The PCIe® Specification: A High Bandwidth Interconnect Solution for AI and ML

Presenter: Richard Solomon, PCI-SIG® Vice President and Compliance Chair

PCI-SIG® Snapshot

- Organization that **defines the PCI Express® (PCIe®) I/O bus specifications and related form factors**
- PCI-SIG began 32 years ago in 1992
- The PCIe specification was first released in 2003
- **960+** member companies located worldwide
- Creating specifications and mechanisms to **support compliance** and **interoperability**

Board of Directors
2024 –2025

AMD
arm
DELL Technologies
IBM
intel
KEYSIGHT TECHNOLOGIES
NVIDIA
Qualcomm
SYNOPSYS®
Silicon to Software™

FMS

PCI-SIG® Roadmap

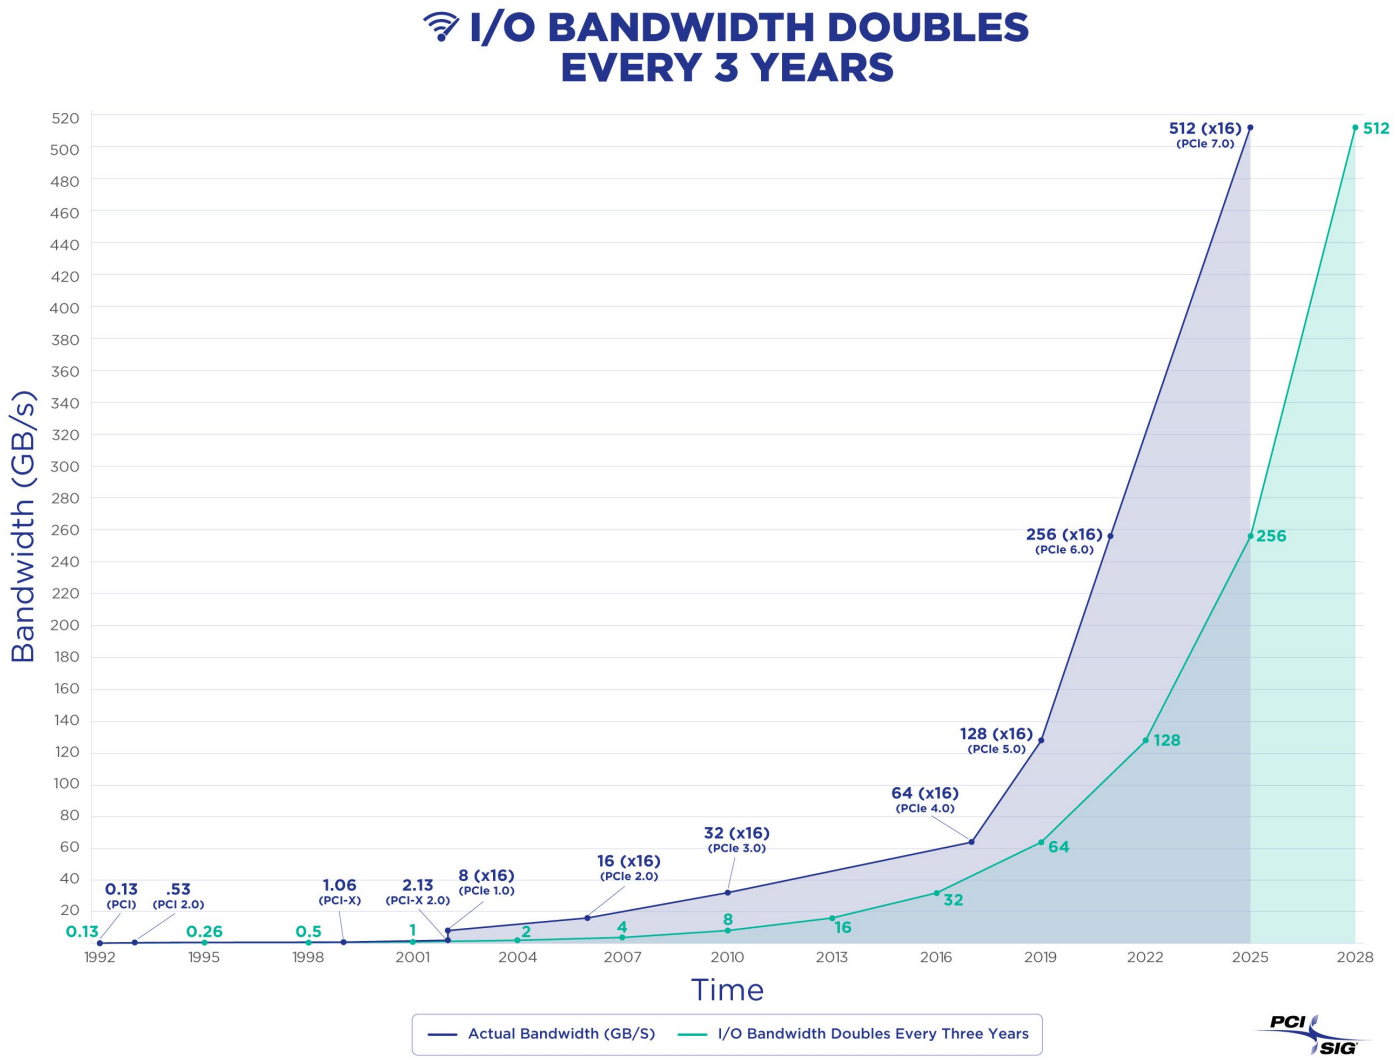

PCI Express® 6.0 Specification & Status

Released in January 2022

Key Features:

- 64 GT/s raw data rate and up to 256 GB/s via x16 configuration; doubles the bandwidth and power efficiency from PCIe 5.0 specification (32GT/s)
- Pulse Amplitude Modulation with 4 levels (PAM4) signaling and leverages existing PAM4 already available in the industry
- Lightweight Forward Error Correct (FEC) and Cyclic Redundancy Check (CRC) mitigate the bit error rate increase associated with PAM4 signaling
- Flit (flow control unit) based encoding supports PAM4 modulation and enables more than double the bandwidth gain
- Updated Packet layout used in Flit Mode to provide additional functionality and simplify processing
- Maintains backward compatibility with all previous generations of PCIe architecture

Adoption:

- Various PCIe 6.0 specification enabled products available in the industry; visit our sponsors for more information

Compliance Program Progress:

- FYI Testing introduced in Q1 2024

PCI Express® 7.0 Specification & Status

PCIe 7.0 specification, version 0.5 is live for PCI-SIG® members; The full PCIe® 7.0 specification is targeted for release in 2025

- **What does Version 0.5 mean?**
  - The official first draft of the specification, incorporating all the feedback we received from members after the release of Version 0.3, is available to PCI-SIG members.

Feature Goals:

- Delivering 128 GT/s data rate and up to 512 GB/s bi-directionally via x16 configuration
- Utilizing PAM4 signaling
- Defining the channel parameters
- Continuing to deliver the low-latency and high-reliability targets
- Improving power efficiency
- Maintaining backwards compatibility with all previous generations of PCIe technology

Revision	Max Data Rate	Encoding	Signaling
PCIe 7.0 (2025)	128.0 GT/s	1b/1b (Flit Mode*)	PAM4
PCIe 6.0 (2022)	64.0 GT/s	1b/1b (Flit Mode*)	PAM4
PCIe 5.0 (2019)	32.0 GT/s	128b/130b	NRZ
PCIe 4.0 (2017)	16.0 GT/s	128b/130b	NRZ
PCIe 3.0 (2010)	8.0 GT/s	128b/130b	NRZ
PCIe 2.0 (2007)	5.0 GT/s	8b/10b	NRZ
PCIe 1.0 (2003)	2.5 GT/s	8b/10b	NRZ

(*Flit Mode also enabled in other Data Rate with their respective encoding)

PCIe® Technology Enables Generative AI

What is Generative AI?

- Generative AI is a type of AI technology that can produce content, including text, images, video, audio and more

Challenges of Generative AI:

- Need for low power, low latency technologies to connect these systems together
- Due to the continuing increase in complexity and scale of these Large Language Models (LLMs), the most advanced Generative AI models can't fit on one GPU, one server, one rack, or even a single data center anymore

PCIe technology is a ubiquitous interconnect serving as the building block of Generative AI

- PCIe technology can make the entire data center look like one pooled resource of compute, memory, storage and more
- As data rates increase, and electrical reaches decrease, CopprLink Internal and External Cables can extend the reach of PCIe signals within Generative AI applications
- The PCI-SIG Optical Work Group is investigating a path for enabling PCIe technology over optical links to ensure any PCIe link will be possible

Key PCIe Features for Generative AI:

- Low Latency
- PAM4 signaling
- FEC
- Flit Mode
- Low Power Mode
- Backwards Compatibility

PCIe® Technology in AI Chipsets

AI Chipsets Available in PCIe Card Format

(Source: ABI Research)

Manufacturer	Description
AMD	Xilinx Alveo U280 features PCIe 4.0 and HBM2 memory bandwidth to provide high-performance, adaptable acceleration for ML inference.
Cambricon Technologies	Cambricon MLU-1000C Series is a smart computing card that can be plugged into a server via a PCIe 3 x16 slot.
Graphcore	The Graphcore C2 IPU PCIe Card runs two of Graphcore's Colossus GC2 IPU processors that can support both AI training and inference workloads.
Groq	The GroqCard accelerator packages a single GroqChip into a standard PCIe 4 x16 form factor for plug-and-play low-latency and scalable performance.
Habana Labs (Intel)	All Habana Labs' AI training and inference chipsets, namely Gaudi, Gaudi 2, Goya, and GRECO, are available in PCIe cards.
Hailo	The Hailo-8R Mini PCIe Module is an AI accelerator module that can be plugged into an existing edge device with a mini PCIe Full-Mini socket to execute in real time and with low-power deep neural network inferencing for a broad range of market segments.
IBM	The IBM Artificial Intelligence Unit (AIU) is a scaled-up Telum chip packaged in a standard PCIe slot for AI training and inference in the data center.
NVIDIA	All NVIDIA GPUs are available in PCIe format.
Qualcomm	The Qualcomm Cloud AI100 is a low-profile PCIe card focusing on inference workload for computer vision and NLP.

- All AI chipset vendors serious about making inroads into the data center ecosystem have offered their AI chipsets in a PCIe card format
- PCIe architecture is a widely adopted chip-to-chip interconnect protocol that reduces interoperability challenges and fully allows end users to leverage heterogenous computing for AI
- Combination of CPU, GPU, and AI accelerators creates an environment of heterogeneous computing that would benefit significantly from a standard interconnect technology

PCIe® Technology AI Market Data Forecasted Through 2030

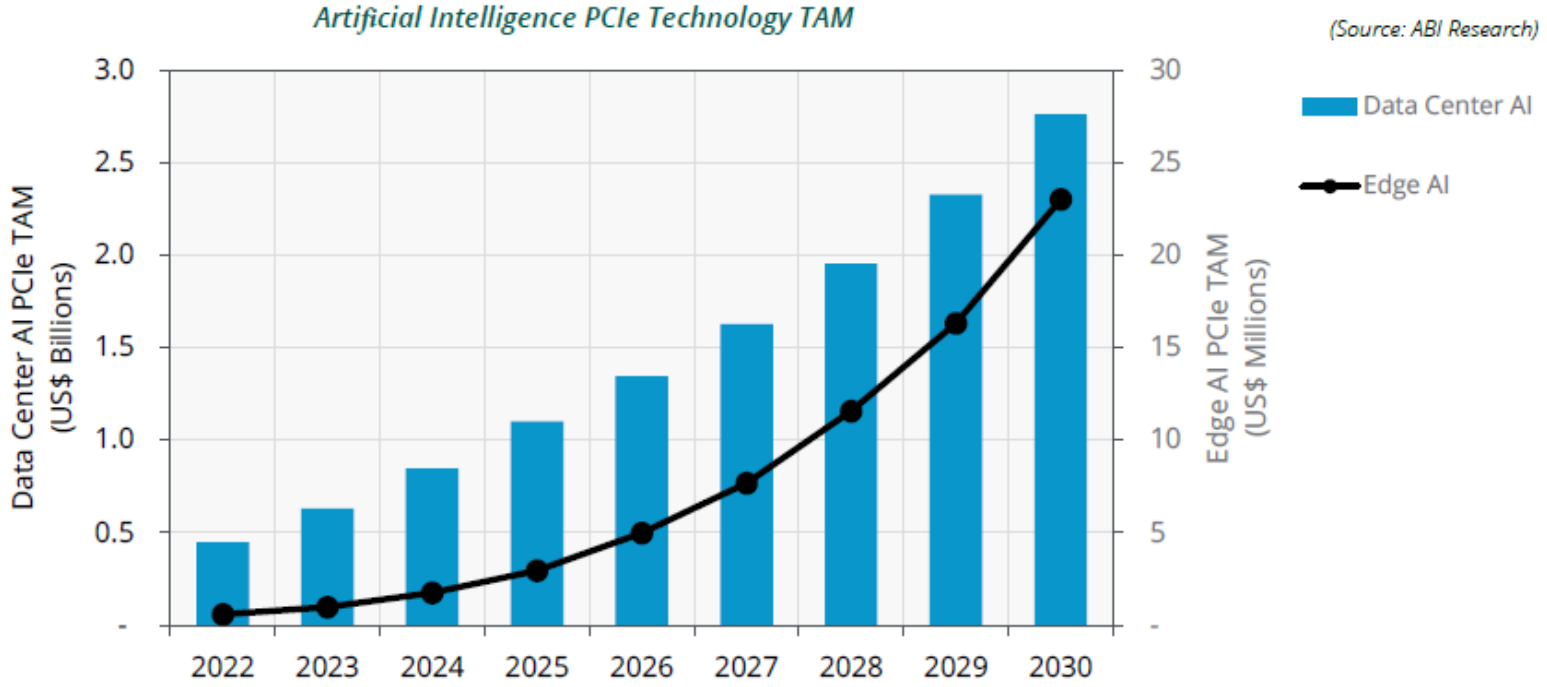

- **AI industry adoption will be high**, as PCIe technology offers decision makers agility through forwards and backwards compatibility, **improving time-to-value and lowering deployment risk**
- Total Addressable Market (TAM) for PCIe technology in AI, (edge and data center deployments), projected to grow to **US $2.784 billion by 2030, at a CAGR of 22%**
- Edge AI market is likely to grow more rapidly, at a CAGR of 50%, as more enterprise verticals deploy edge servers and AI continues to democratize

Engage with PCI-SIG® on Social Media

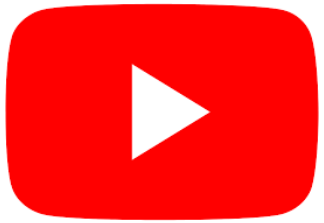

Backup

PCIe® Technology Usage is Growing in AI

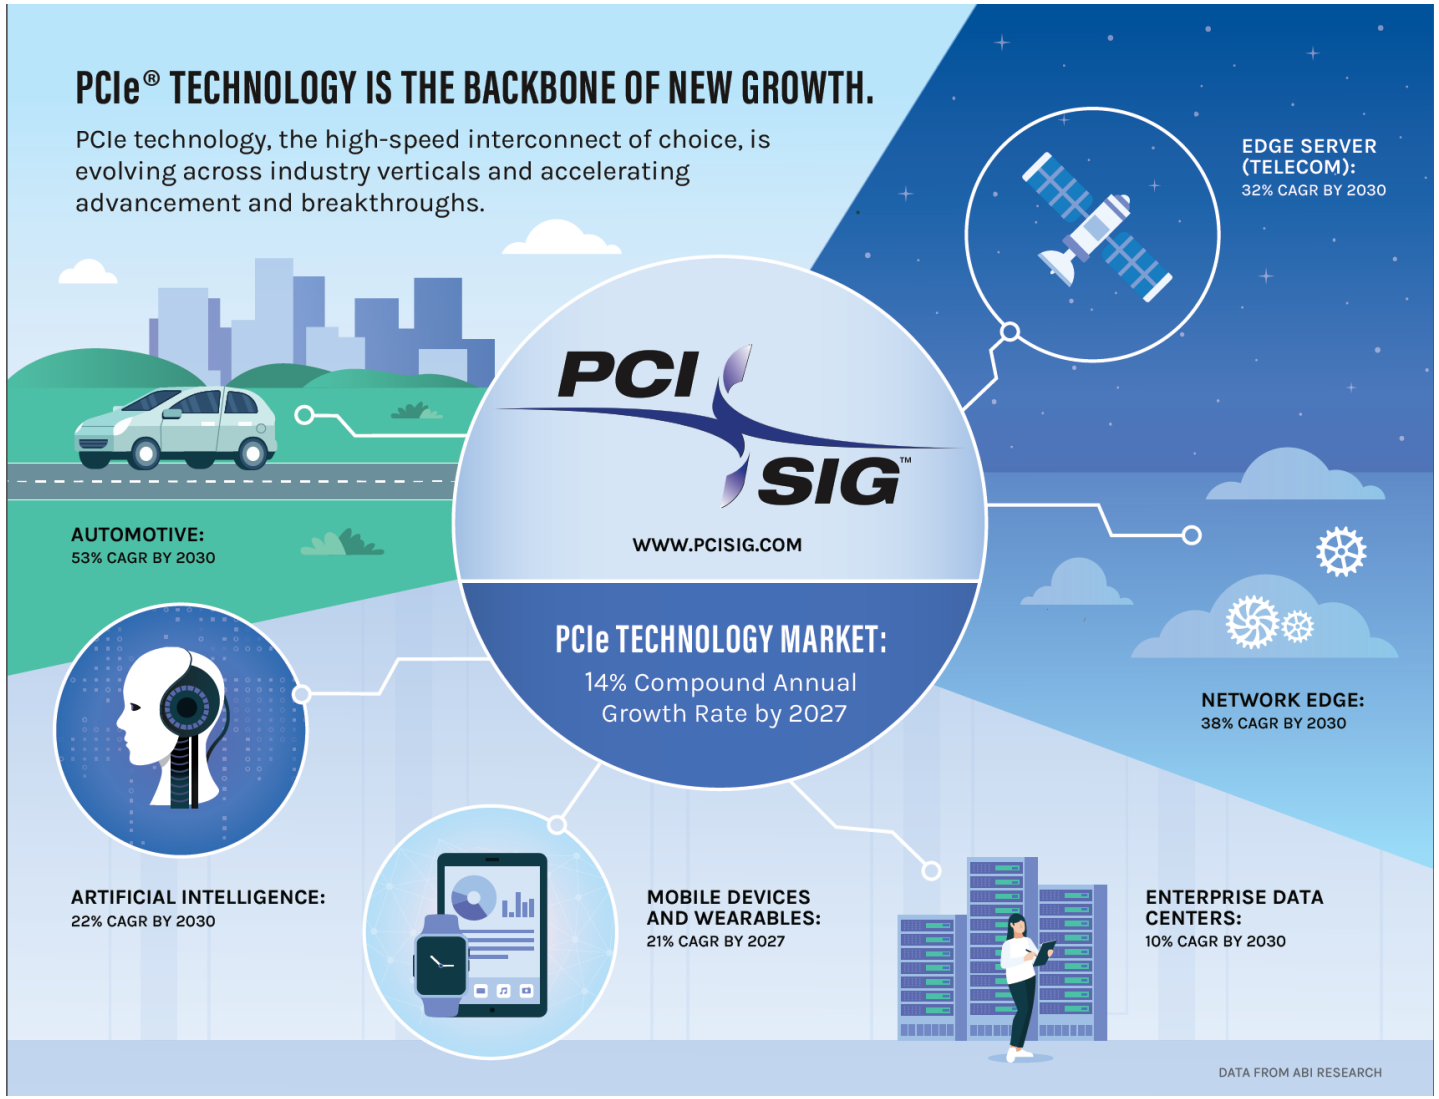

Looking Toward the Future
PCI Express® Architecture Adoption

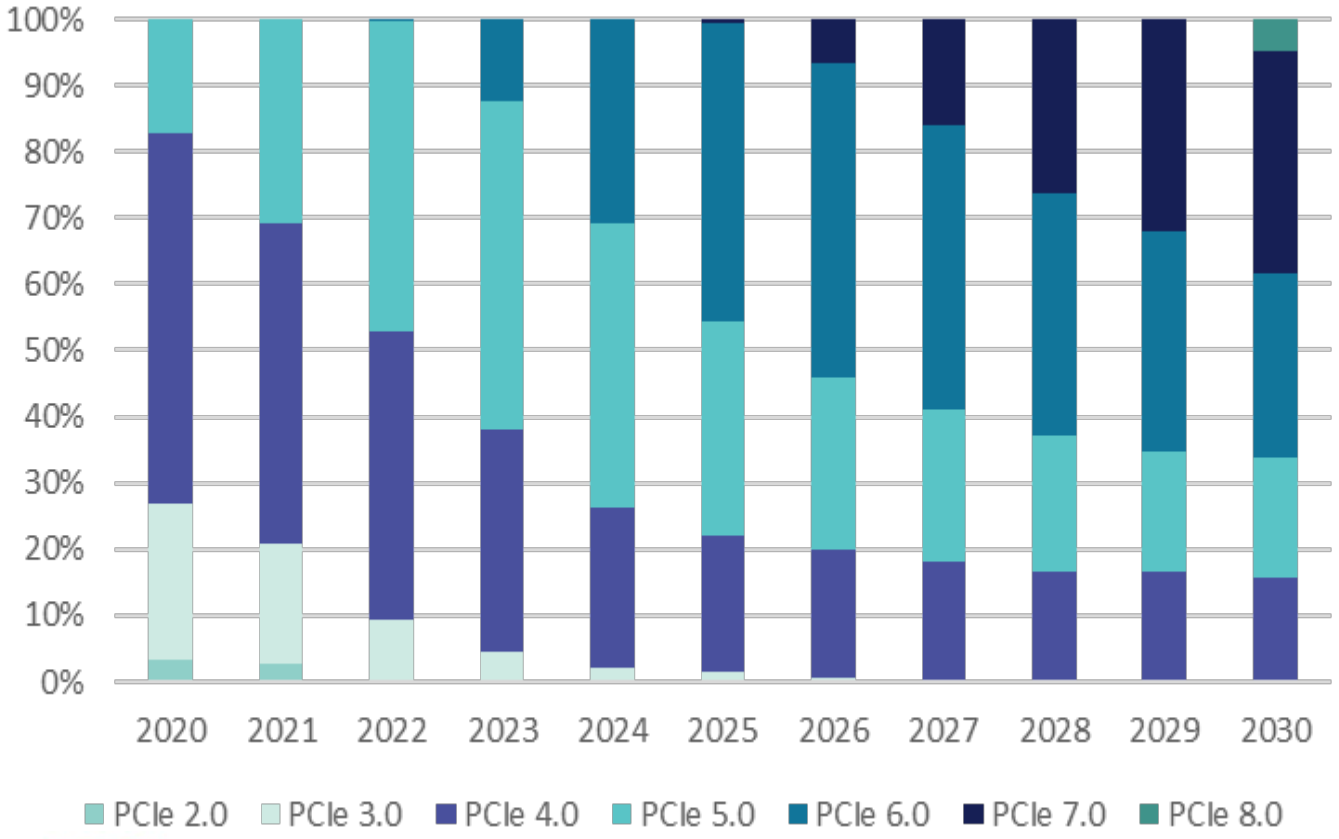

* Data provided by Synopsys, Inc.; used with permission

- Estimated usage across all market segments & applications
- PCIe® 8.0* architecture is a placeholder for a future generation after PCIe 7.0 architecture
  - *Not an endorsement or announcement by PCI-SIG®

Future: PCI Express® Architecture Usage by Market Segment

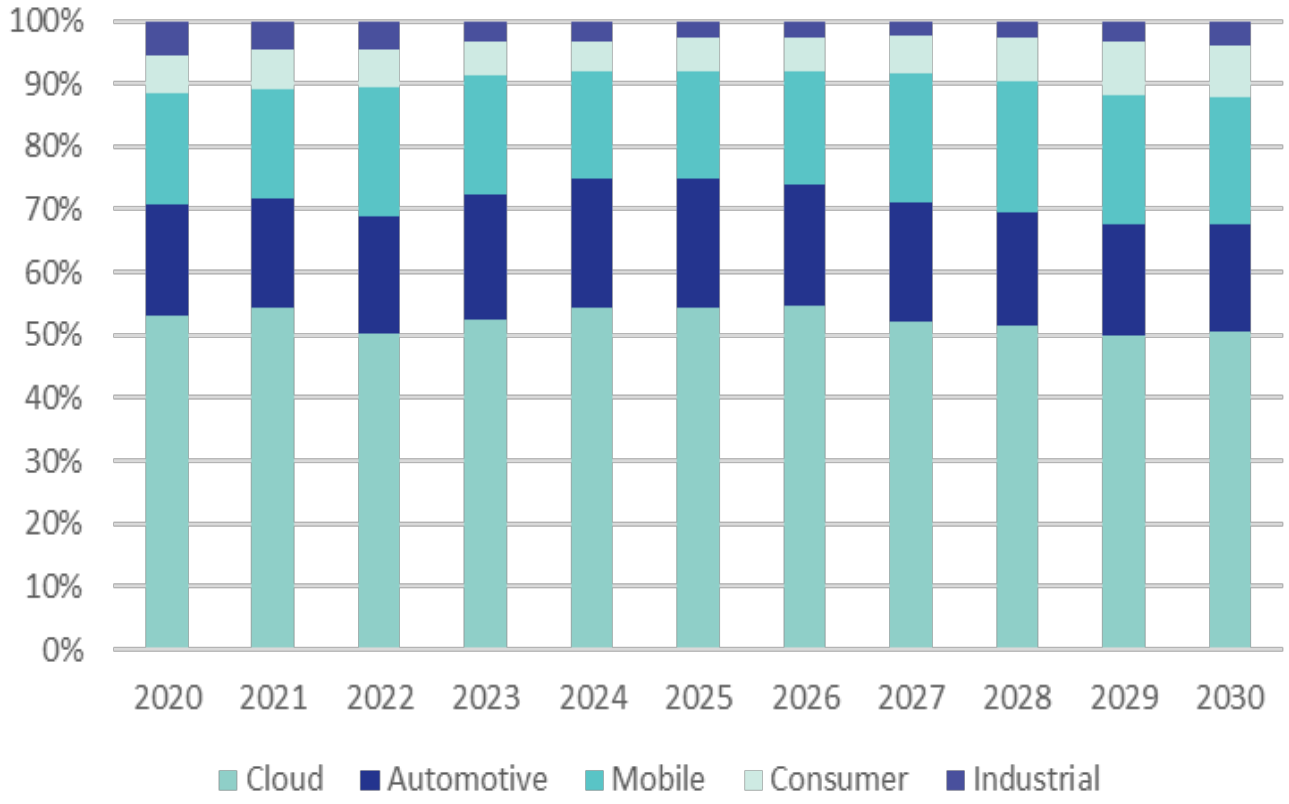

- Estimated usage across all market segments and applications based on market trends, data req's, etc.

* Data provided by Synopsys, Inc.; used with permission

PCI Express® Technology in AI Provided by Amphenol

CHALLENGES

- AI architectures are increasingly decentralized and composable
  - Require high data bandwidth, low latency, standard interconnect protocol
  - Link CPUs, GPUs, and specialized accelerators, NICs, storage, within and between racks

IMPLEMENTATION

- AI driven demand for greater interconnect bandwidth is accelerating adoption of PCIe 6.0 at 64 GT/s
- PCIe slots can accommodate various types of expansion cards, expanding capacity to move data between multiple hosts and multiple devices in parallel
- Low-power modes like the L0p allows chipset vendors and application providers to reduce power requirements

BENEFITS OF PCIE

- PCIe devices are easily discovered, programmed and managed using standard SW for instant deployment for AI
- Adoption of PCIe enables universal interop between hosts and devices
- For large training models, PCIe-based interconnects such as CopprLink enable maximum communication BW between CPUs and GPUs/AI accelerators
- Security features like CMA, TDISP and IDE

FUTURE

- Next generation of PCIe 7.0 will provide increased bandwidth at higher speeds of 128GT/s
- PCIe CopprLink Internal Specification supports up to 1m within a platform node between CPUs and GPUs/AI, which will support large training models