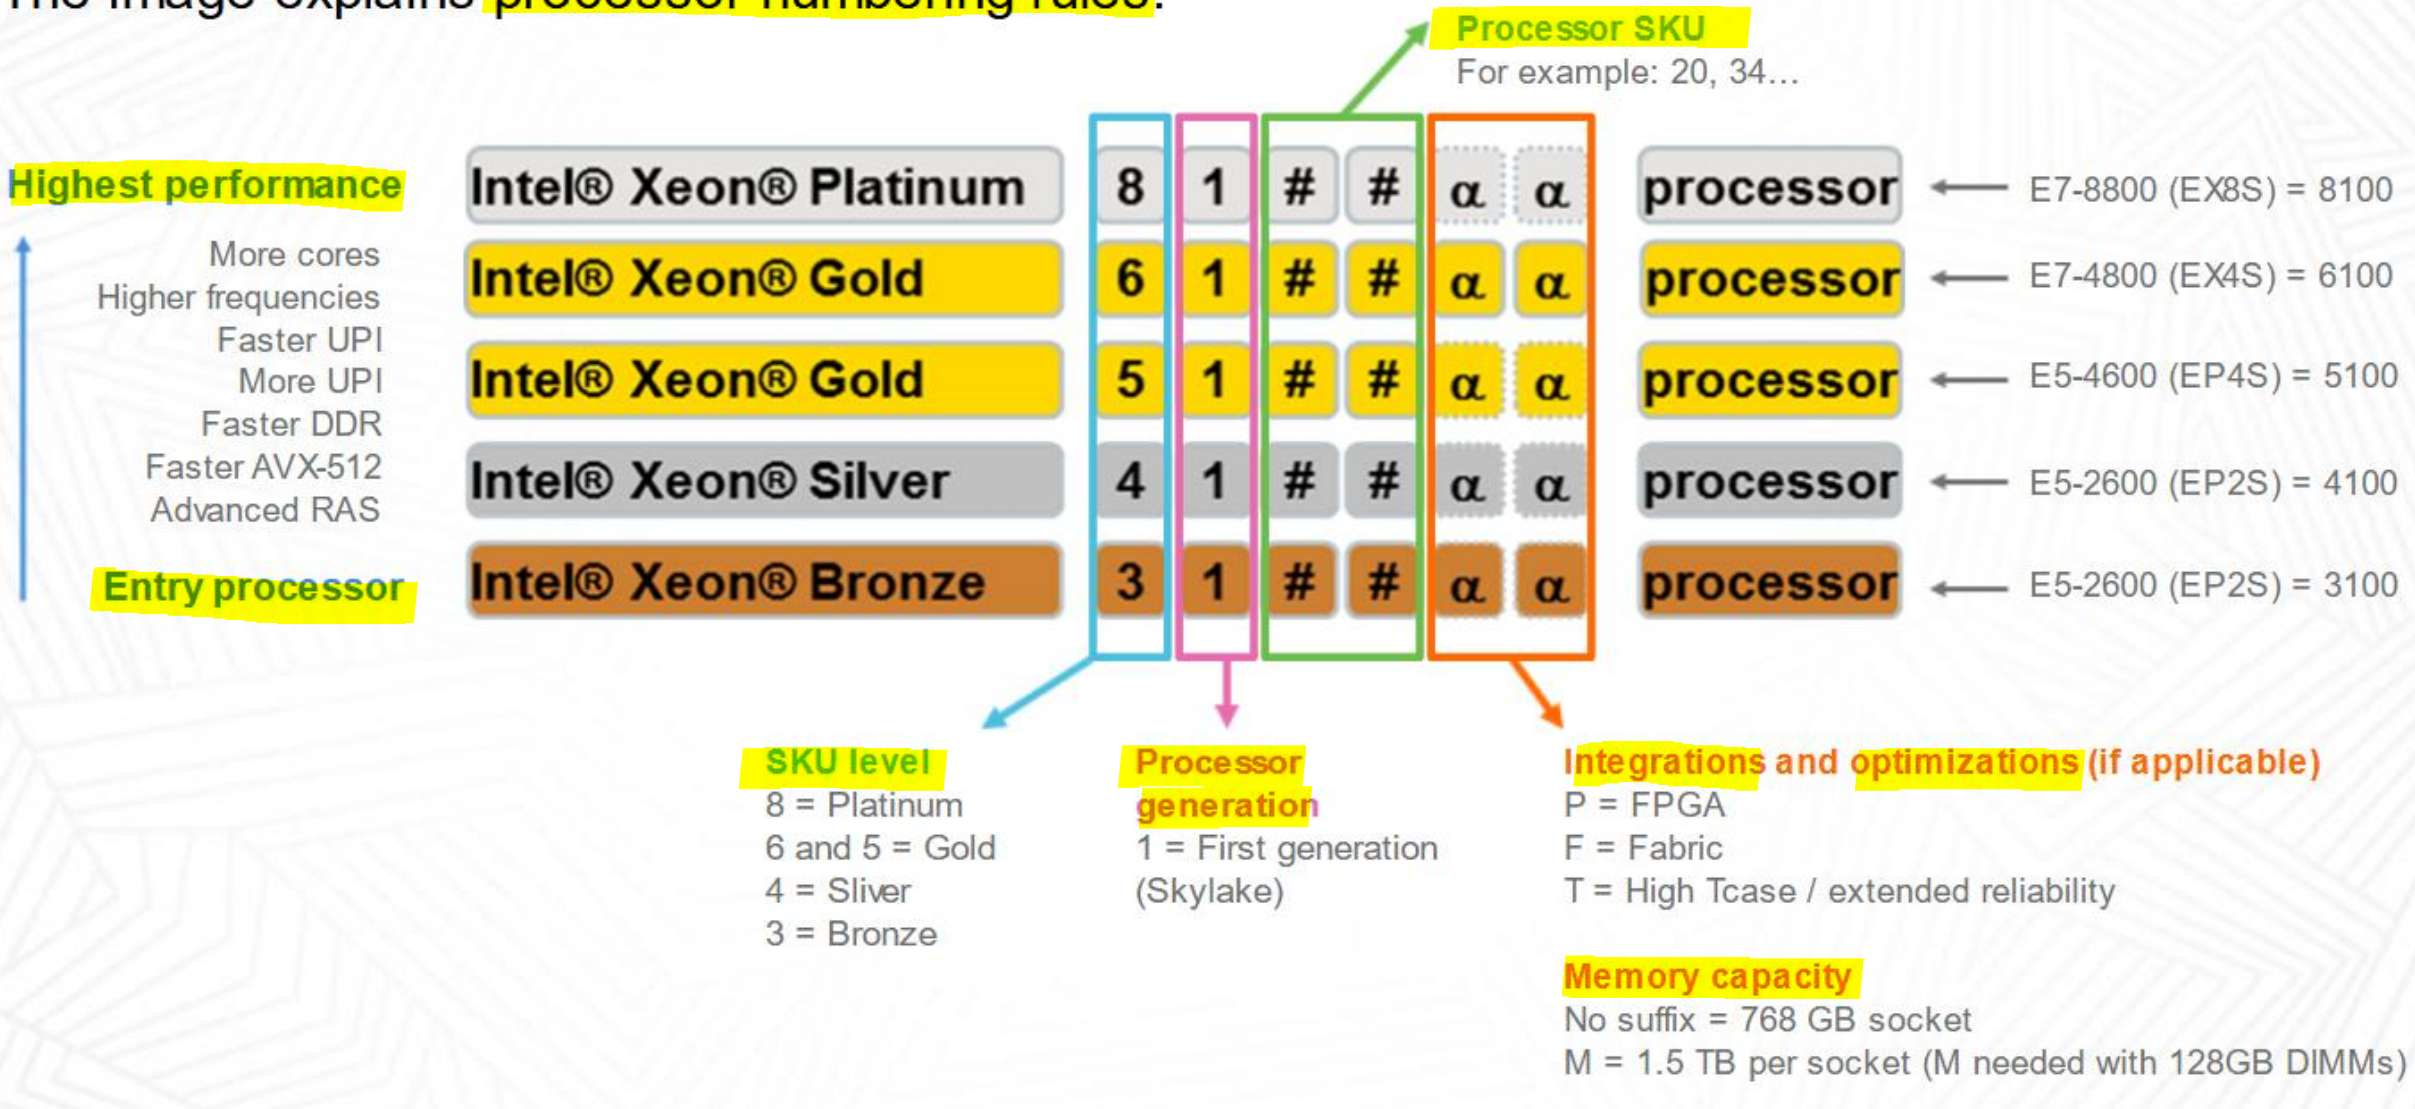

Processor levels and features

This is a feature comparison table of the processor scalable family.

Four-socket supported processors

Feature	81xx (Platinum)	61xx (Gold)	51xx (Gold)	41xx (Silver)	31xx (Bronze)
Number of Intel UPI links	Three	Three	Two	Two	Two
UPI speed	10.4 GT/s	10.4 GT/s	10.4 GT/s	9.6 GT/s	9.6 GT/s
Supported topologies	3 UPI → 2S-2UPI, 2S-3UPI, 4S-2UPI, 4S-3UPI, 8S-3UPI	3 UPI → 2S-2UPI, 2S-3UPI, 4S-2UPI, 4S-3UPI	2S-2UPI, 4S-2UPI	2S-2UPI	2S-2UPI
Node controller support	Yes	Yes	No	No	No
Number of memory channels	Six	Six	Six	Six	Six
DDR4 speed	2666	2666	2400*	2400	2133
Memory capacity	768 GB, 1.5 TB (select SKUs**)	768 GB, 1.5 TB (select SKUs**)	768 GB	768 GB	768 GB
RAS capability	Advanced	Advanced	Advanced	Standard	Standard
Intel Turbo Boost Technology	Yes	Yes	Yes	Yes	No
Inter Hyper-Threading Technology	Yes	Yes	Yes	Yes	No
Intel AVX-512 ISA support	Yes	Yes	Yes	Yes	Yes
Intel AVX-512 – # of 512b FMA units	Two	Two	One*	One	One
Number of PCIe lanes	48	48	48	48	48

* Gold processor #5122 supports 2666 DDR4 and two 512-bit FMA units.

** SKUs that support 1.5 TB per socket memory capacity are shown here.

Processor heat sink module

The processor heat sink module (PHM) refers to the subassembly where the heat sink and processor are clipped together prior to installation.

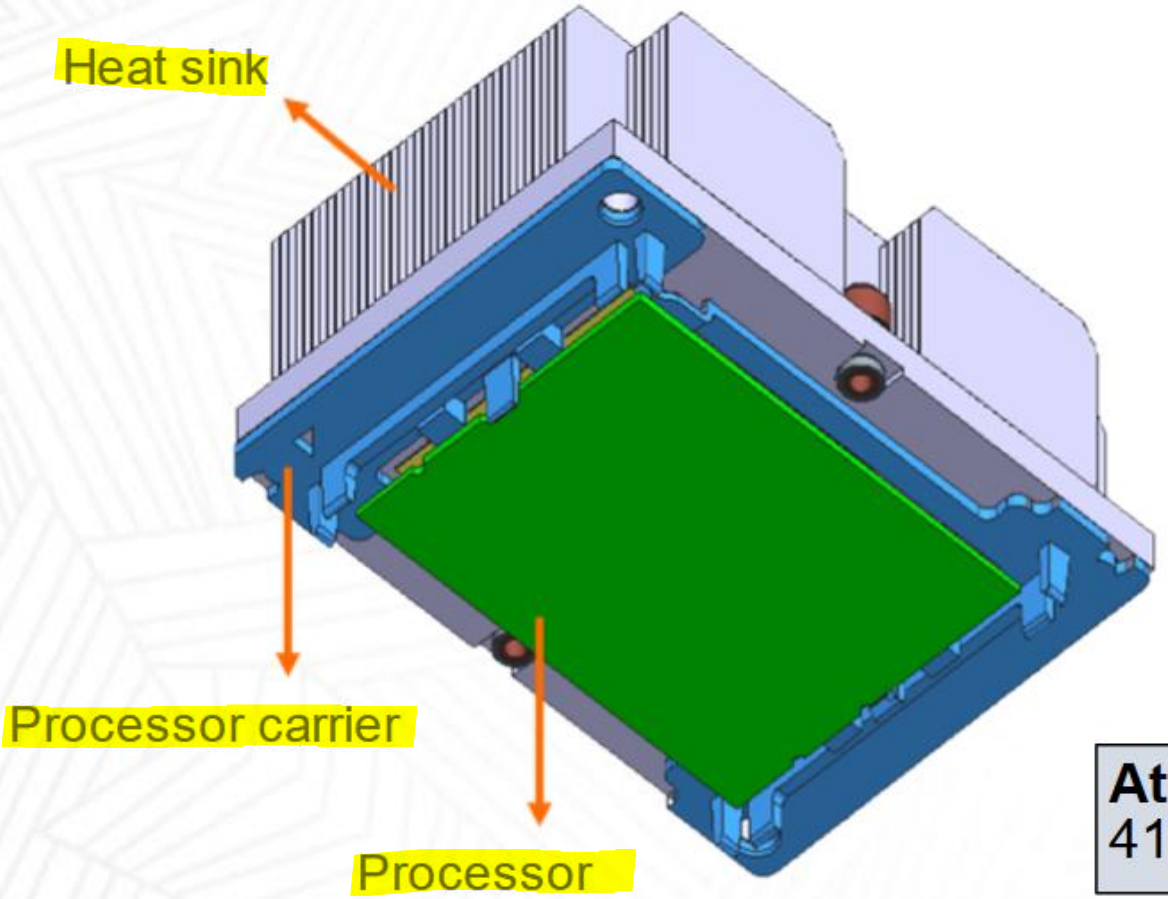

Click CPU FRUs and PHM assembly to view the details.

SKX standard

SKX-F

PHM to system board assembly

Attention: Use the gray thermal grease (part number 41Y9292) when replacing the processor.

Intel C620 series chipsets Southbridge (Lewisburg PCH)

Intel Xeon scalable processors with Intel C620 series chipsets deliver new levels of consistent performance. Here are the key features of Intel C620 series chipsets:

- Integrated Intel X722 Ethernet, up to Quad 10 Gbit (x8 uplink)
- Node Manager 4.0
- Direct media interface (DMI) x4 link is now Gen3, supporting up to 8 Gbit
- ThinkSystem servers use enhanced serial peripheral interface (eSPI) and not lower pin count (LPC) interface

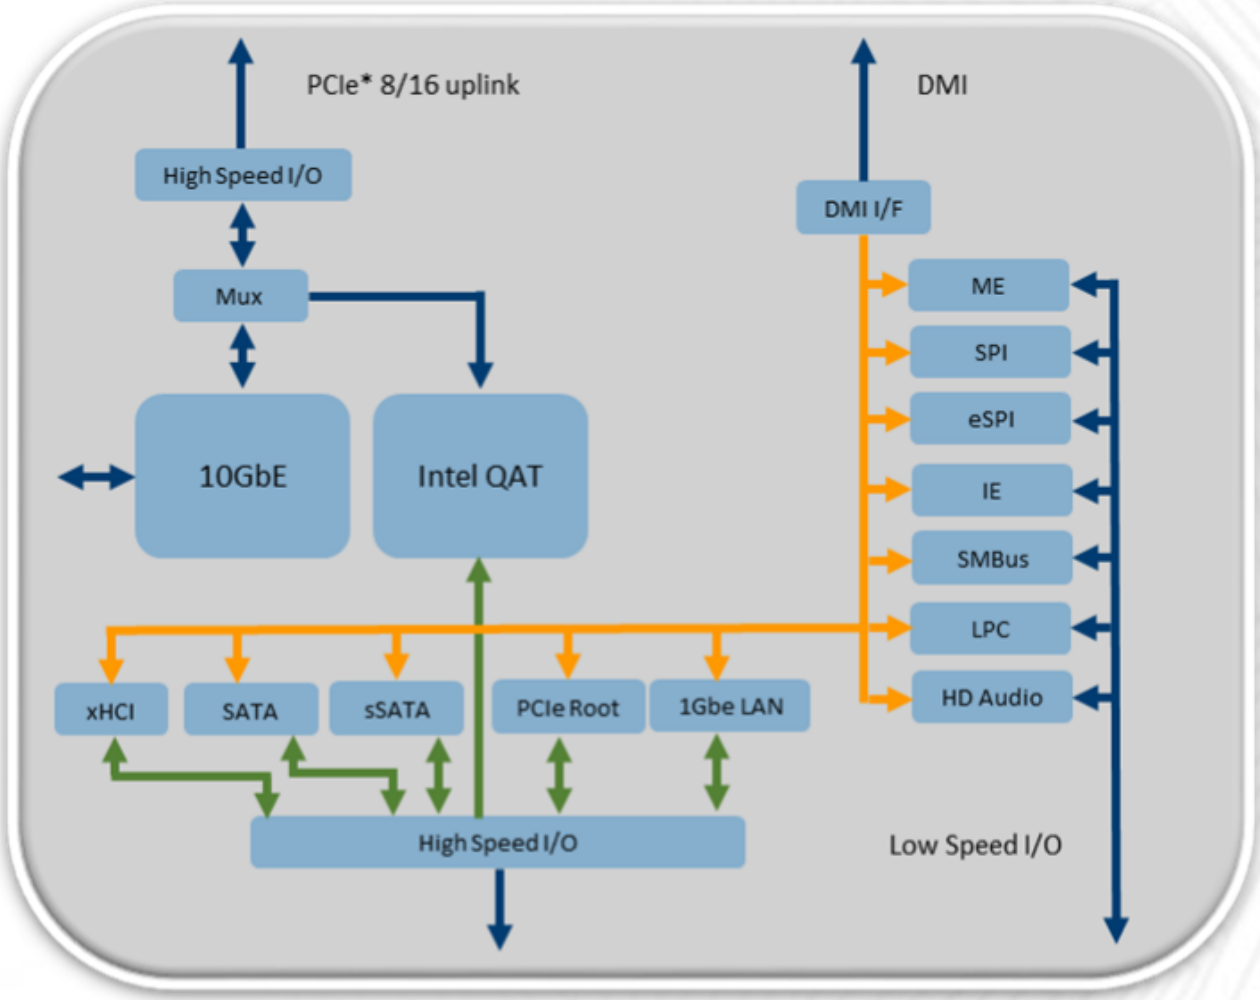

Typical 2S-2UPI configuration

This graphic shows two-socket platform configuration. In this example, a DIMM population is shown. Look up Apache Pass customer collateral for specific rules on DDR4/Apache Pass DIMM populations.

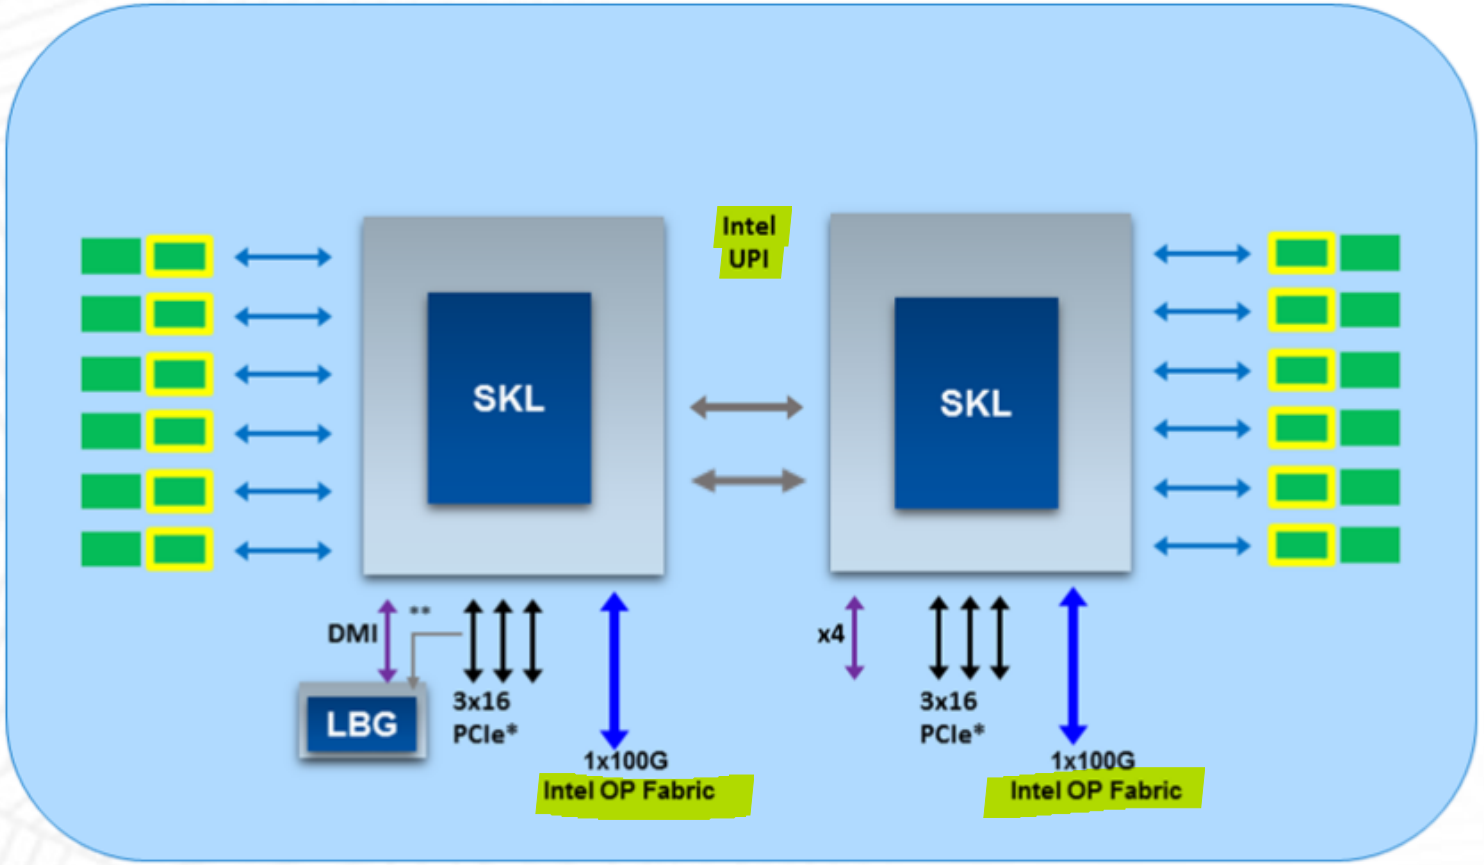

DDR4 DIMMs

DDR4/Apache Pass

** PCIe uplink connection for Intel QuickAssist Technology and Intel Ethernet

Information:

LBG represents Lewisburg chipset.

SKL represents Skylake processor.

Typical 4S configuration

This graphic shows four-socket platform configuration.

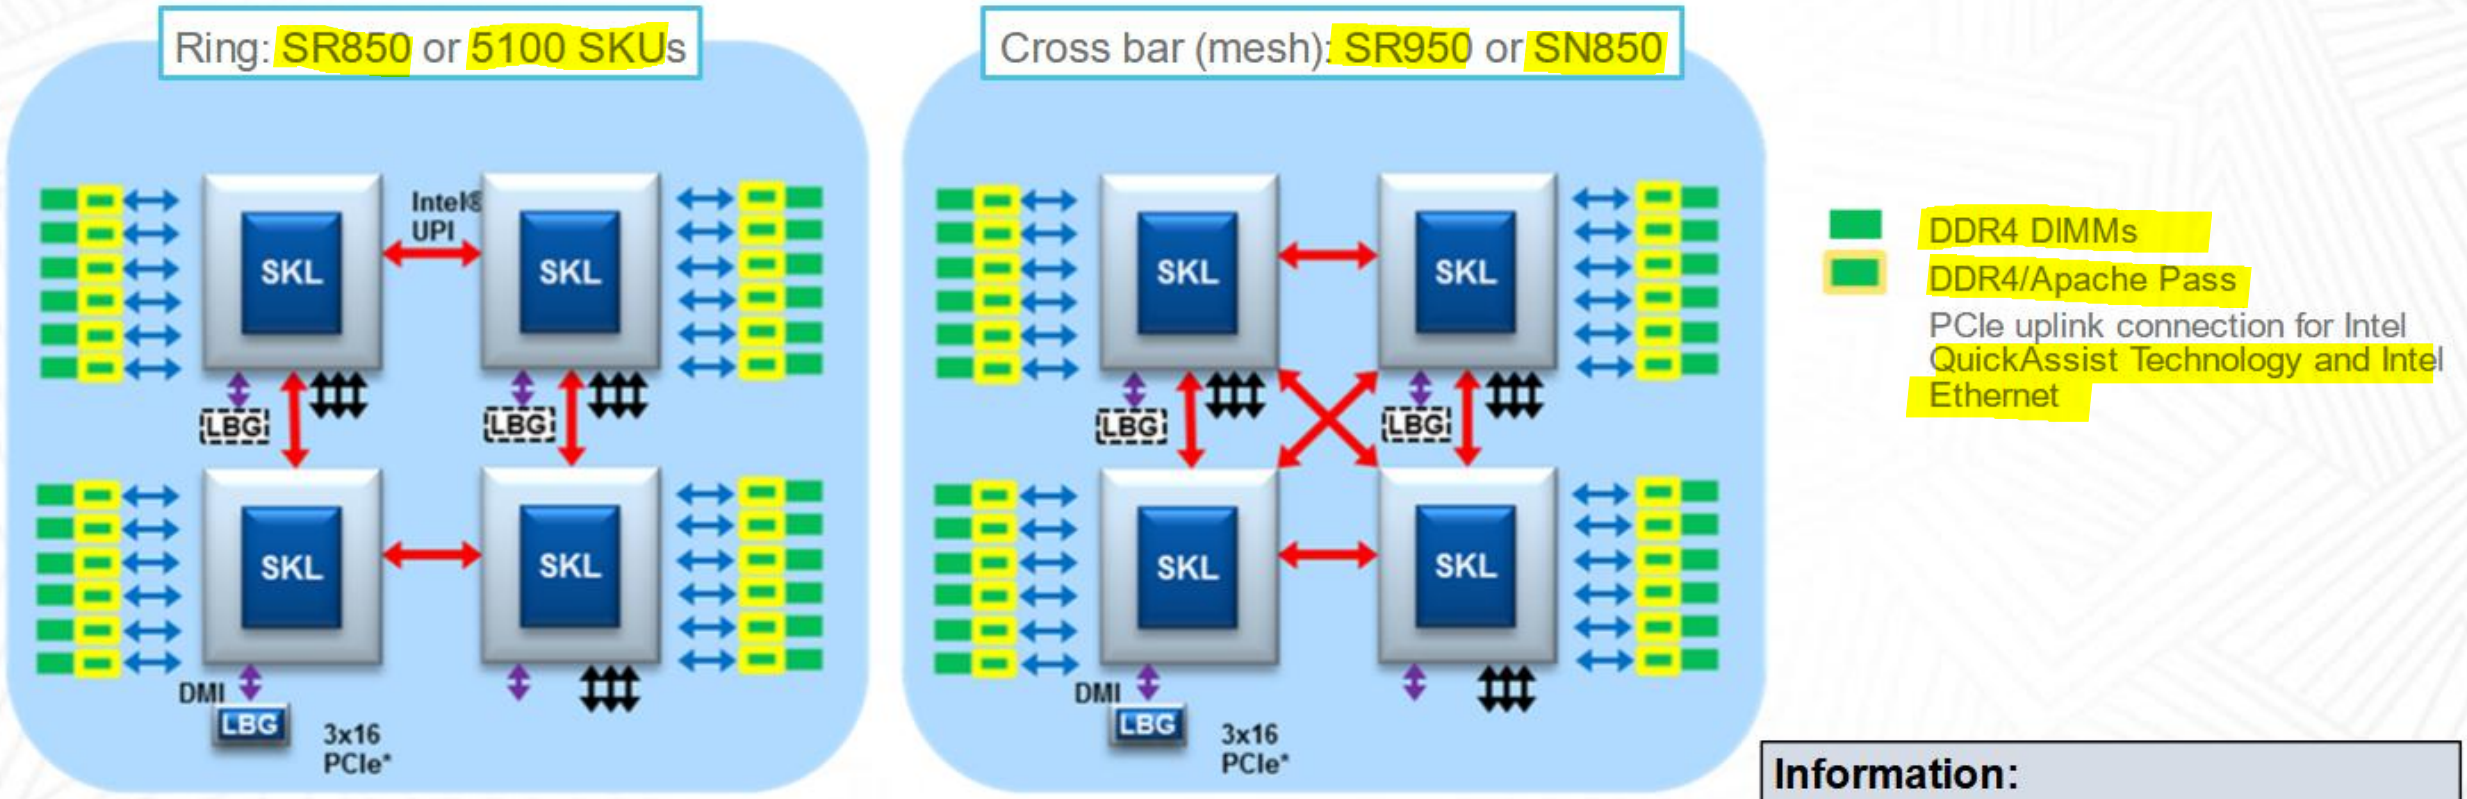

Information:

LBG represents Lewisburg chipset.

SKL represents Skylake processor.

Typical 8S configuration (SR950)

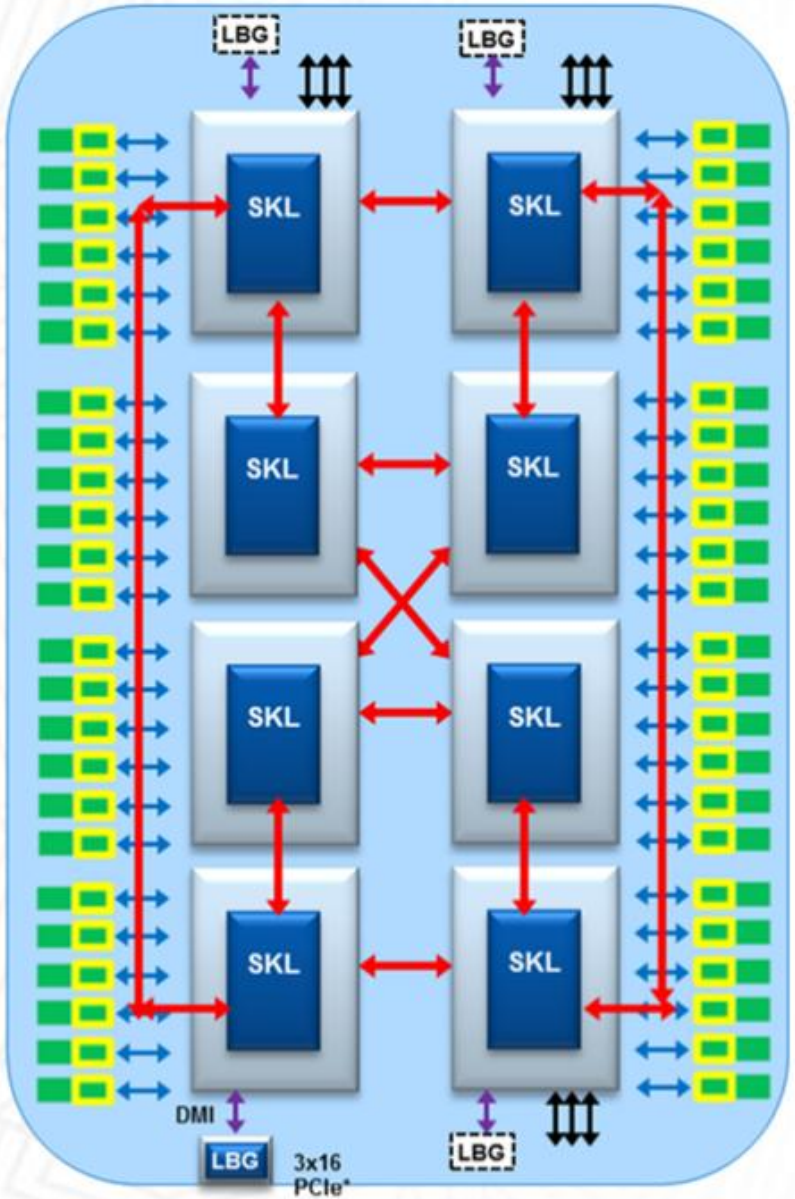

This graphic shows eight-socket platform configuration.

DDR4 DIMMs

DDR4/Apache Pass

LBG Optional
PCIe uplink connection for Intel QuickAssist Technology and Intel Ethernet

Information:

LBG represents Lewisburg chipset.

SKL represents Skylake processor.

RAS features

RAS features on 5100/6100/8100 SKUs

High-value RAS features on 5100/6100/8100 SKUs

- Adaptive Double Device Data Correction with Multiple Regions (ADDDC-MR) (Memory Self-Healing)
  - Requires two ranks of x4 DIMMs per DDR4 channel for best MR features.
  - Requires closed-page memory.
- Machine Check Architecture
  - Recovery for a limited set of faults
  - Error containment for improved reporting and isolation
- UPI Faildown (link width reduction, but server keeps running)
  - Link can retrain from 20 data lanes to eight data lanes
- Address Based Mirroring (6 GB per CPU when the option is selected)
  - Provides RAID1 protection on critical areas of memory like OS or Hypervisor

Note: Use Adaptive open-page if maximum performance mode is needed. However, you will not get full RAS capability while in maximum performance mode.