Open. Together.

OCP SUMMIT

PCI Express®: Enabling Performance in HPC

Richard Solomon
Vice President
PCI-SIG®

OPEN COMMUNITY®

Agenda

- PCI-SIG Overview
- PCI Express Overview
  - PCI Express 4.0
  - PCI Express 5.0
- Benefits of PCI Express in HPC
- Summary

PCI-SIG® Snapshot

Organization that **defines the PCI Express® (PCIe®) I/O bus specifications and related form factors.**

750+ member companies located worldwide.

Creating specifications and mechanisms to support compliance and interoperability.

- Australia
- Austria
- Belgium
- Brazil
- Bulgaria
- Canada
- China
- Czech Republic
- Denmark
- Finland
- France
- Germany
- Hong Kong
- Hungary
- India
- Ireland
- Israel
- Italy
- Japan
- Malaysia
- Norway
- Russia
- Singapore
- Slovak Republic
- South Korea
- Sri Lanka
- Sweden
- Switzerland
- Taiwan
- The Netherlands
- Turkey
- United Kingdom
- United States

BOARD OF DIRECTORS 2018-2019

AMD, DELL, Hewlett Packard Enterprise, IBM, intel, KEYSIGHT TECHNOLOGIES, NVIDIA, Qualcomm, SYNOPSYS® Silicon to Software™

I/0 BANDWIDTH DOUBLES

Every 3 Years

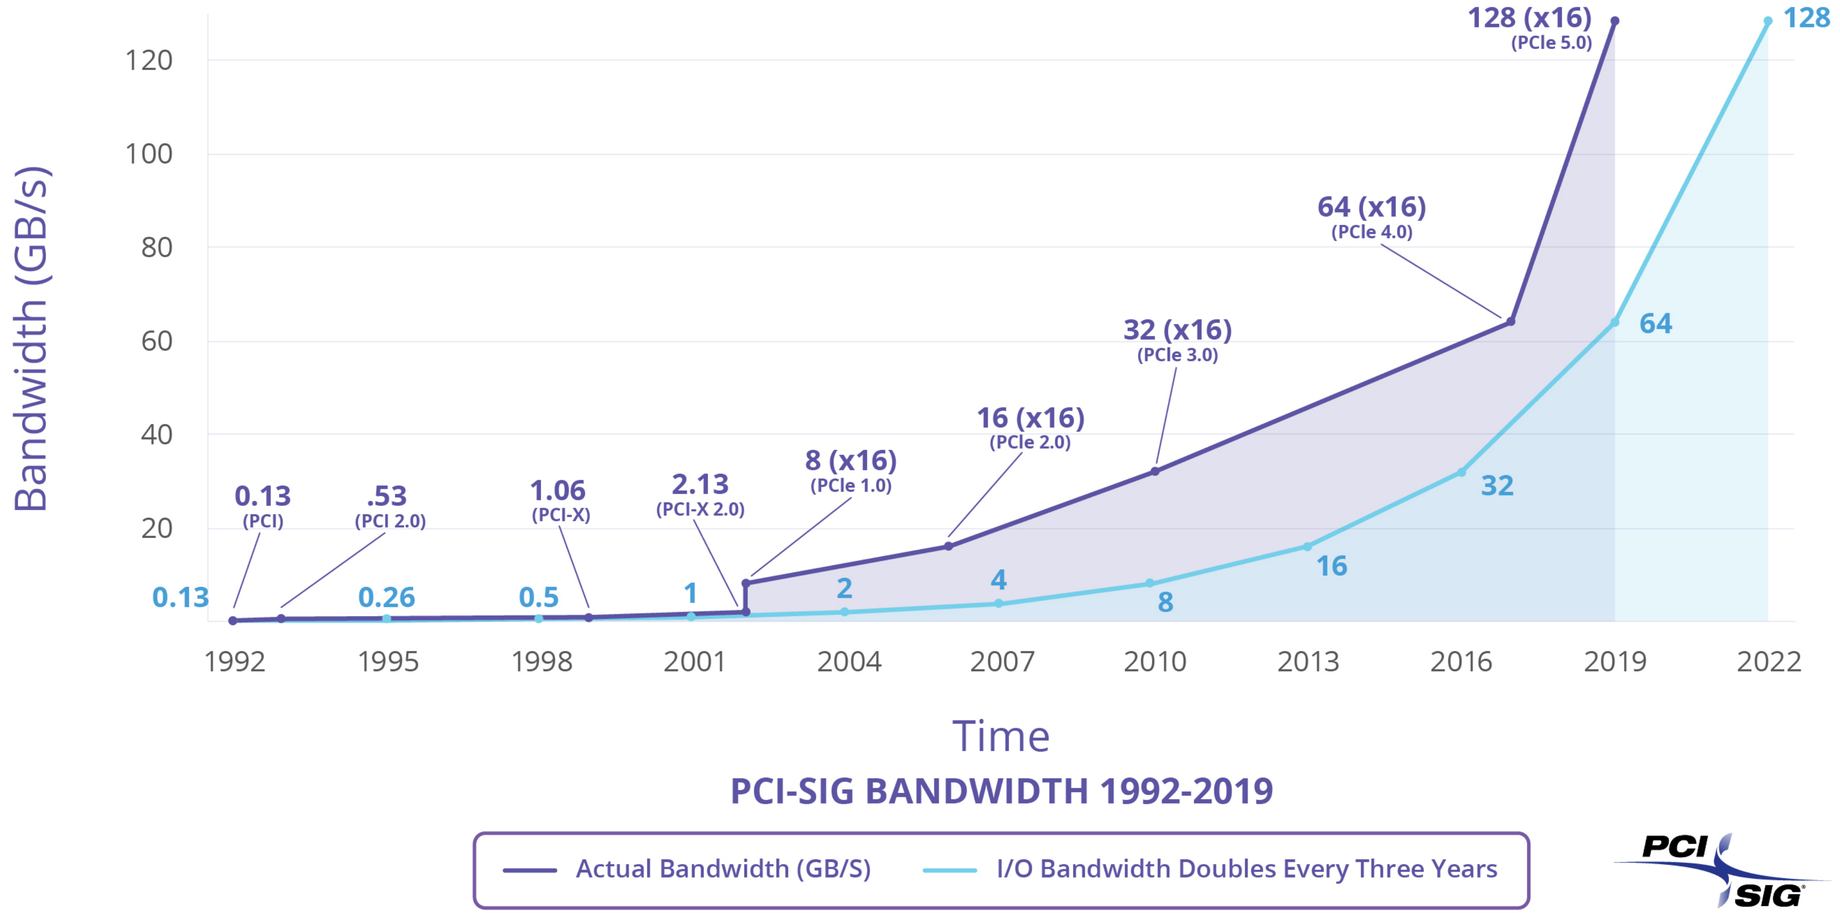

Market Segments Addressed by PCIe®

Artificial Intelligence
- High-performance
- High-bandwidth

Enterprise Servers
- Redundancy/failover
- Ubiquity
- Power savings

Cloud
- Scalable architecture
- Increased performance
- Reduced TCO

Automotive
- High-performance
- Reliability
- Availability
- Serviceability

PC/Mobile/IoT
- Faster performance
- Power efficiency

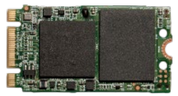

Storage
- Faster data transfer
- Better user experience
- Ubiquity

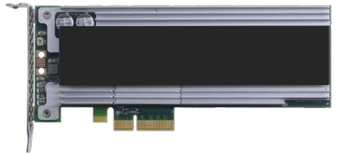

PCIe 4.0 Specification & Status

Adoption Is Well Under Way

Key Features:

- Delivers 16 GT/s
- Maintains backward compatibility with PCIe 3.x, 2.x, and 1.x
- Implements:
  - Extended tags and credits
  - Reduced system latency
  - Lane margining
- Superior RAS capabilities
- Scalability for added lanes and bandwidth
- Improved I/O virtualization and platform integration

Compliance Status:

- PCI-SIG Launched Official FYI Testing for PCIe 4.0 in December 2018
- >50 participants in pre-FYI testing
- Formal Compliance testing targeted for Q3 2019

Adoption:

- Numerous vendors with 16GT/s PHYs and controllers in silicon
- Test equipment from multiple vendors
- Several member companies have publicly announced & exhibited PCIe 4.0 products

PCI Express 5.0 Specification

PCI Express 5.0, Rev 0.9 Released in November 2018; Rev 1.0 on Target for Q1 2019

Key Features:

- Delivers 32 GT/s
- Maintains backward compatibility with PCIe 4.0, 3.x, 2.x, and 1.x

Changes limited primarily to speed upgrade

- Protocol already supports higher speed via extended tags and credits
- Electrical changes to improve signal integrity and mechanical performance of connectors
- CEM connector backwards compatible for add-in cards

	RAW BIT RATE	LINK BW	BW/ LANE/WAY	TOTAL BW X16
PCIe 5.0	32GT/s	32Gb/s	~4GB/s	~128GB/s
PCIe 4.0	16GT/s	16Gb/s	~2GB/s	~64GB/s
PCIe 3.x	8.0GT/s	8Gb/s	~1GB/s	~32GB/s
PCIe 2.x	5GT/s	4Gb/s	500MB/s	16GB/s
PCIe 1.x	2.5GT/s	2Gb/s	250MB/s	8GB/s

Benefits of Using PCIe in HPC

Bandwidth and performance

- PCIe 4.0 delivers 16 GT/s
- PCIe 5.0 delivers 32 GT/s

Flexibility/interoperability

- PCIe 4.0 and PCIe 5.0 maintain backwards compatibility

Data integrity

- Extensive logging and error reporting mechanisms (i.e. link cyclic redundancy check (LCRC) and end-to-end cyclic redundancy check (ECRC))

Shared I/O

- PCIe includes single-root IO virtualization (SR-IOV) requirements in the end-points, switches and root complexes
- I/O controllers can be shared among multiple hosts on a PCIe fabric by several different schemes

RAS Features

PCIe architecture supports a very high-level set of Reliability, Availability, Serviceability (RAS) features

- All transactions protected by CRC-32 and Link level Retry, covering even dropped packets
- Transaction level time-out support (hierarchical)
- Well defined algorithm for different error scenarios
- Advanced Error Reporting mechanism
- Support for degraded link width / lower speed
- Support for hot-plug

Lane Margining

Problems with higher data rates:

- Limited transmission distances
- Degraded signal integrity
- Lower manufacturing yield

PCIe 4.0 added lane margining at the receiver

- Allows the system to determine how close to “the edge” each lane is operating under real conditions

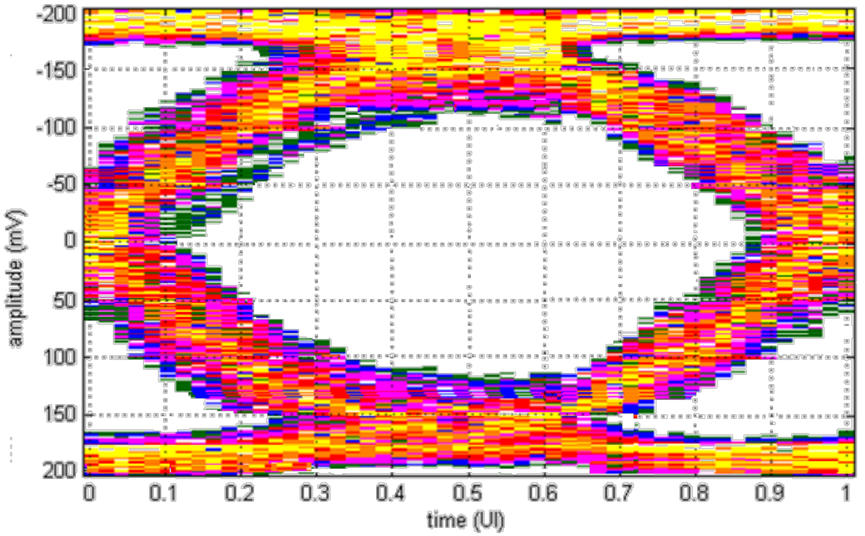

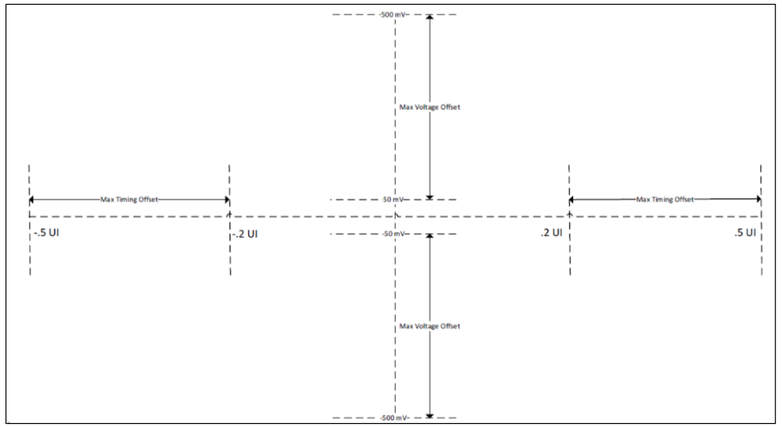

PCIe in OCP Use Cases

- OCP 3.0 Pinout and PCIe Bifurcation
- OCP Mezzanine card v0.5, original standard – features PCIe 3.0 interface
- Learn more about OCP server projects utilizing PCIe technology on the Wiki: https://www.opencompute.org/wiki/Server/Mezz

Summary

- PCIe brings many benefits to HPC, with its high bandwidth, flexibility, data integrity and more
- PCIe 4.0 adoption is well under way
- PCIe 5.0 is on track for completion in Q1 2019
- PCI-SIG continues to maintain its leadership position in delivering high-performance, low power I/O that meet demanding markets like servers

Back-Up Slides

Form Factors for PCI Express®

M.2

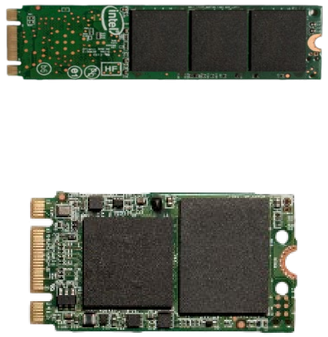

42, 80, and 110mm lengths, smallest footprint of PCIe connector form factors, use for boot, for max storage density, for PXI/AXIe ecosystem

U.2

2.5in makes up the majority of SSDs sold today because of ease of deployment, hotplug, serviceability, and small form factor Single-Port x4 or Dual-Port x2

CEM Add-in-card

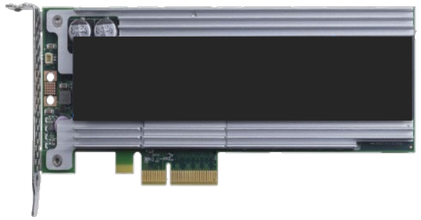

Add-in-card (AIC) has maximum system compatibility with existing servers and most reliable compliance program. Higher power envelope, and options for height and length

How Does Adding PCIe Affect the Data Center?

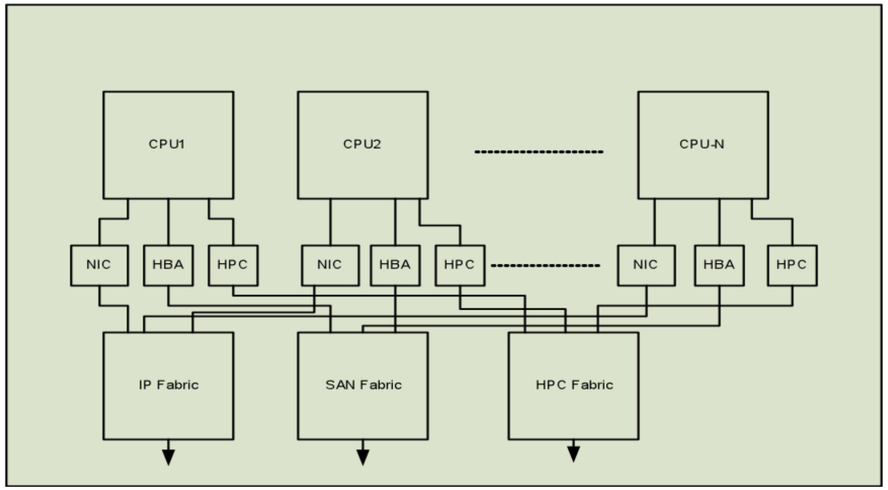

Typical Data Center I/O Interconnect

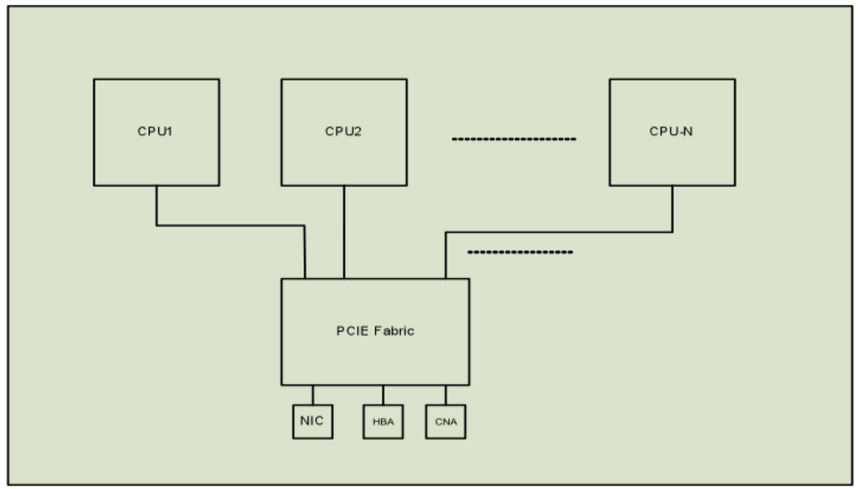

PCI Express-based Server Cluster

Source: A Case for PCI Express as a High-Performance Cluster Interconnect


Open. Together.
OCP Global Summit | March 14–15, 2019
OCP
SUMMIT