OIF – February 2025

Alan Weckel
alan@650group.com
+1 650-600-7139

About 650 Group

Company Overview

We started 650 Group to provide our subscribers and customers with a unique perspective on the industries we research. Our team has decades of experience researching our focus industries and in roles at companies similar to the ones we cover. We research the data center, communications and Information Technology markets.

- Launched in 2017 by Chris DePuy and Alan Weckel
- Founded in Silicon Valley
- Trusted source of research for system vendors, component manufacturers, service providers, sell-side, buy-side, and standards bodies

Company Overview

What We Track

Cloud
- IaaS, SaaS, Colo, Search and Social
- CAPEX, Equipment Trends

Telecom Equipment
- Broadband Access, Telecom Core
- NFV, Mobile RAN, SP Routing, Optical Transport

Data Center Equipment
- Data Center Switching, Servers, Storage
- Merchant Silicon, DCI, Security

Enterprise Networking Equipment
- Switching, WLAN, Security
- Enhanced NAC, Unified Access, SD-WAN

Semiconductors and FABs
- Campus, DC, Cloud, IoT Semis
- ASIC trends in Switching, WLAN

Who Uses Our Research

Buy Side / Sell Side

Cloud Providers

Telco SPs

System Vendors

Semiconductor Suppliers

Component Manufacturers

FABs

Setting the Stage for AI/ML

DC Semiconductors:
DC Semiconductor Revenue – Total

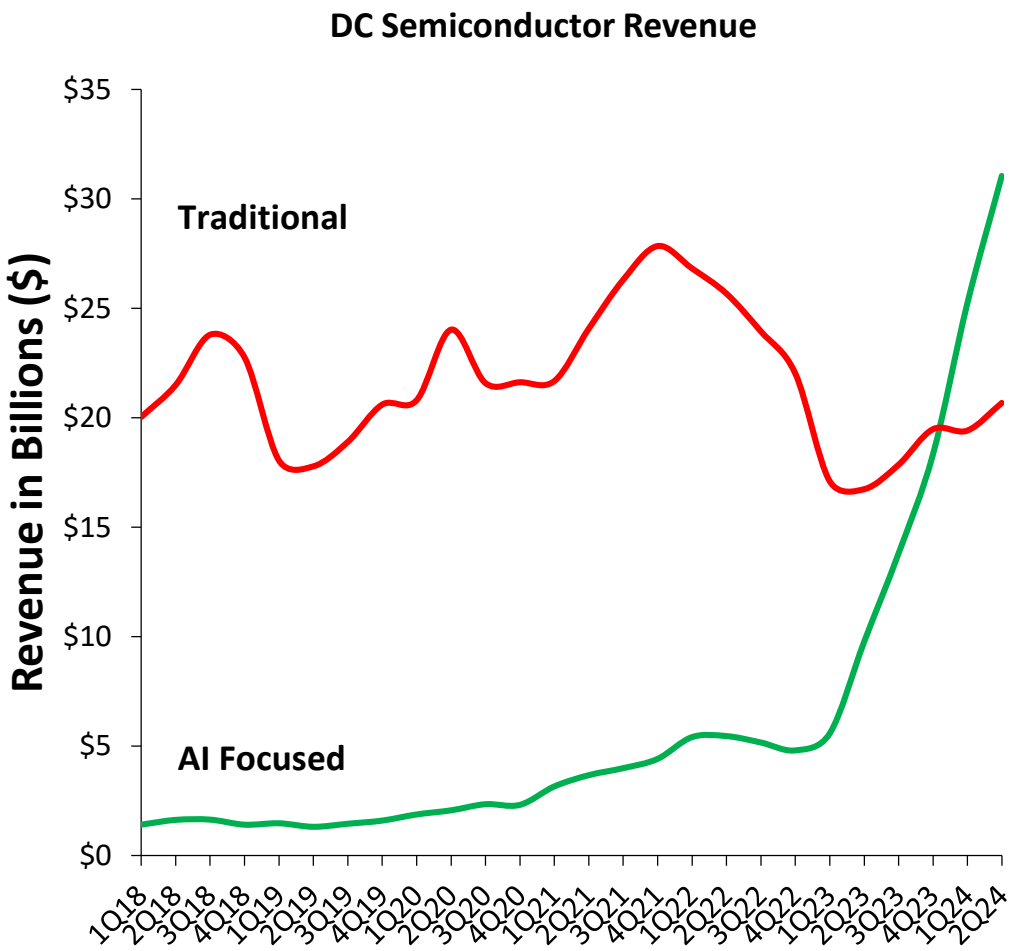

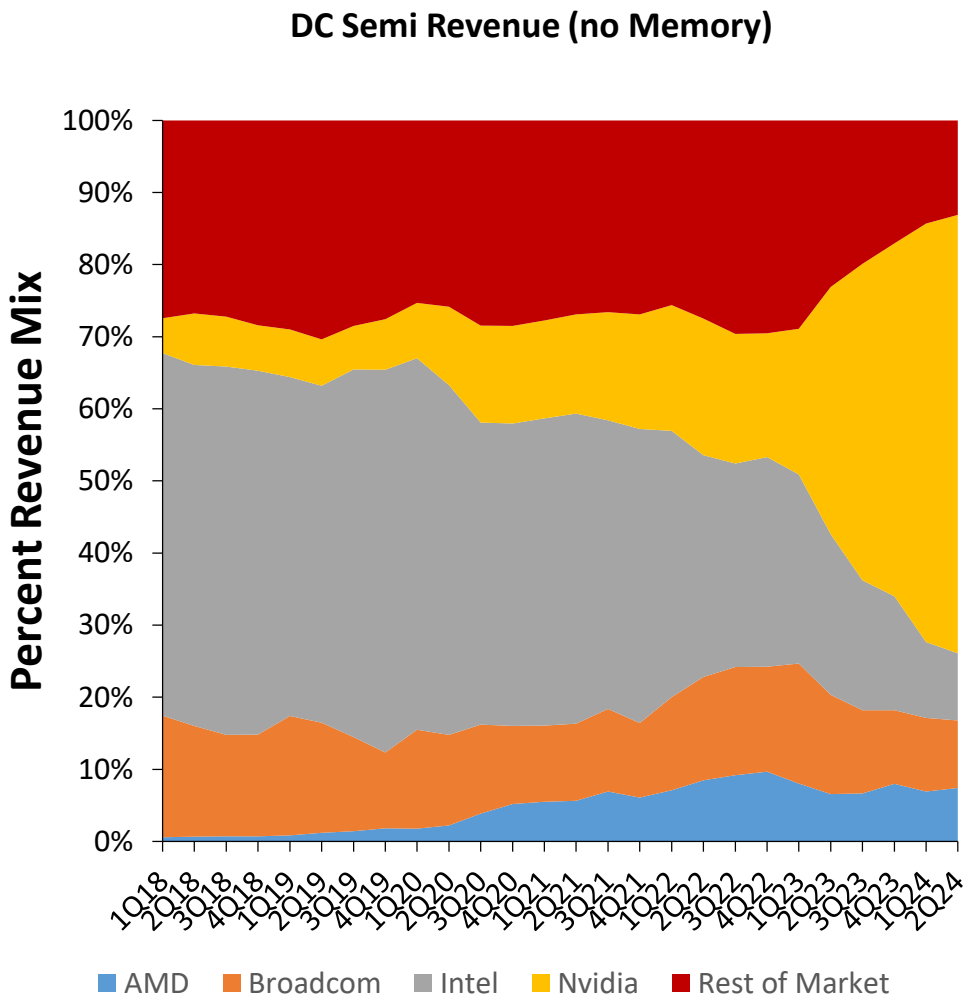

AI Waves

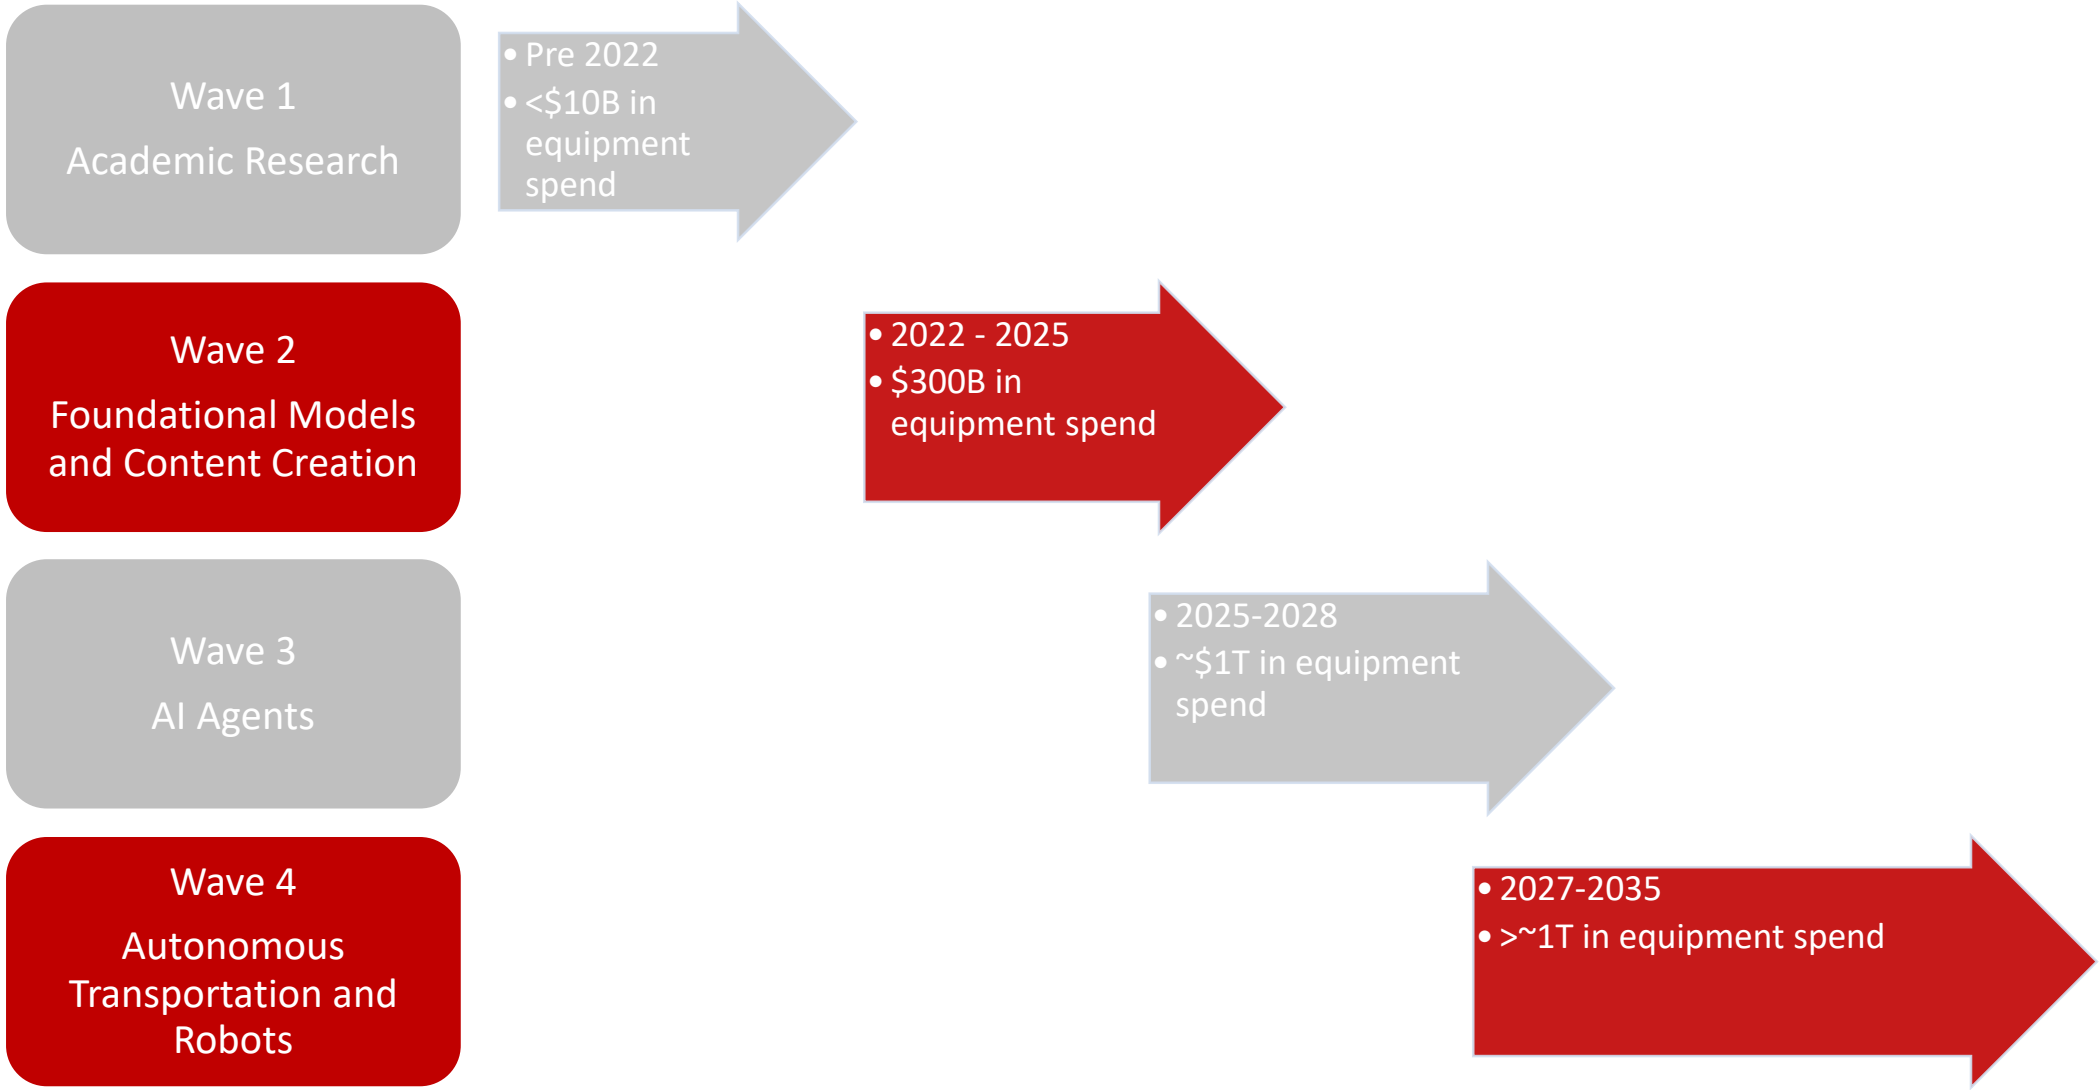

DC Semiconductors: Likely 2024 GPU/ASIC Revenue Split

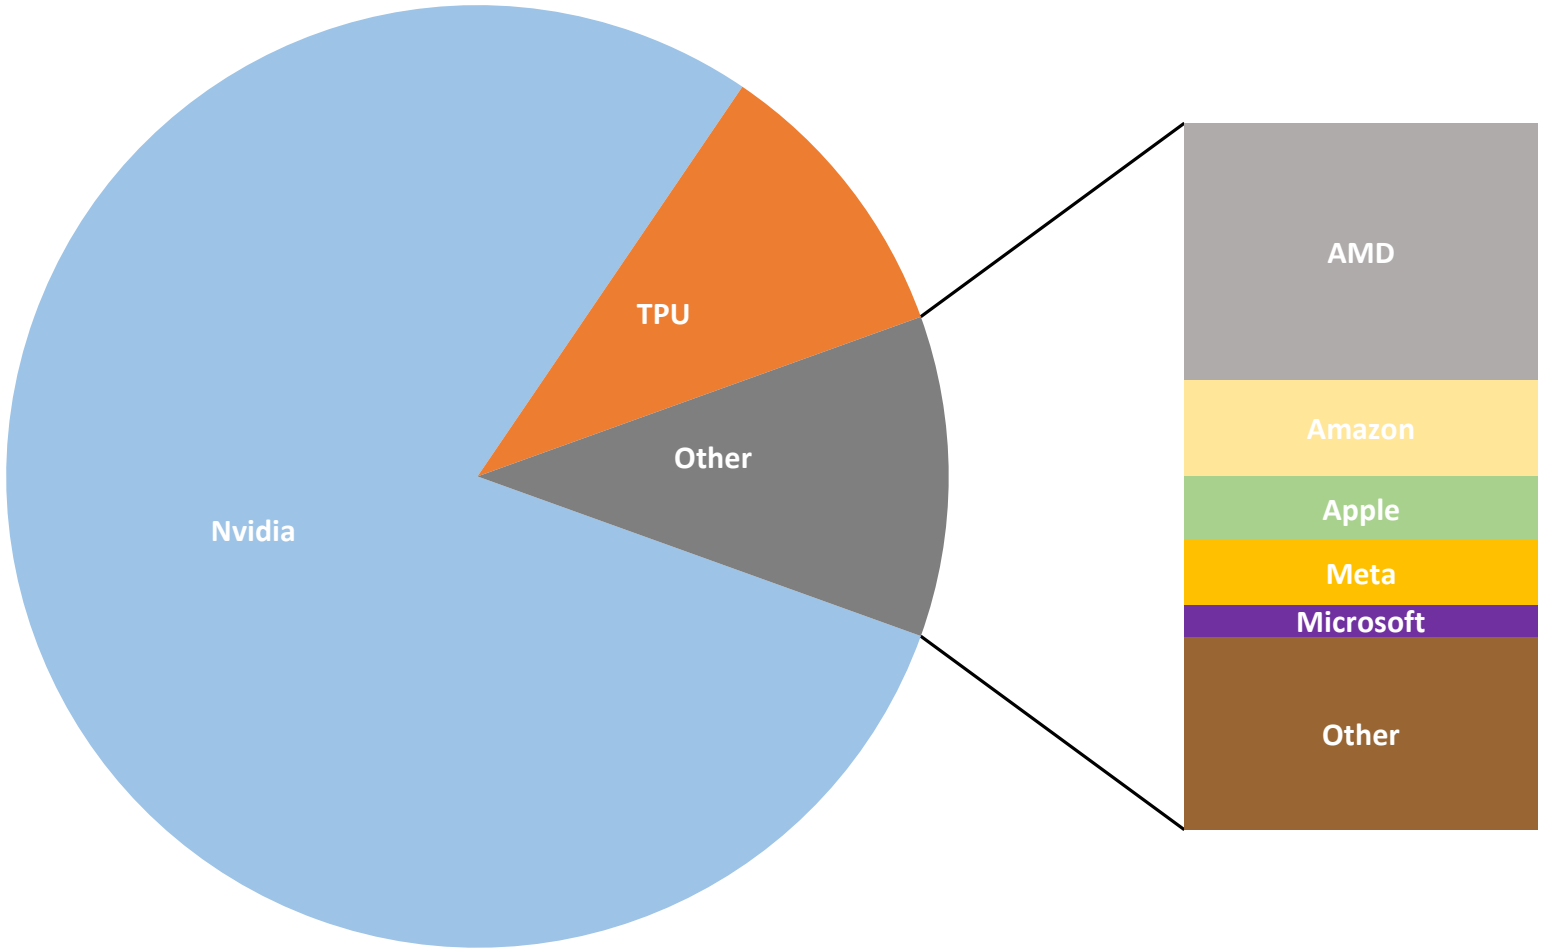

*Based on company guidance and 650 Group estimates; Many ASIC projects are 4Q24 dependent and may shift into 2025 or be recognized for revenue in 2025

DC Related Power:
Data Center Power (North America) (Preliminary)

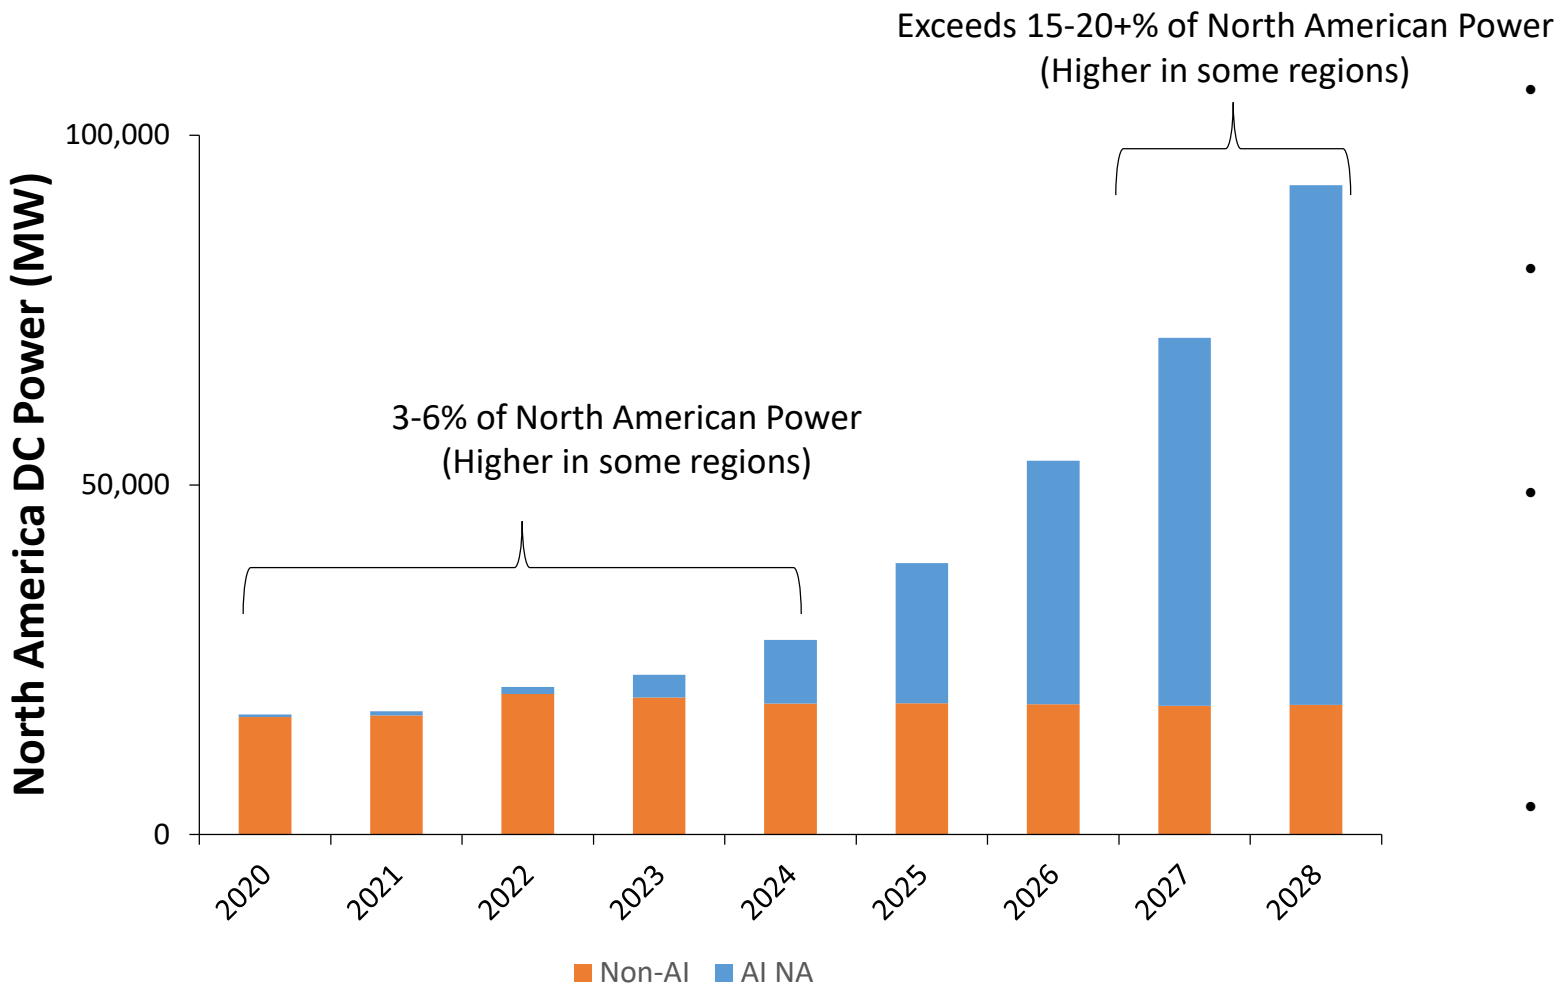

- Tier-2 Training and Inference will need purpose-built lower power ASICs (lowers blue bar)
  - The right ASIC for the right workload
- May move workloads to other continents where power is more readily available (lowers blue bar)
  - Similar to how most of Japan's DCs sit in the Pacific Northwest
- X86 Server refresh can push down Non-AI (lowers orange bar)
  - 1M older servers can be replaced with ~600K to get the same level of compute
  - Only a one-time savings, but can cover almost all of one years shortfall in new power generation
- Liquid Cooling reduces power consumption (lowers blue bar)
  - Cold Plate cooling is the technology for current data centers
  - Emersion cooling is not ideal in many facilities

DC Related Cooling: Market Transition with AI

Traditional x86 Server Rack (5-15 kW per rack)

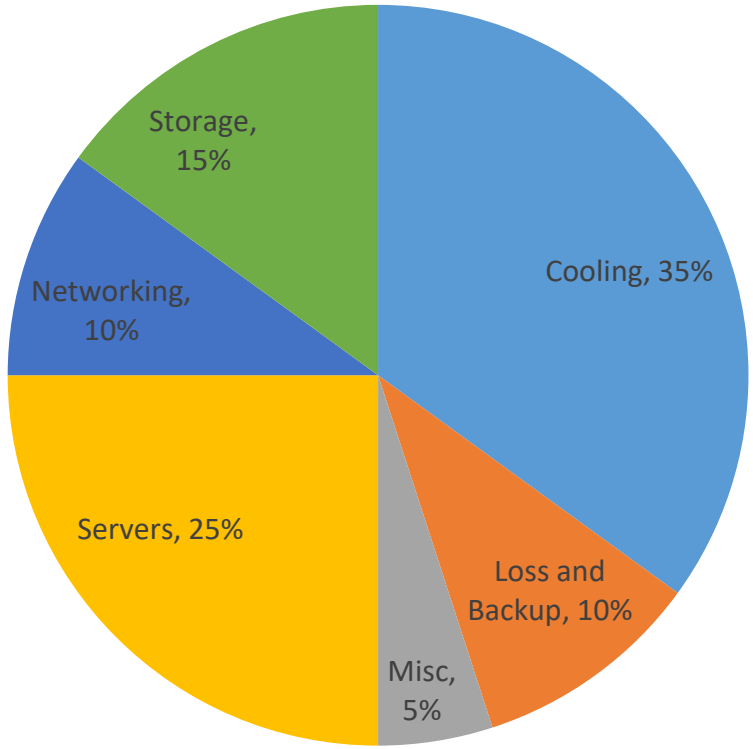

AI/ML Server Rack (50-100 kW per rack)

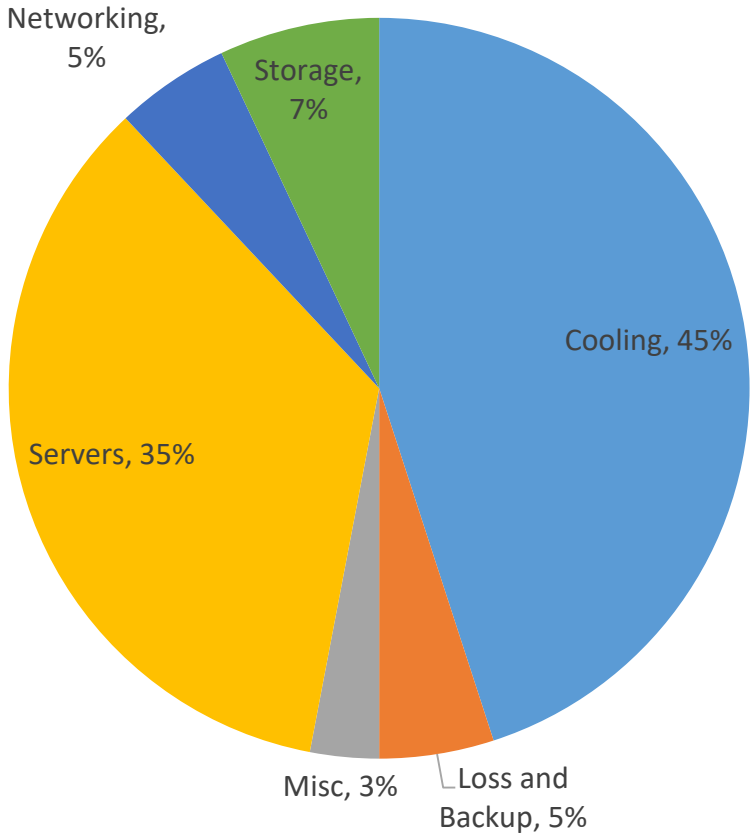

Cloud Report: Capital Intensity

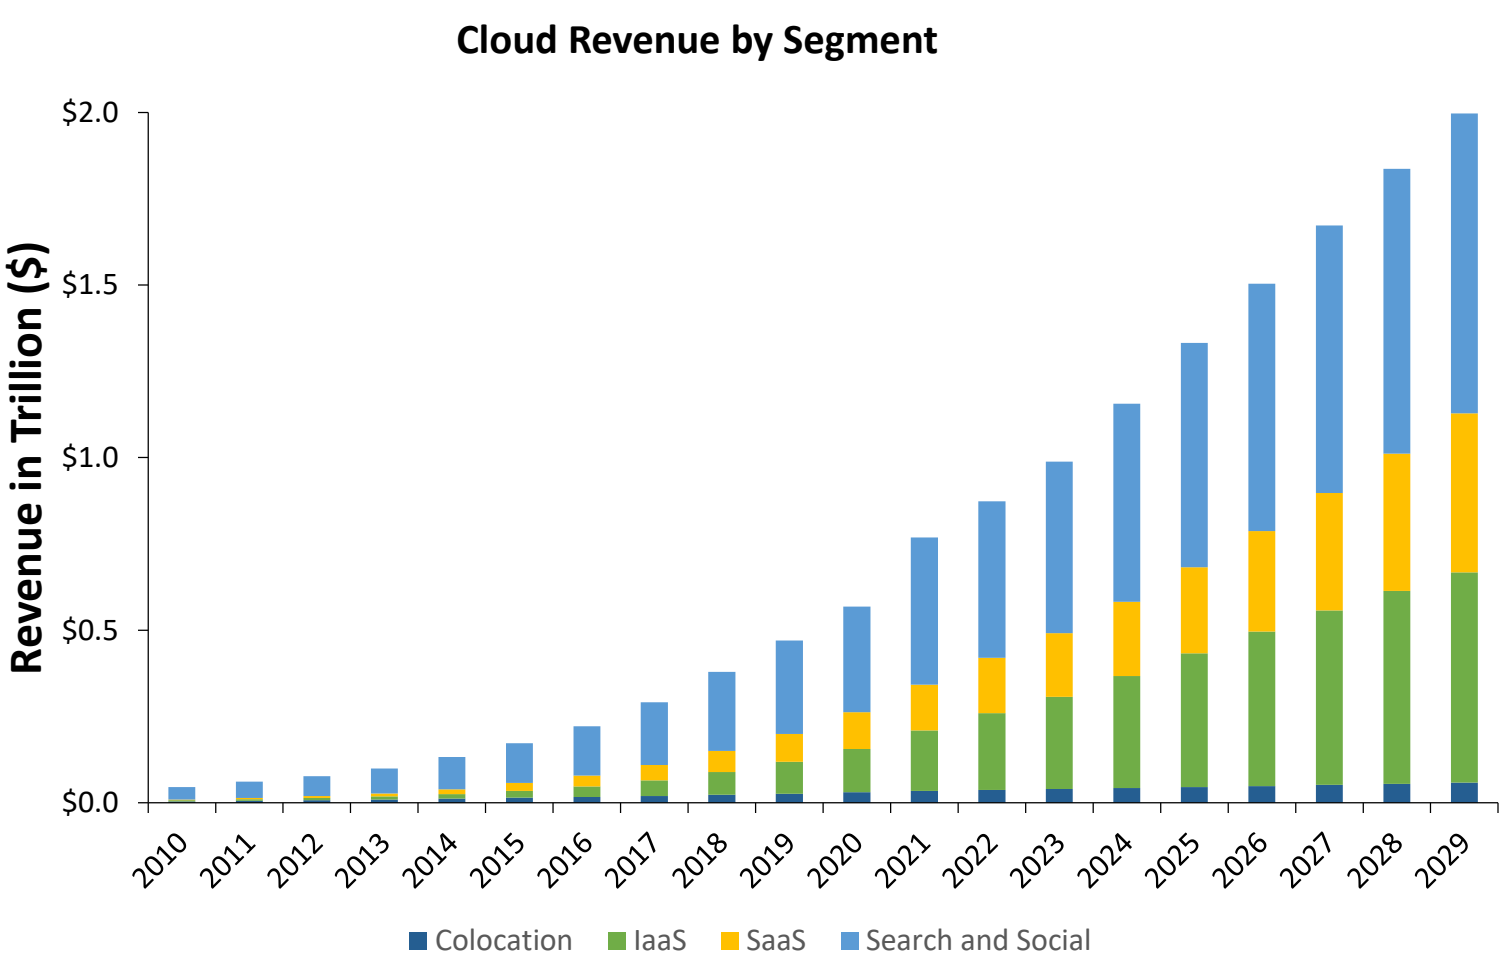

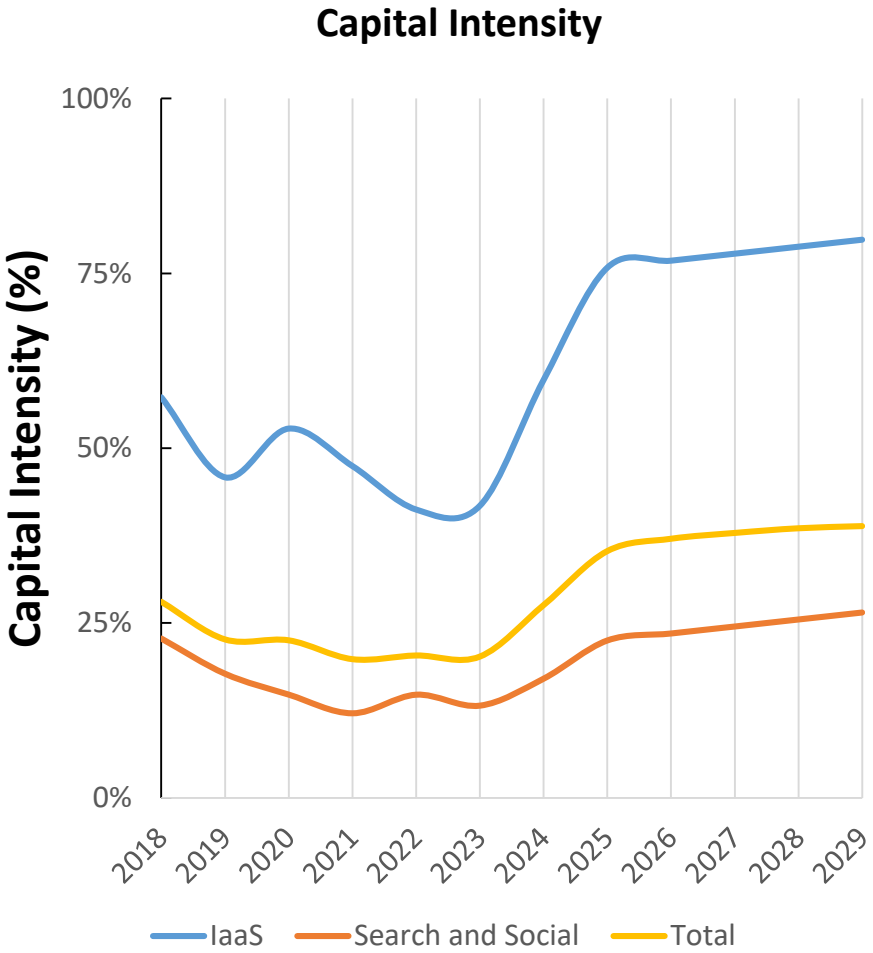

*Using Cloud CAPEX to exclude things like Amazon warehouses, Apple manufacturing, and Google's Waymo

Total Data Center (IT) Equipment: Total Market Installed Base

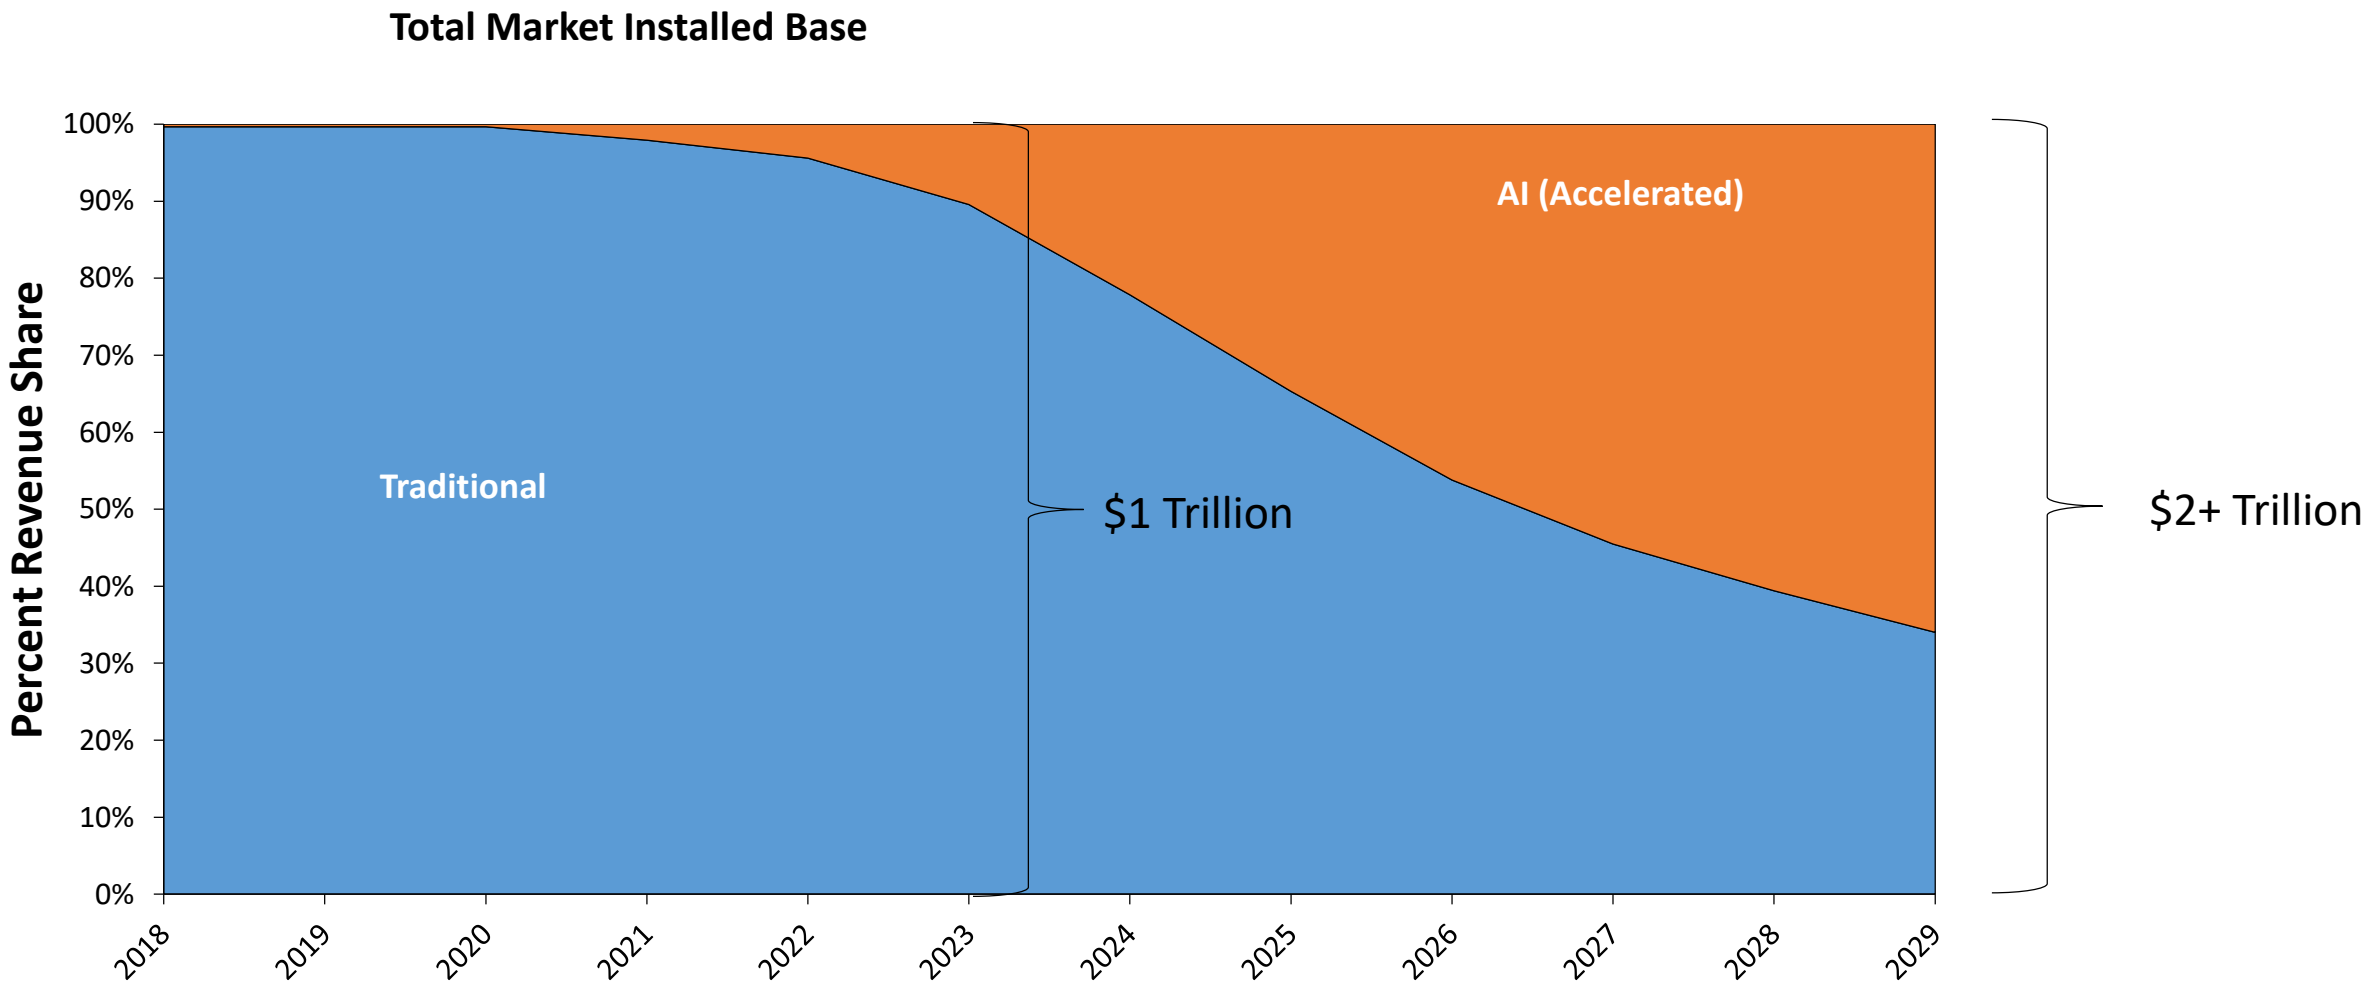

*Including Compute, Storage, and Networking. Excluding Security, ADCs, and Cabling

Server and Smart NICs: Server Shipments (preliminary)

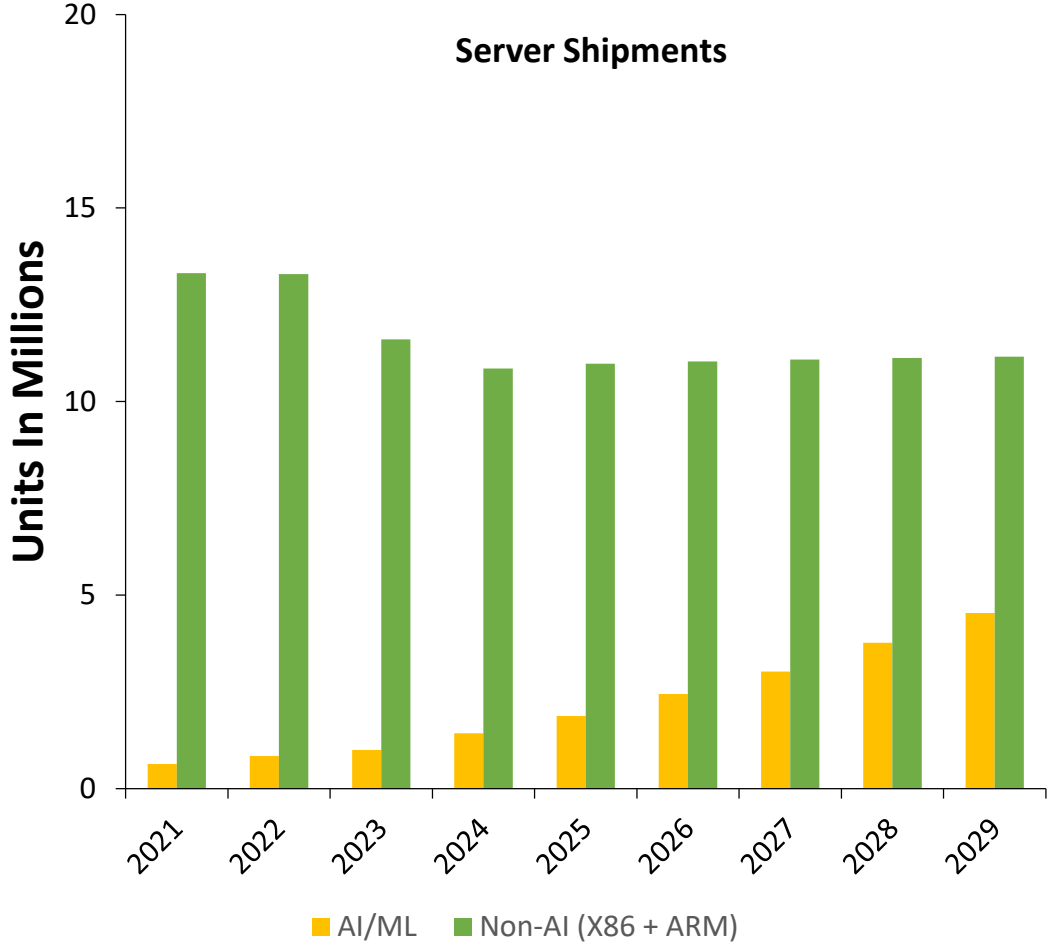

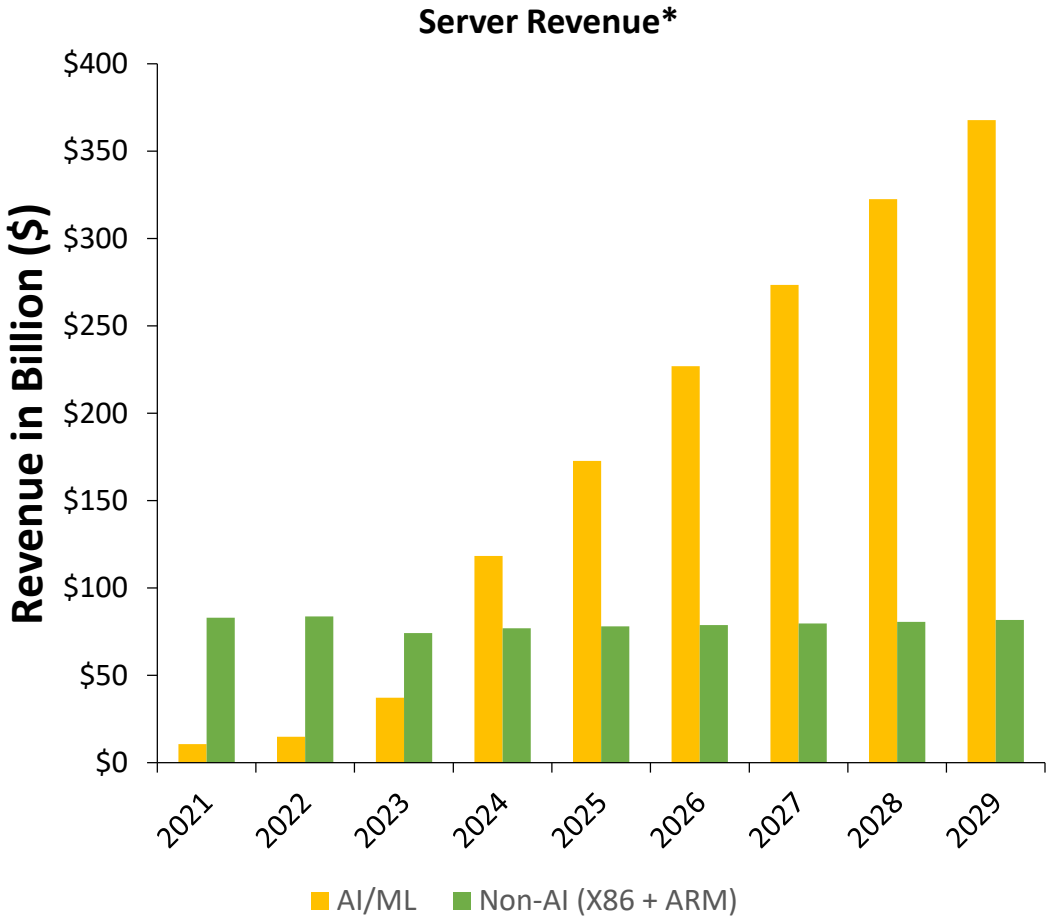

*Includes direct purchase and consignment GPU card revenue

Server and Smart NICs:
Hyperscalers reach new scale, drive customization

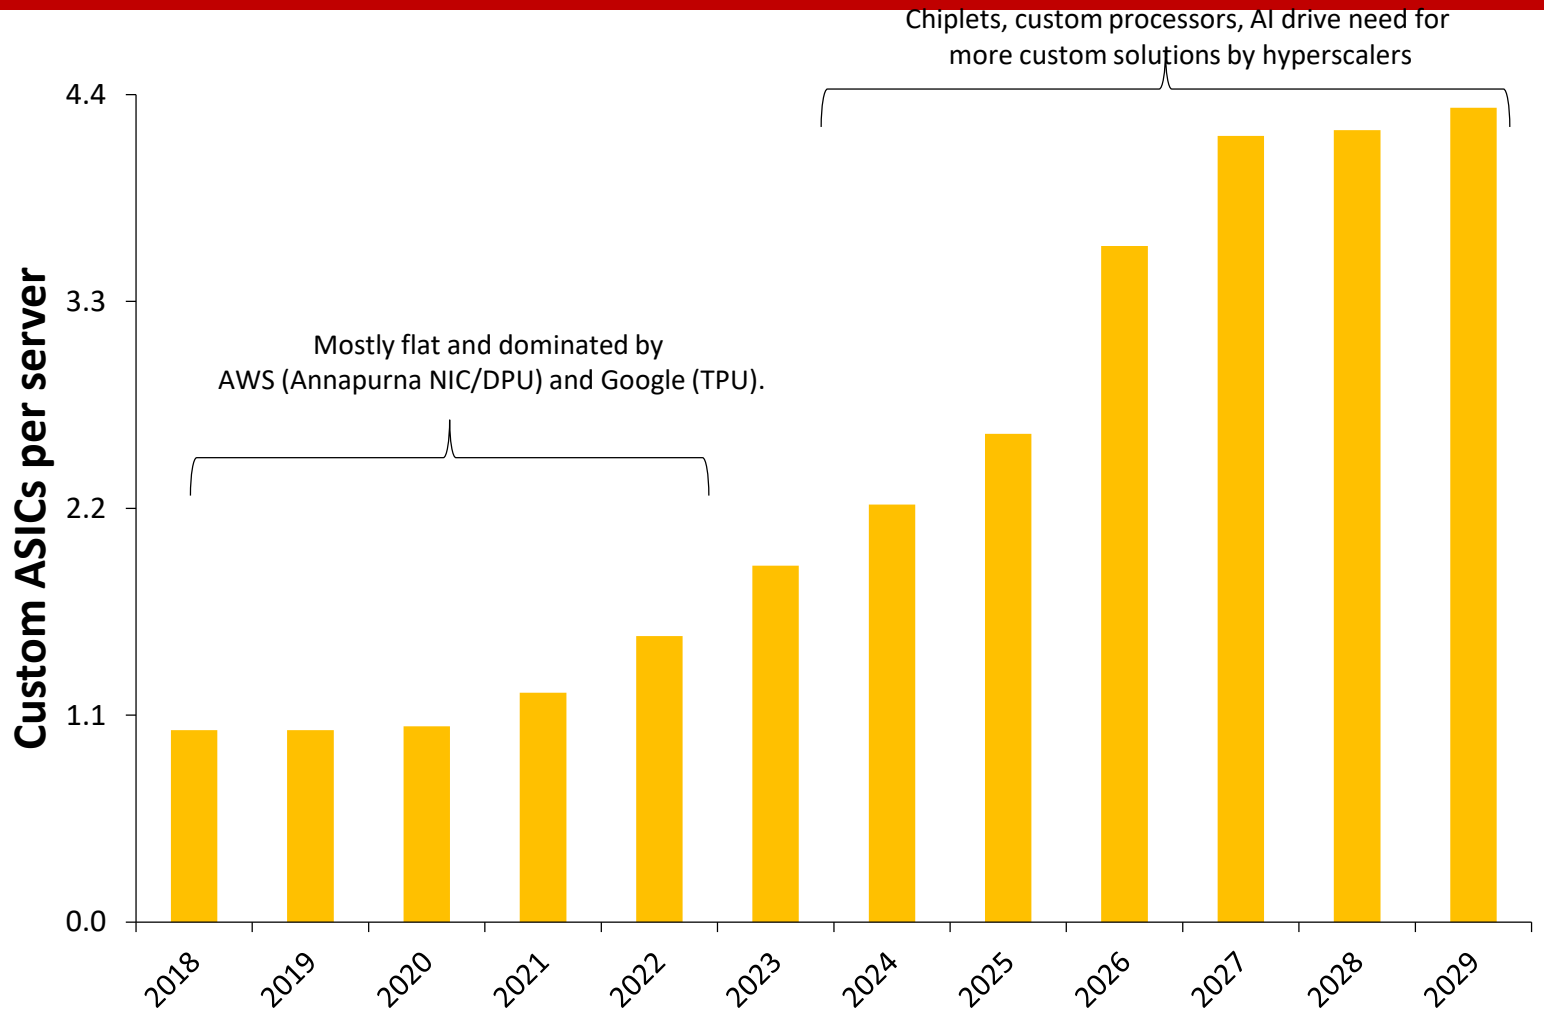

- Custom ARM and RISC V development
- Purpose built accelerators for verticals or workloads
  - 5G/6G
  - AI/ML
  - Enterprise/Consumer
- Fully custom and semi-custom approaches
  - Chiplets allow more customization
  - 100% in-house designs will be rare
  - IP licensing opportunity
- New companies forming
- Additional semiconductor consolidation

*A multi-chip solution, or multi-chiplet solution is counted a 1. This chart only reflects products that go in a server.

Cloud Report:
DC Equipment CAPEX

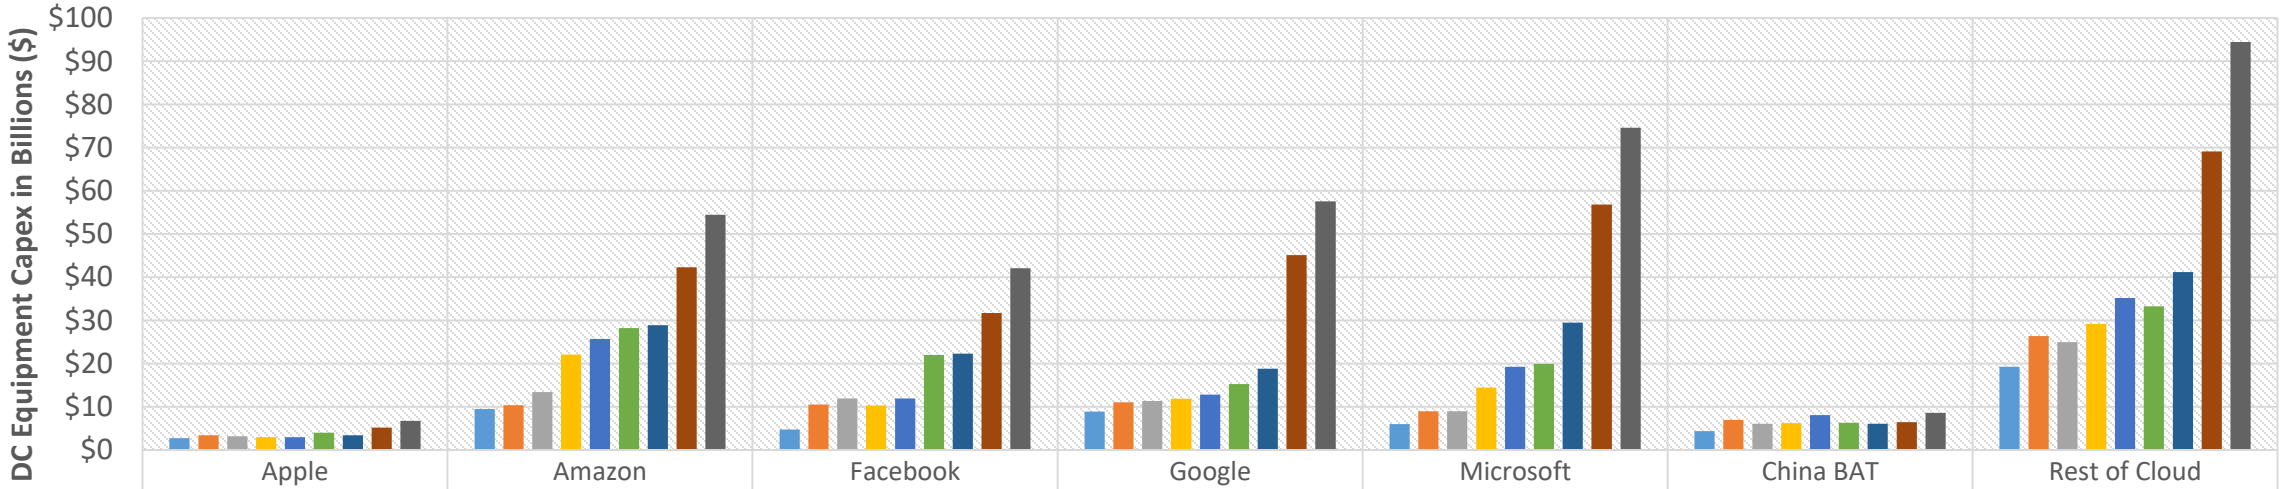

	Apple	Amazon	Facebook	Google	Microsoft	China BAT	Rest of Cloud
2017	$2,702	$9,476	$4,692	$8,860	$5,952	$4,362	$19,251
2018	$3,360	$10,310	$10,493	$10,980	$8,947	$6,947	$26,321
2019	$3,120	$13,385	$11,875	$11,296	$8,902	$6,014	$24,933
2020	$2,948	$22,084	$10,285	$11,795	$14,436	$6,203	$29,186
2021	$2,952	$25,707	$11,880	$12,811	$19,272	$8,049	$35,144
2022	$3,996	$28,219	$21,961	$15,252	$19,906	$6,264	$33,241
2023	$3,375	$28,867	$22,287	$18,828	$29,434	$6,060	$41,169
2024E	$5,152	$42,319	$31,692	$45,103	$56,815	$6,400	$69,133
2025E	$6,684	$54,417	$42,053	$57,589	$74,622	$8,600	$94,501

Meta CAPEX:
DC Equipment CAPEX

Meta's CAPEX Contribution

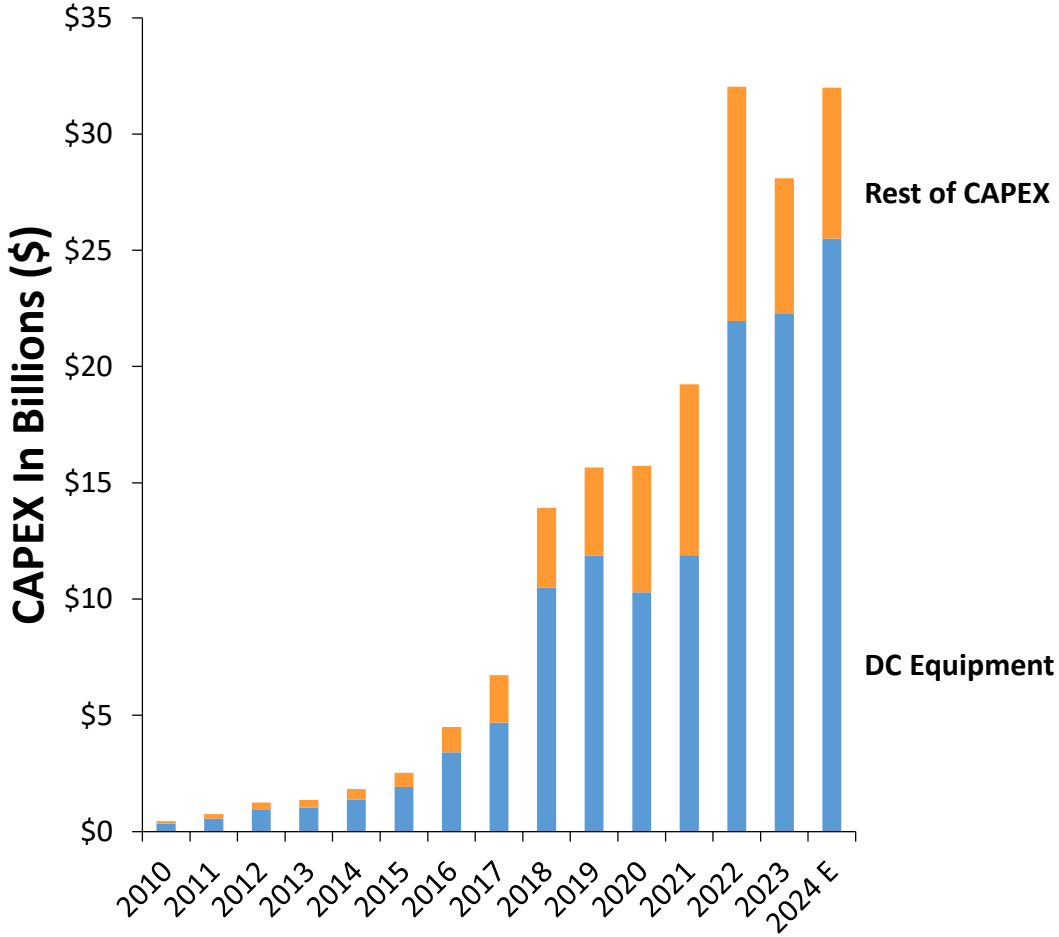

Meta Company CAPEX Guide vs. Reality

	Original Guide	Final	Miss (Peak)
2019	$19 B	$15.6 B	-18%
2020	$17-19 B	$15.7 B	-17%
2021	$21-23 B	$19.2 B	-18%
2022	$29-34 B	$32 B	~correct
2023	$34 B	$28.1 B	-17%
2024 E	$30-37 B		

*Facebook historically guides much higher than final yearly CAPEX number.

Ethernet Switch - Data Center

Ethernet Switch:
Total 400 and 800 Gbps Port Shipments

650 GROUP
MARKET INTELLIGENCE RESEARCH

Total Market

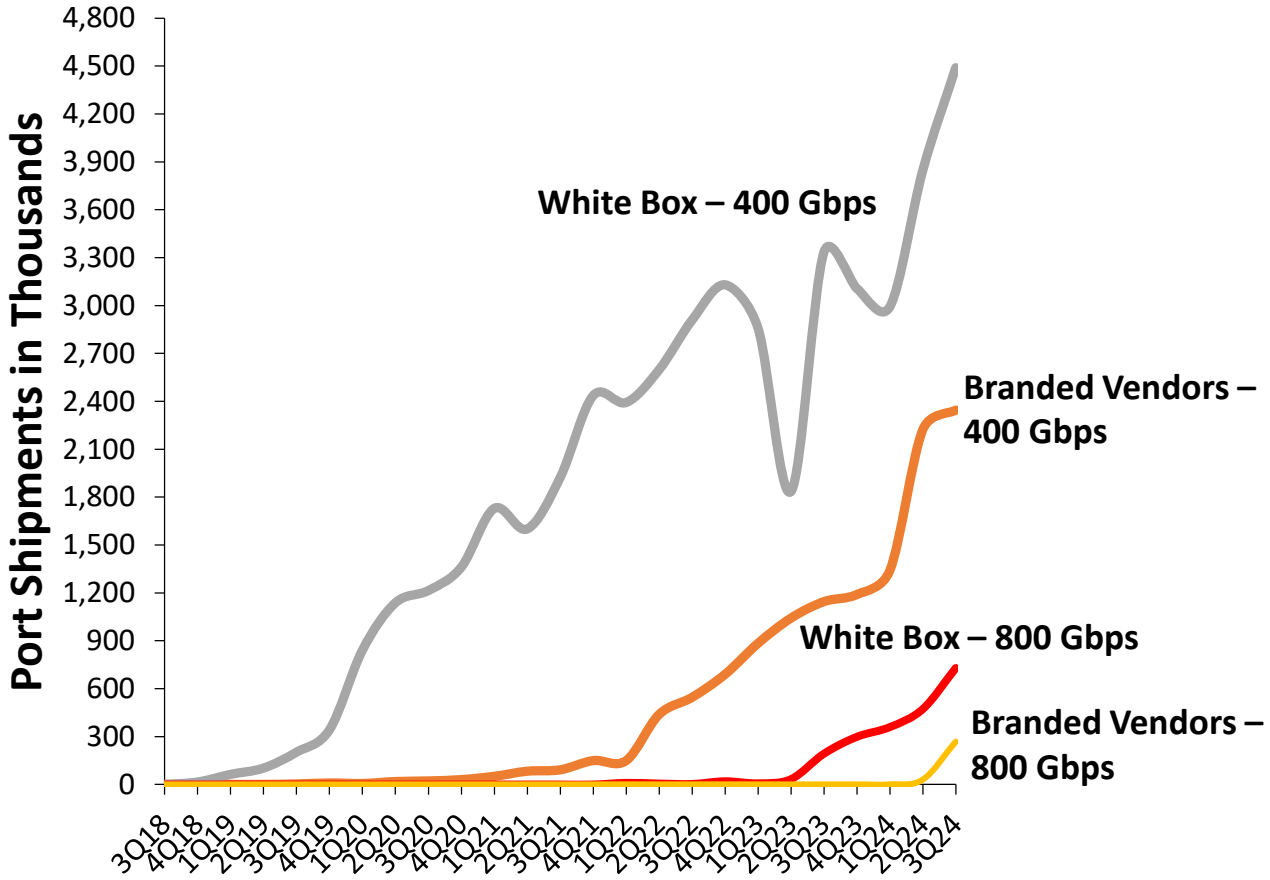

Branded Vendors

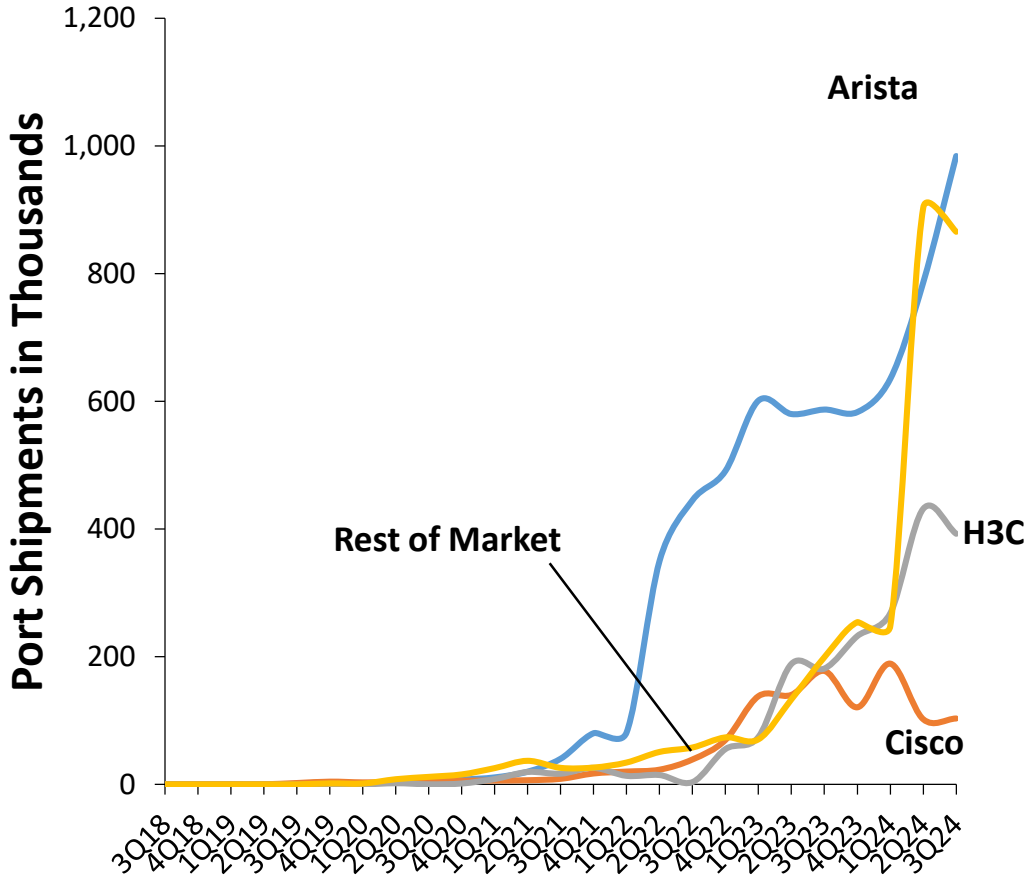

Ethernet Switch – Data Center: Hyperscaler Speed Migration (2023)

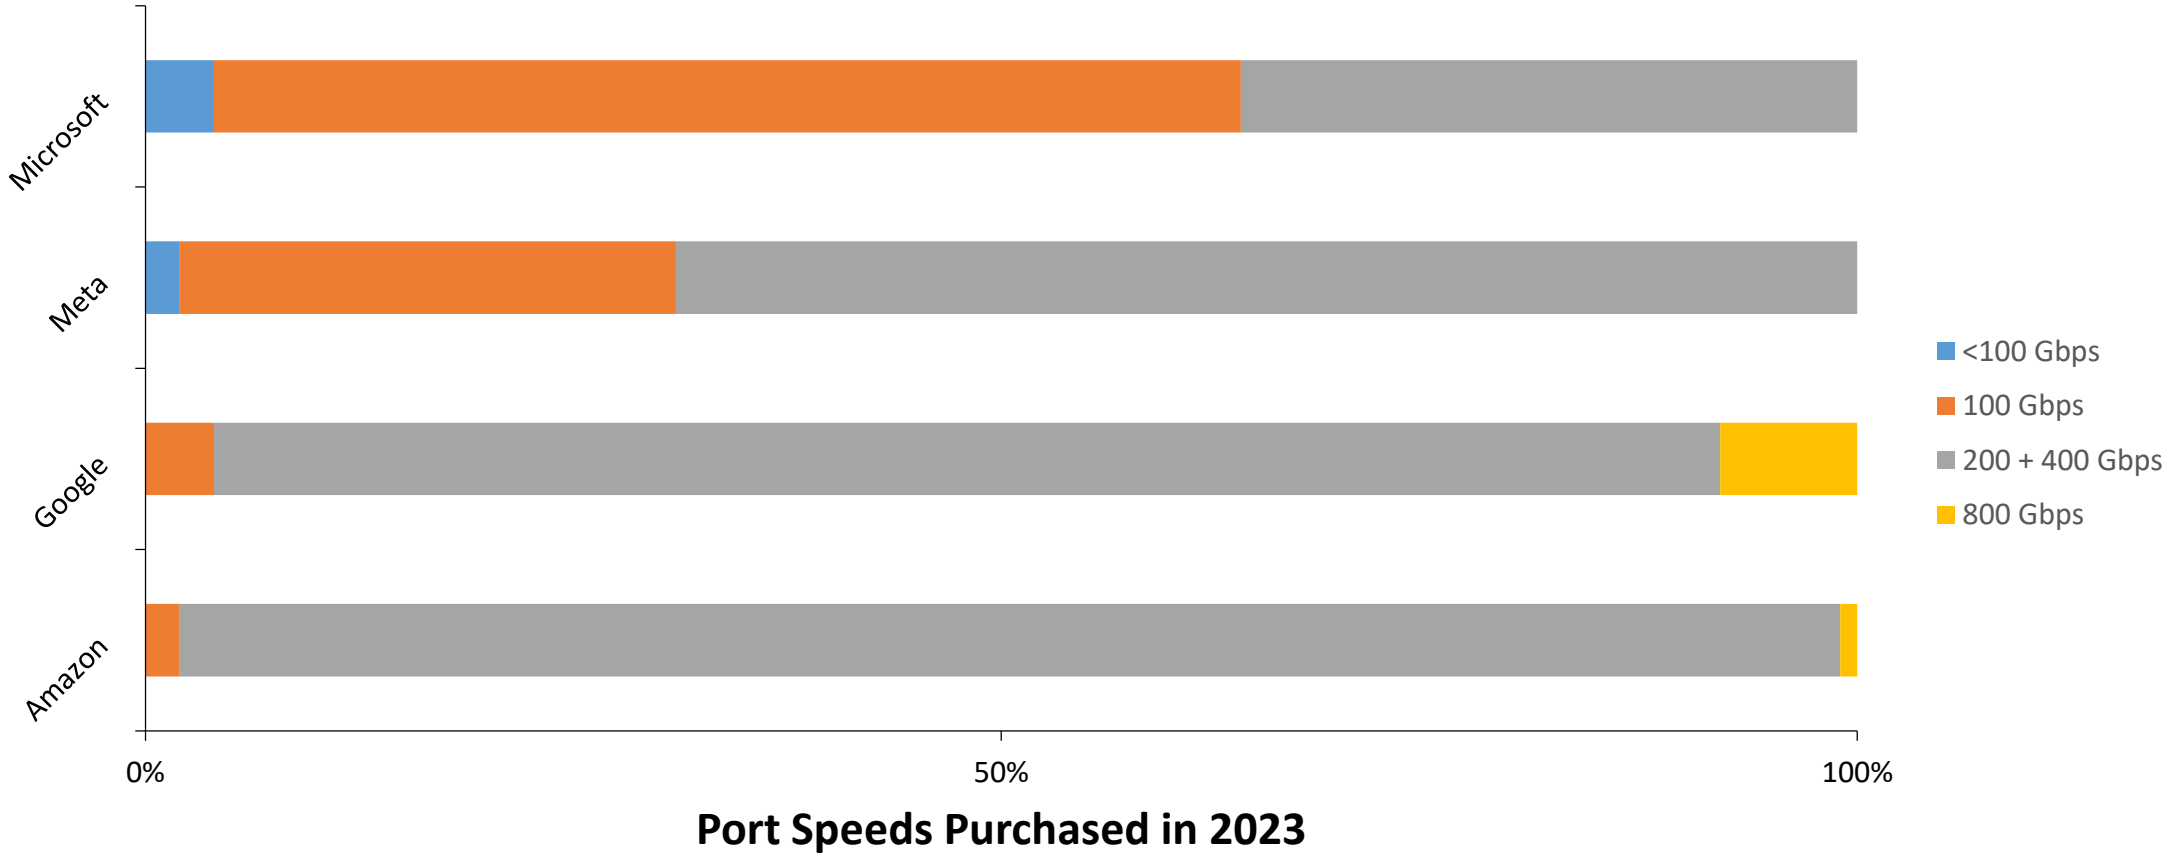

Ethernet Switch – Data Center: Total Market Revenue

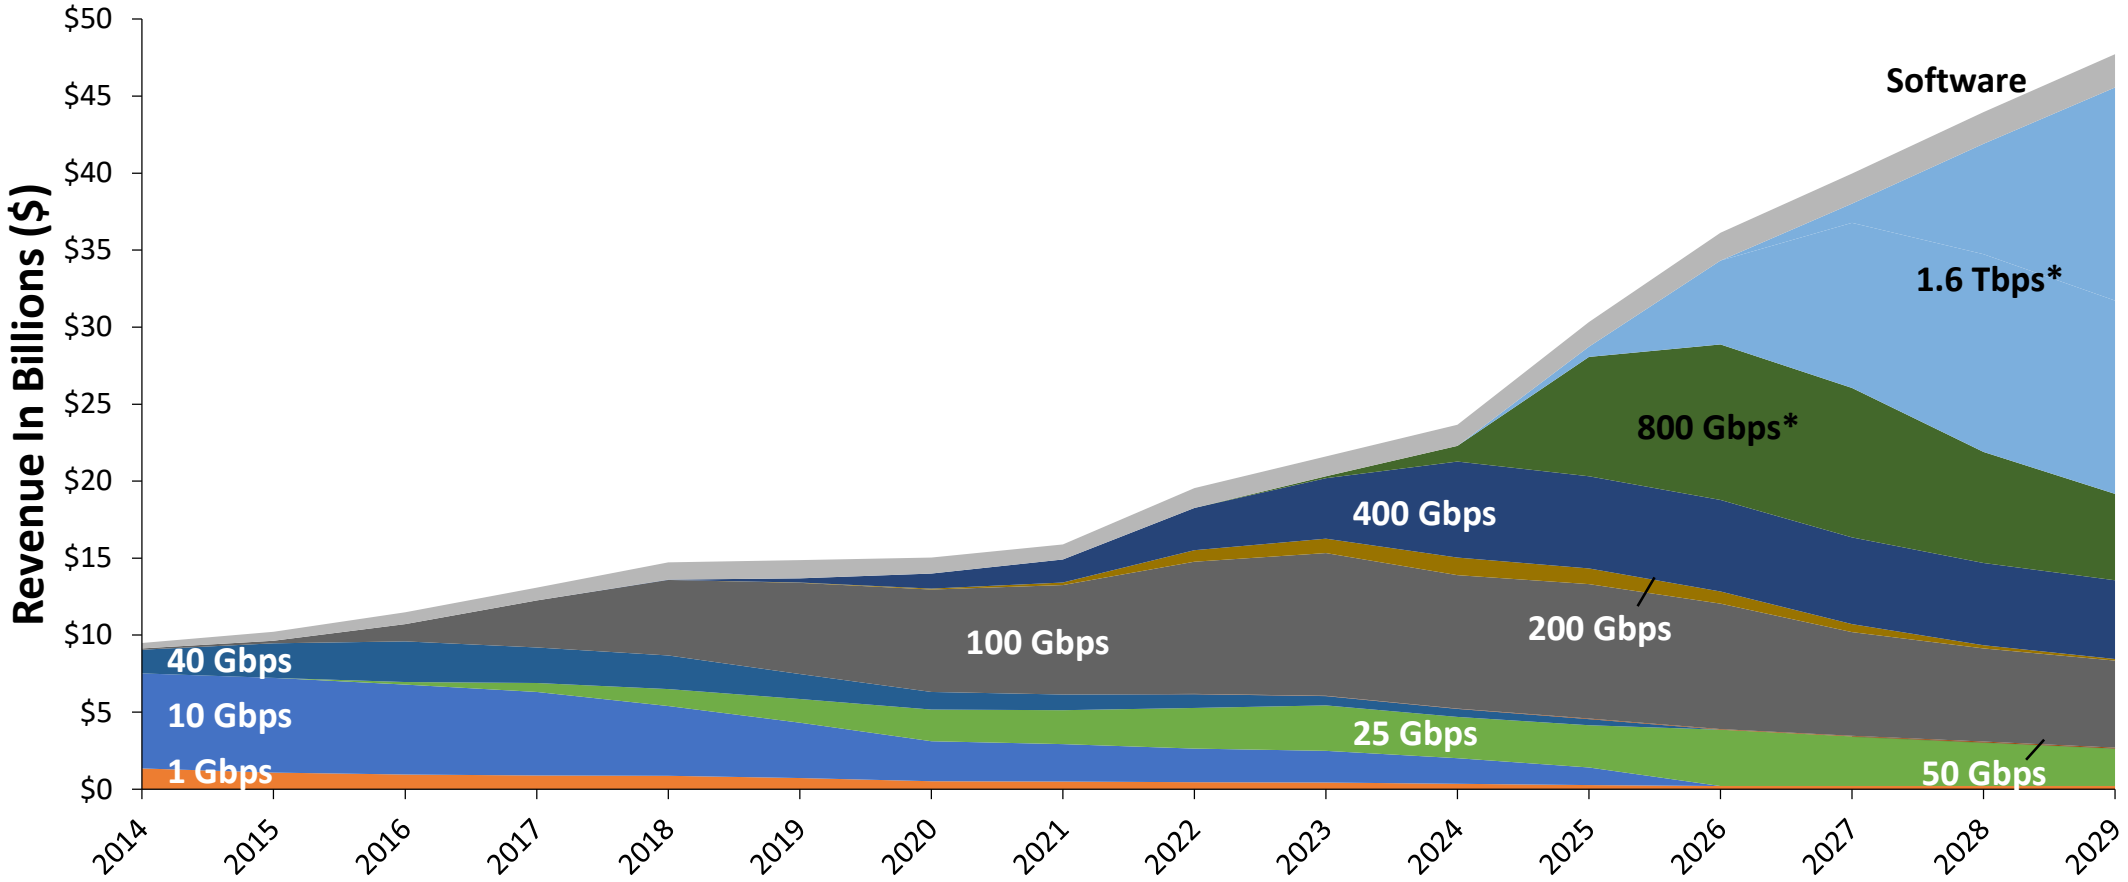

Merchant Silicon – Data Center Switching: Bandwidth Shipping by SERDES speed

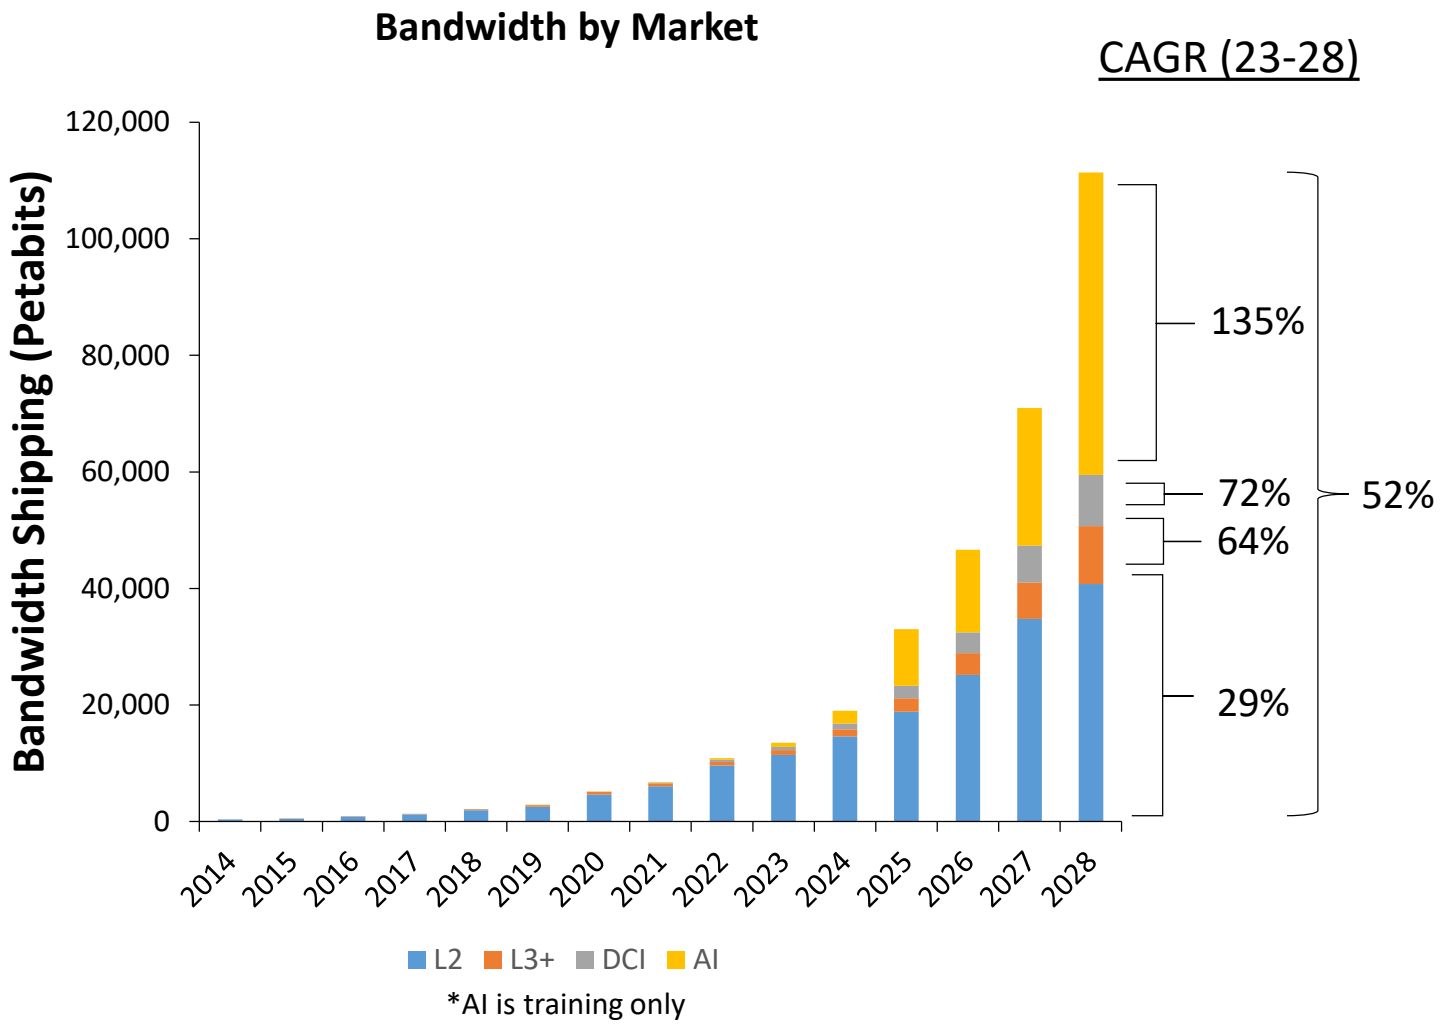

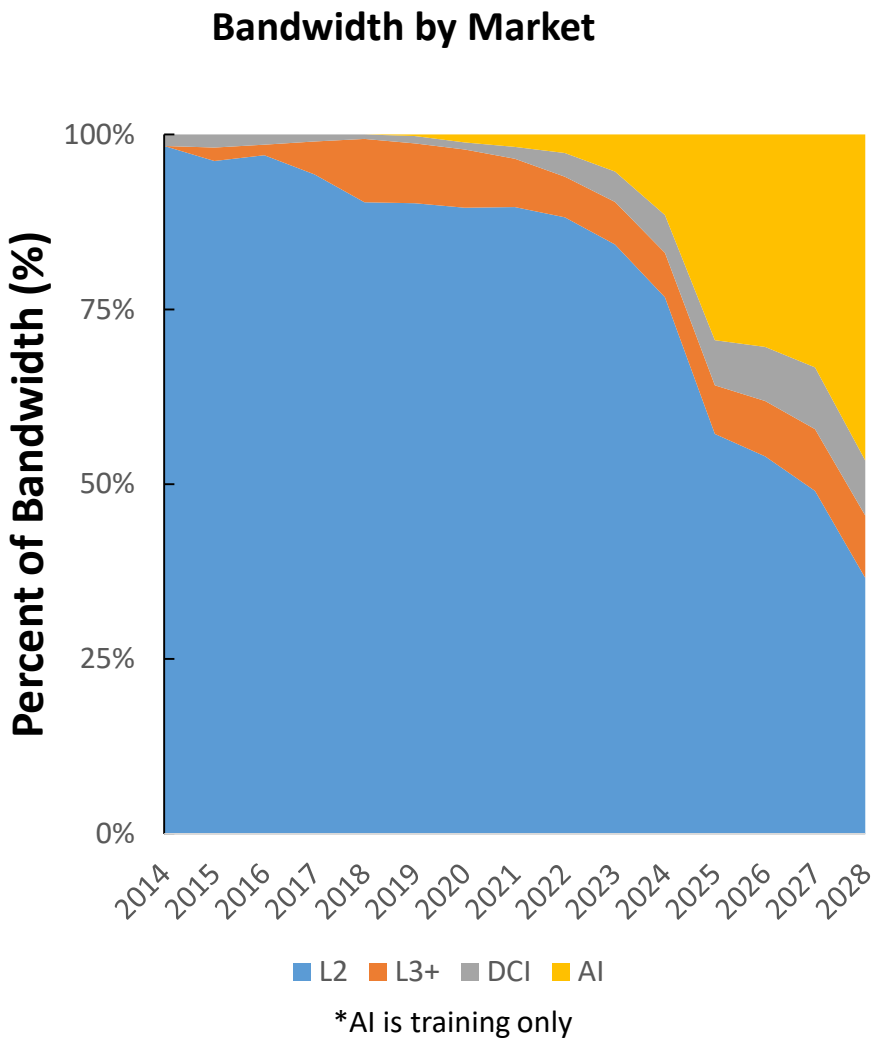

End Markets – Data Center Networking: : AI and HPC Networking Transitions

650 GROUP
MARKET INTELLIGENCE RESEARCH

2022 (x56) -> 2024 (x112) Traditional Cloud Server			
Location	Bandwidth	Technology	Cabling
NIC	25-100G	Ethernet	DAC
Top-of-Rack	100-400G	Ethernet	DAC
Aggregation	100-400G	Ethernet	Fiber

2024 (x56) -> 2024 (x112) Nvidia AI Cloud Server			
Location	Bandwidth	Technology	Cabling
NIC	400-800G	Ethernet	DAC/Active Copper
Top-of-Rack	400-800G	Ethernet	DAC/Active Copper
Aggregation	400-800G	Ethernet	Fiber
NIC	400G-3.2T	InfiniBand	DAC/Active Copper
Top-of-Rack	400G-3.2T	InfiniBand	DAC/Active Copper
Aggregation	400G-3.2T	InfiniBand	Fiber
Embedded	400-900G	NVLink	Copper
Switch	400-900G	NVLink	Rack – Copper Fiber – Multi-rack

2022 (x56) -> 2024 (x112) Ethernet AI Cloud Server			
Location	Bandwidth	Technology	Cabling
NIC	400-800G	Ethernet	DAC/Active Copper
Top-of-Rack	400-800G	Ethernet	DAC/Active Copper
Aggregation	400-800G	Ethernet	Fiber
NIC	400G-3.2T	Ethernet	DAC/Active Copper
Top-of-Rack	400G-3.2T	Ethernet	DAC/Active Copper
Aggregation	400G-3.2T	Ethernet	Fiber

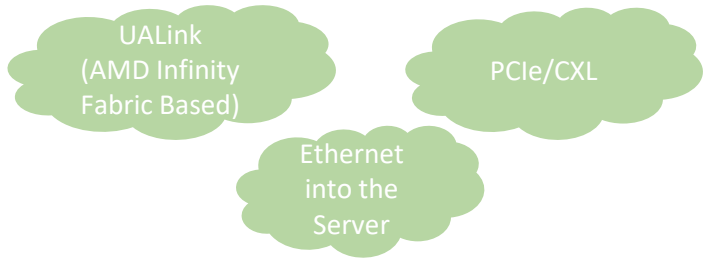

Frontend (External) (1X)

Backend (Internal Scale-Out) (10 X)

Backend (Internal Scale-Up) (100X)

*Optimal fiber required

Ethernet Switch – Data Center: Total Market Revenue

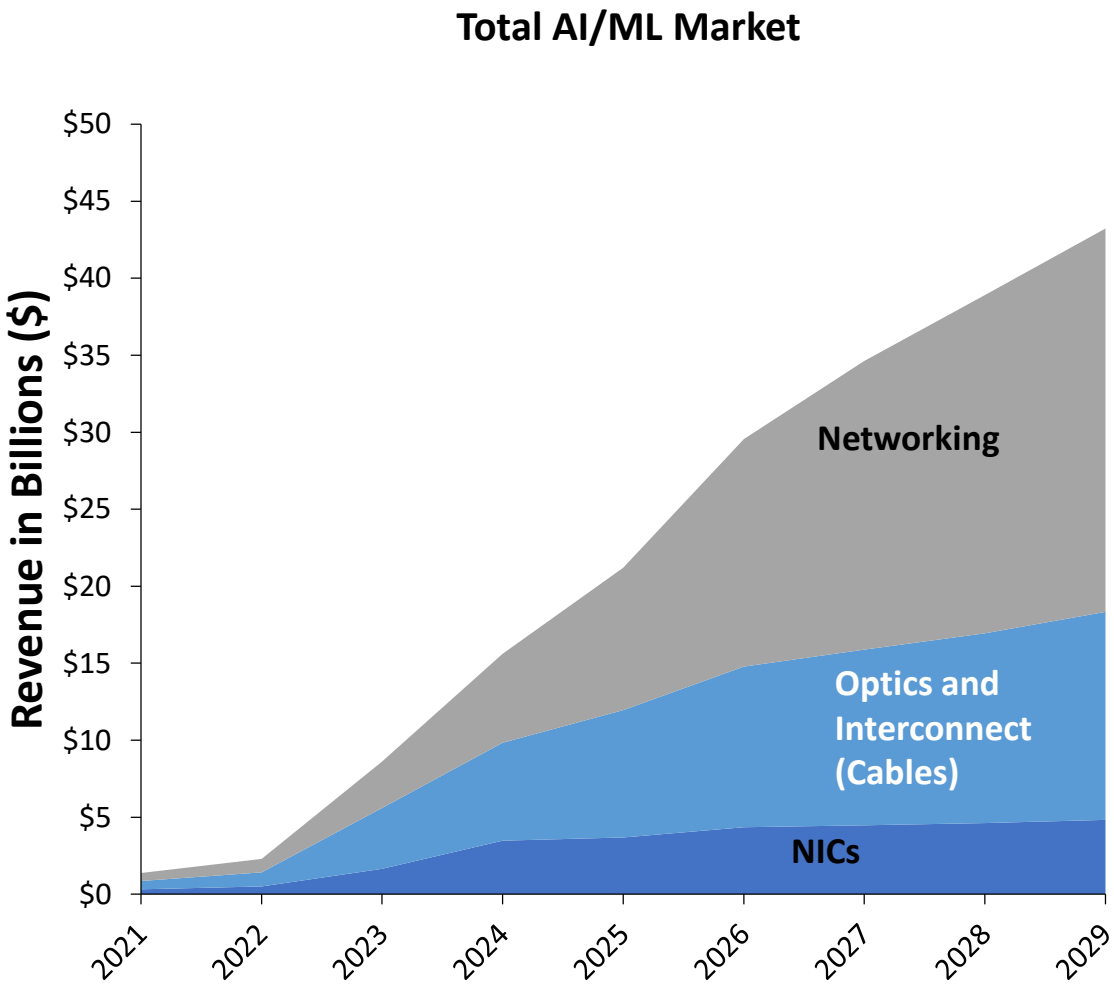

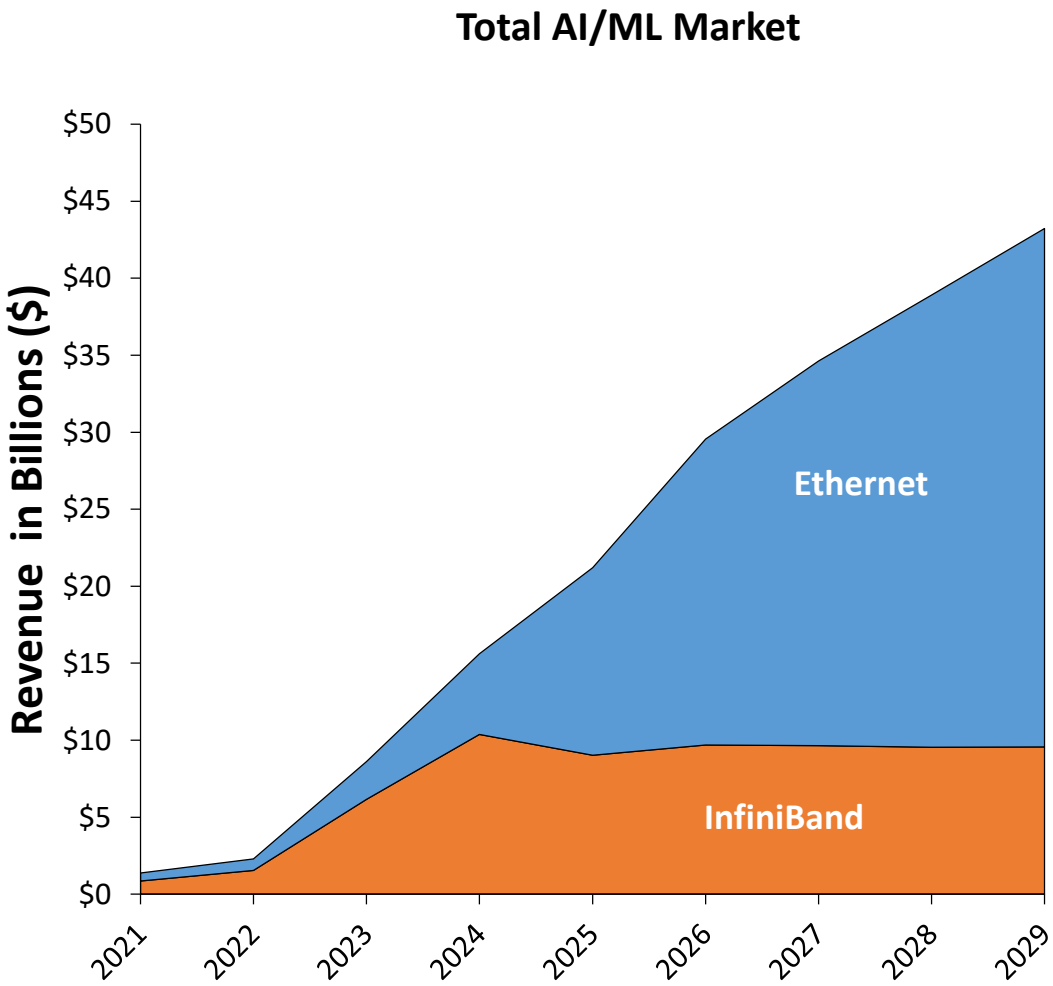

Ethernet Switch – Data Center: Total Market Revenue

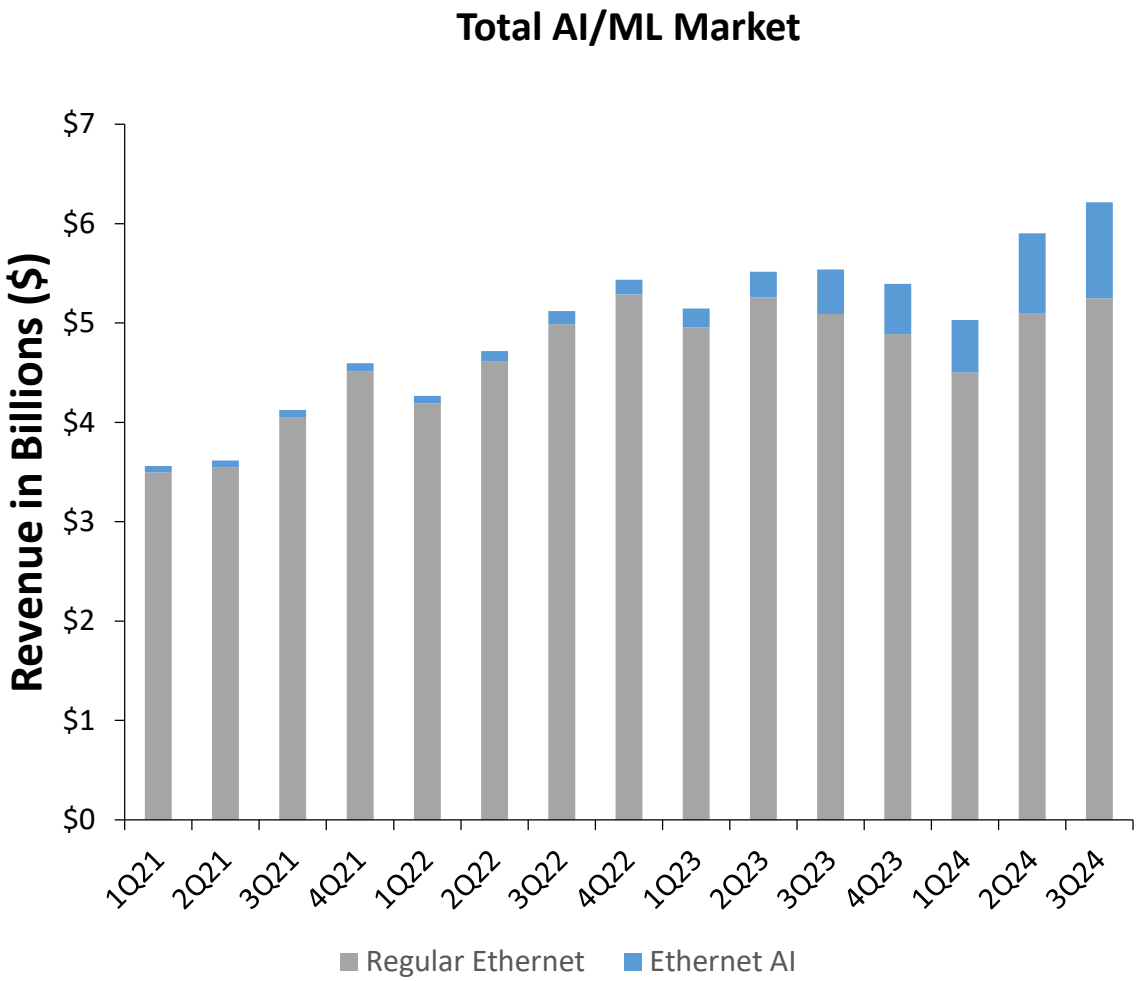

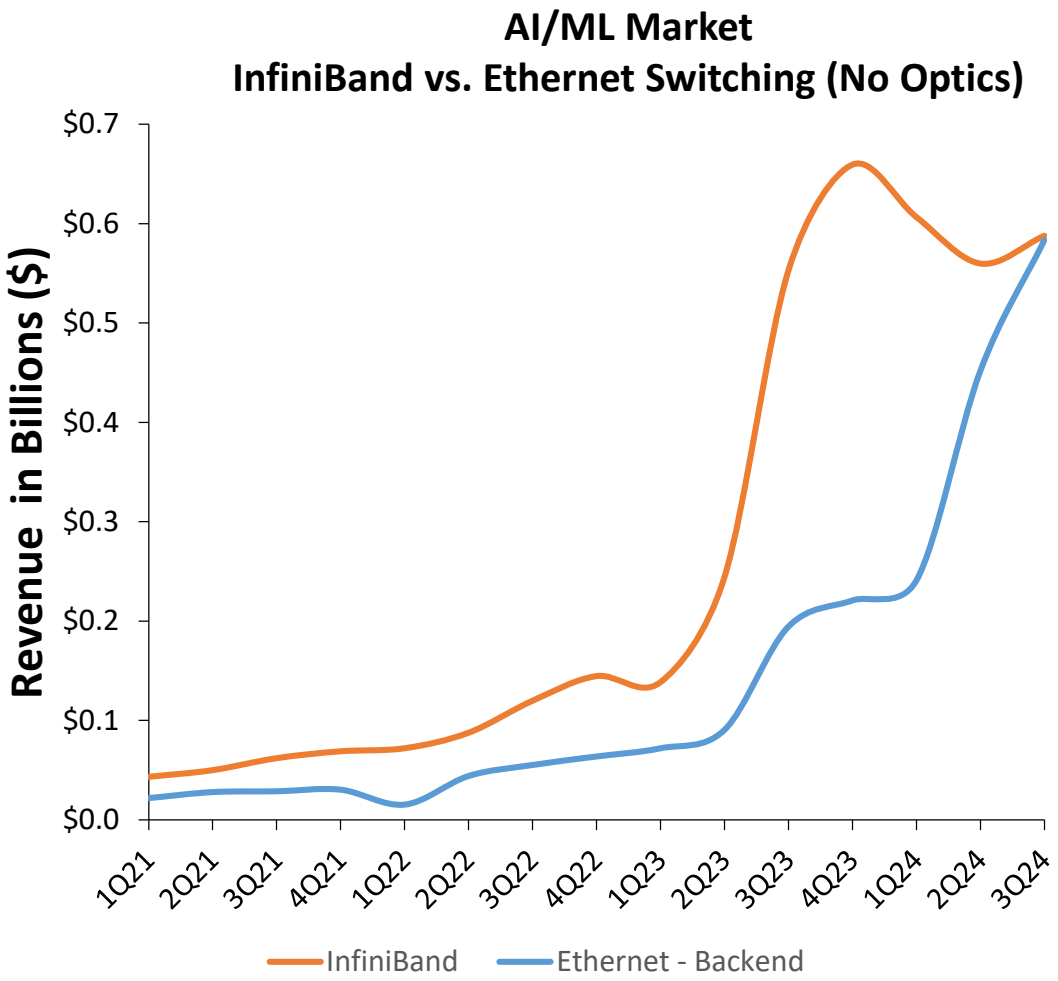

Ethernet Switch – Data Center: AI Networking Market Revenue

2023 - $8.6B

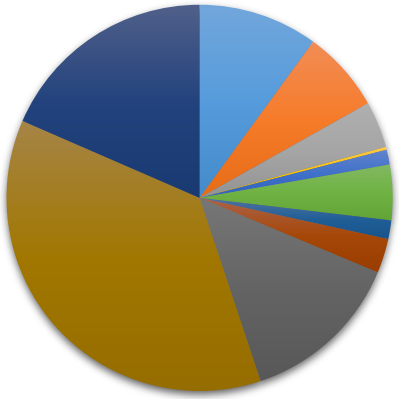

2029 - $43.2B

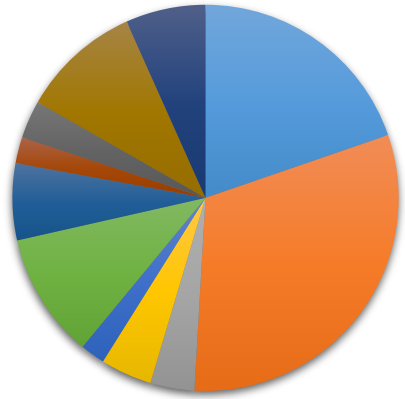

- Ethernet - Frontend
- Ethernet - Backend
- Ethernet NICs - Frontend
- Ethernet NICS - Backend
- Ethernet Copper Connectivty (ACC, AEC, DAC)
- Direct Purchase Optical Tranceivers - Ethernet Backend
- Direct Purchase Optical Tranceivers - Ethernet Frontend for Ethernet
- Direct Purchase Optical Tranceivers - Ethernet Frontend for InfiniBand
- InfiniBand NICs
- InfiniBand Optical Tranceivers and Copper Connectivity
- InfiniBand Switches

Ethernet Switch – Data Center: AI Networking Market Revenue

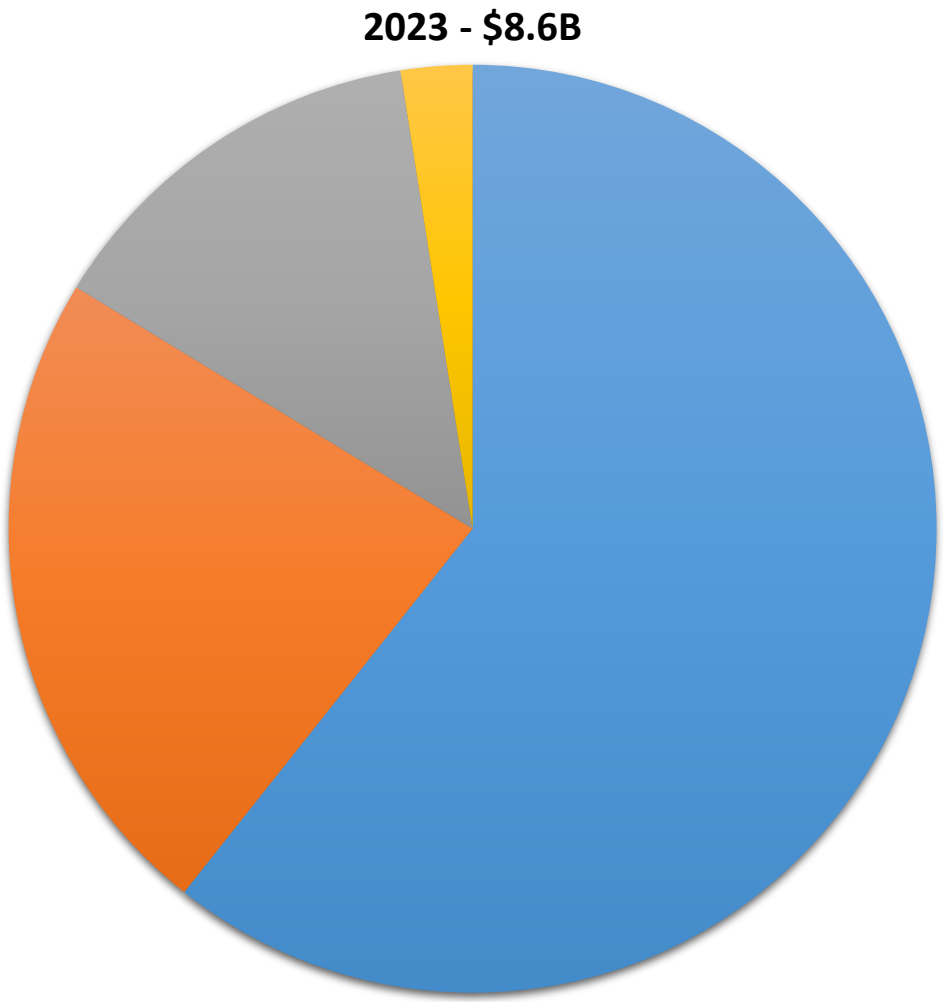

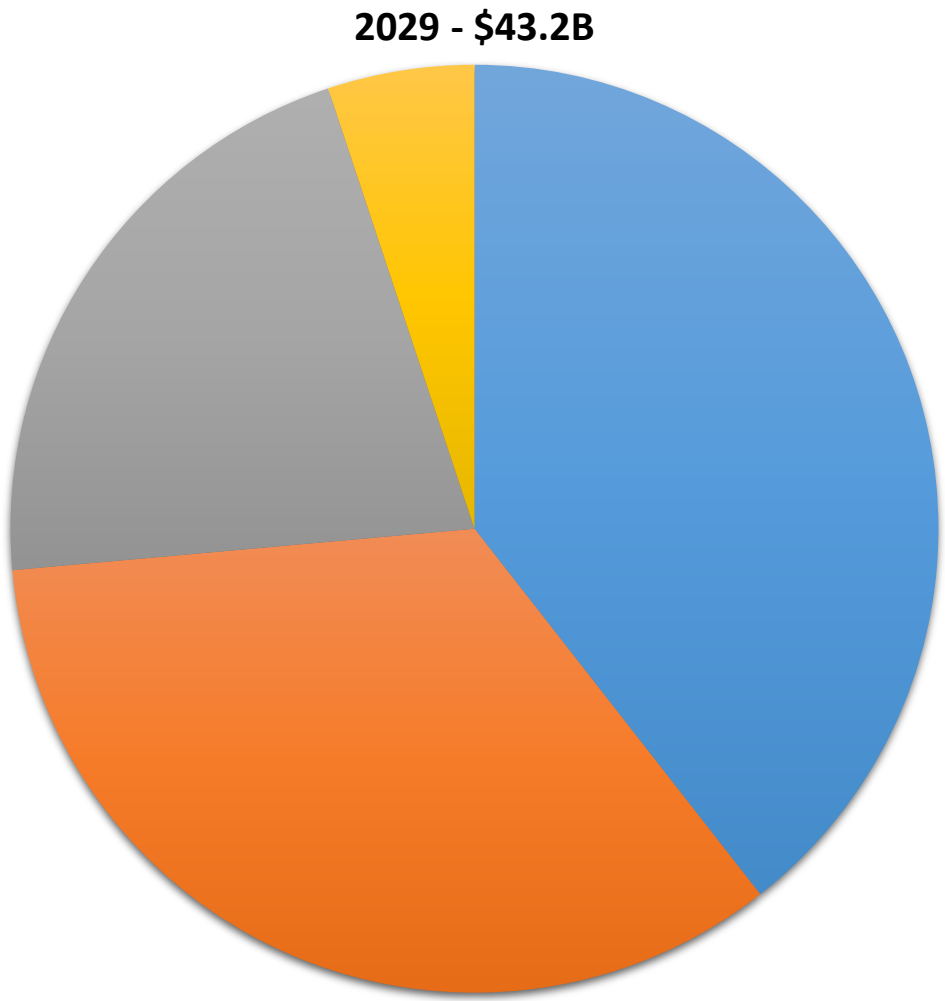

Ethernet Switch – Data Center: Total Market Revenue

650 GROUP
MARKET INTELLIGENCE RESEARCH

Cloud Market

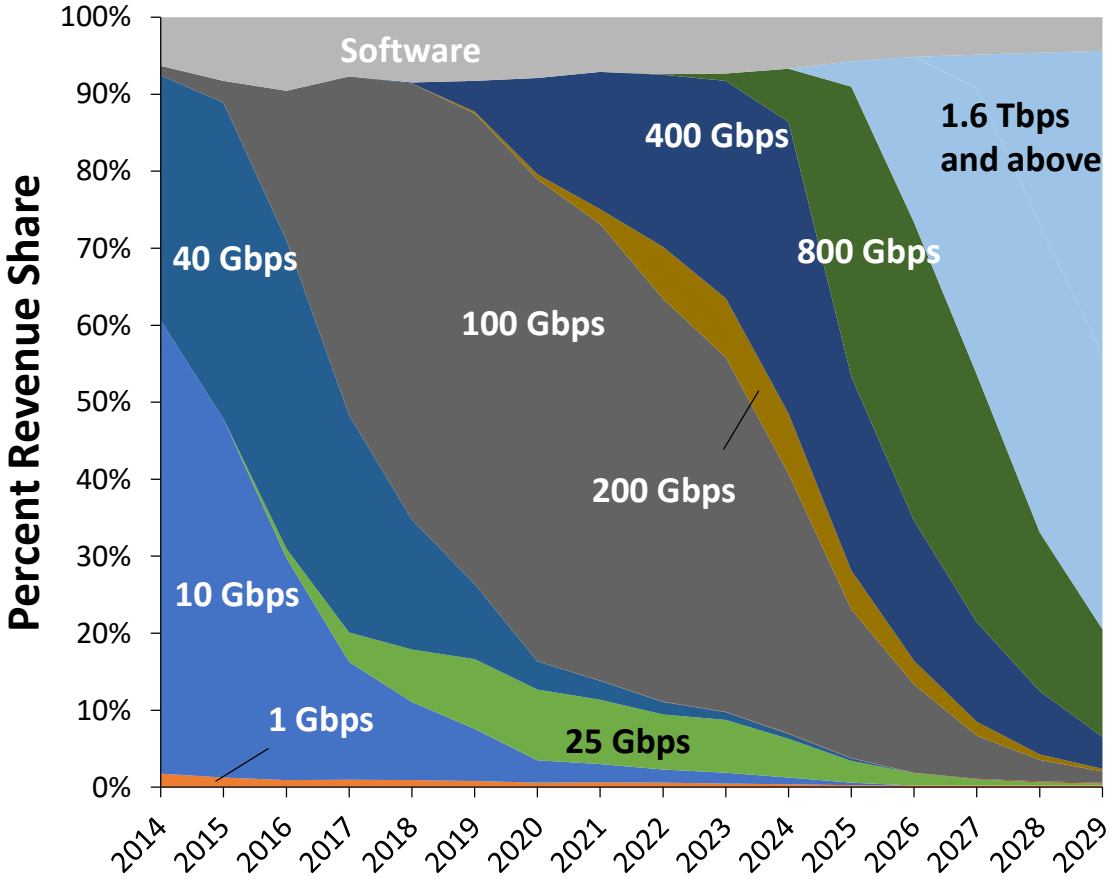

Enterprise and Telco Market

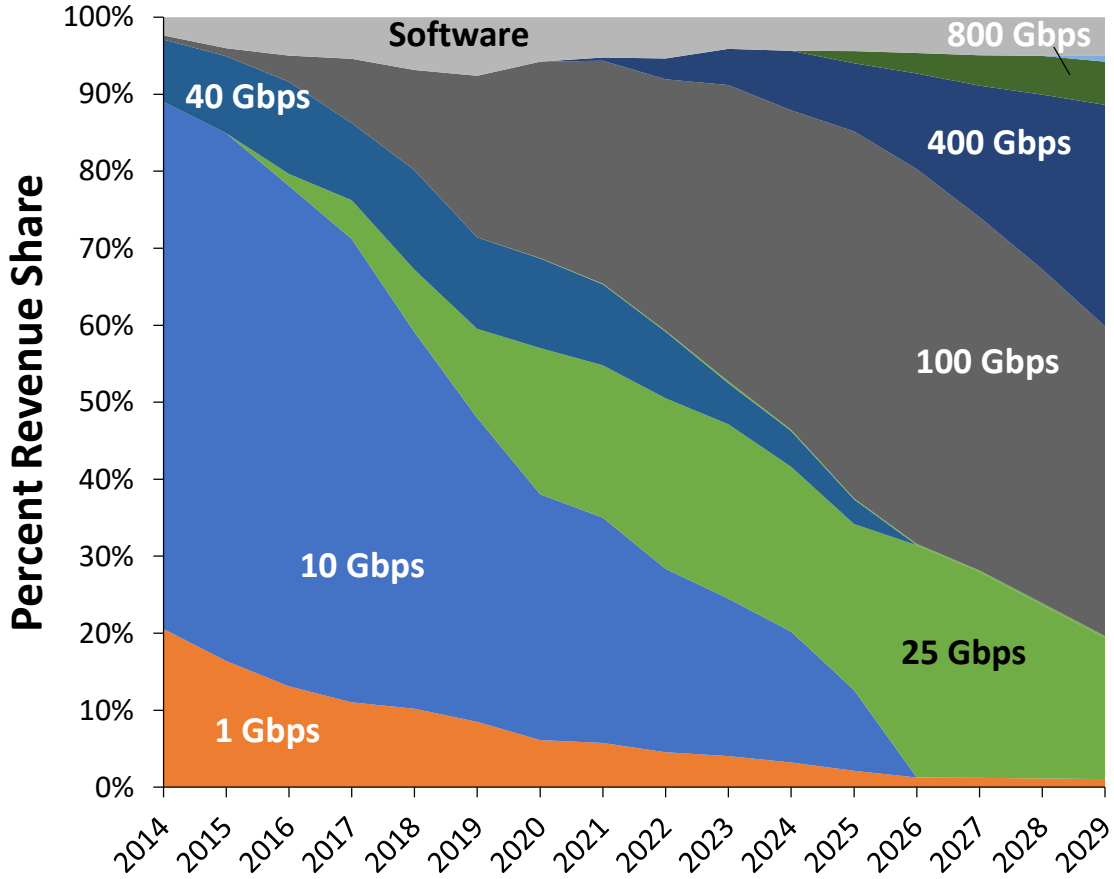

Merchant Silicon – Data Center Switching

Merchant Silicon – Data Center Switching: Total SERDES Shipments

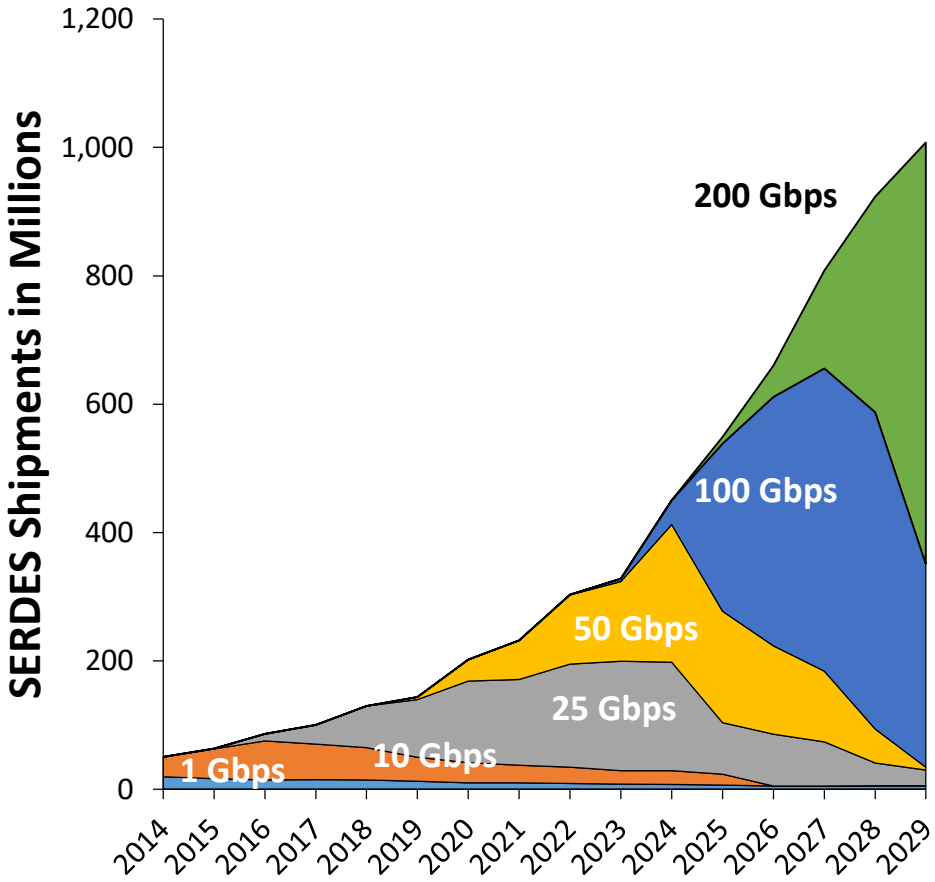

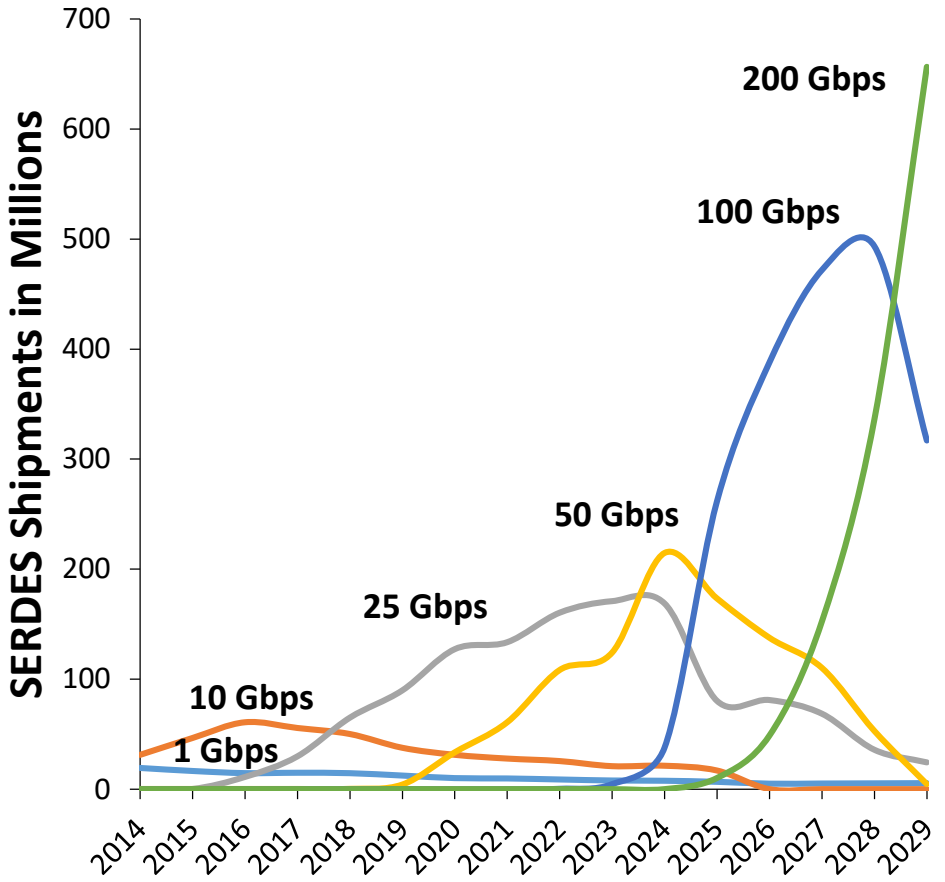

Merchant Silicon – Data Center Switching: ASIC Usage

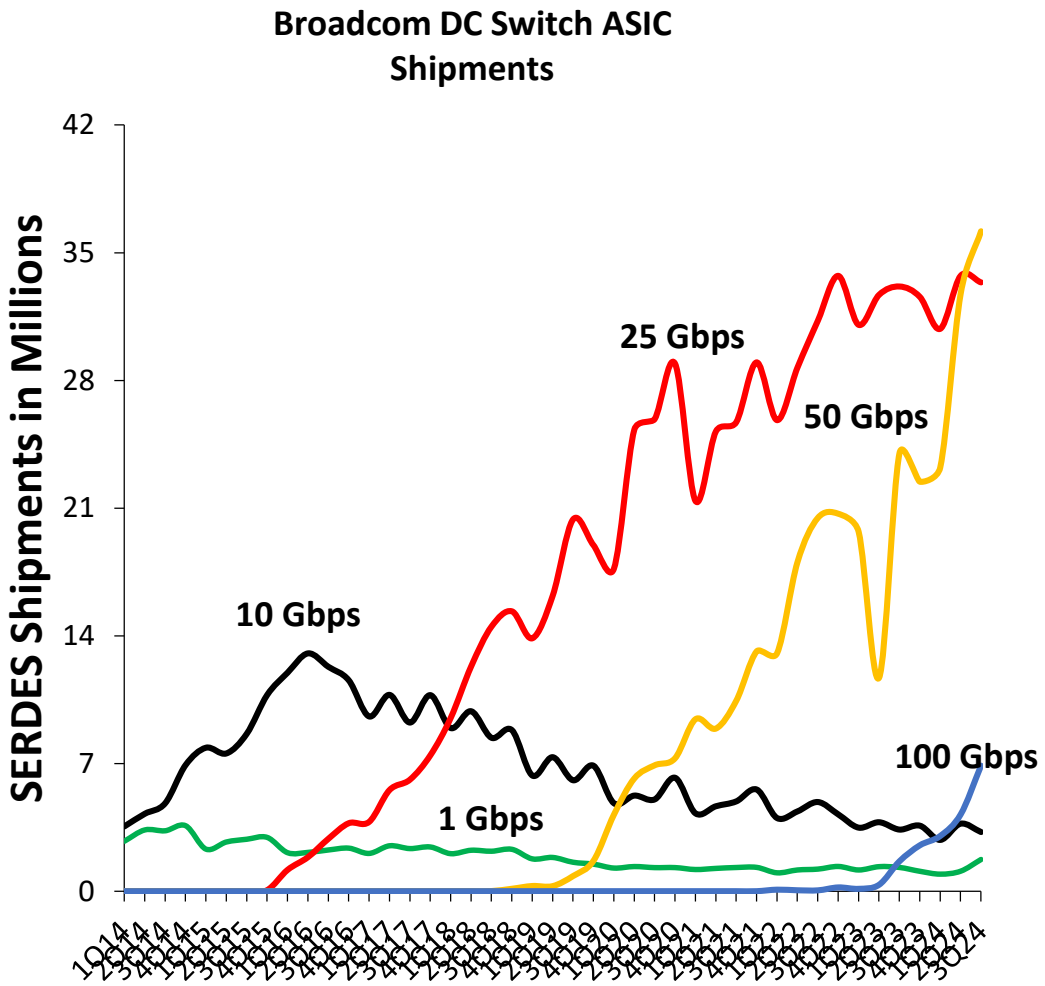

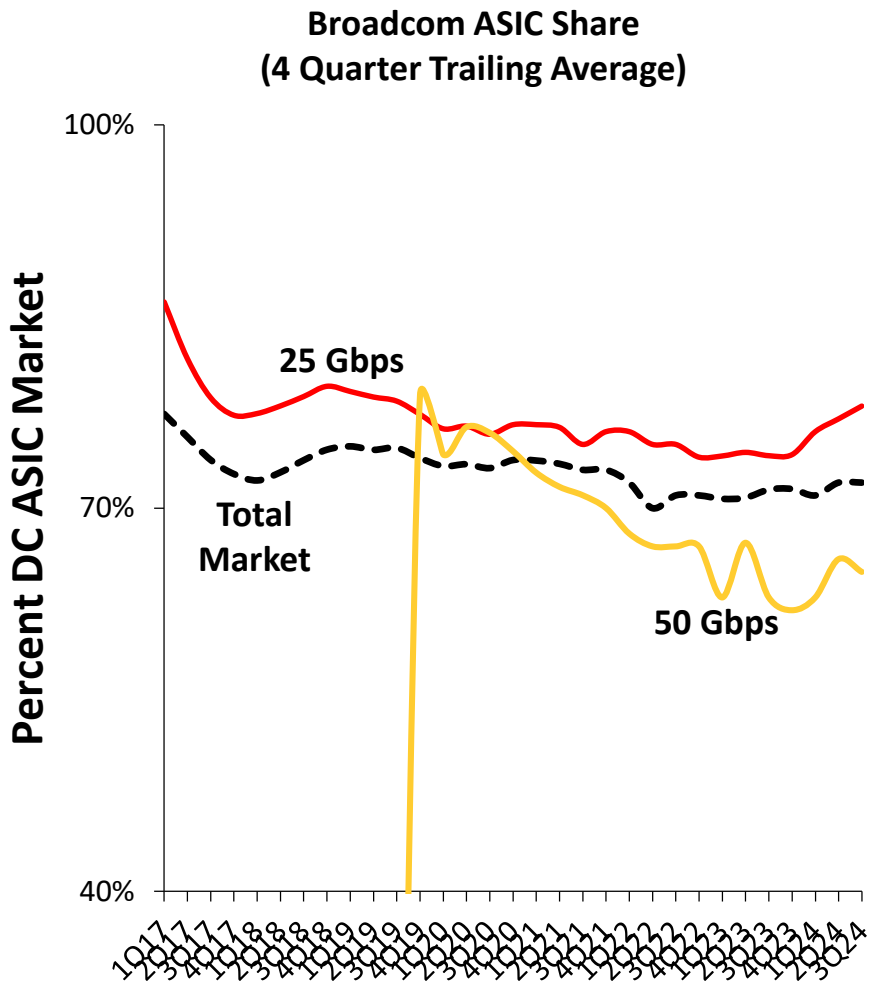

*Does not include backplane ports, only external ports

Merchant Silicon – Data Center Switching: 51.2 Tbps and 102.4 Tbps Announcements

51.2 Tbps

	Form Factor	Process Node	Announcement	Original Shipment Date	Actual Shipment Date
NVIDIA Spectrum 4	512X112G	4nm (NVIDIA Line)	March 2022	Late 2022 / Early 2023	Later Summer 2024
Broadcom Tomahawk 5-100	512X112G	5nm	August,2022	Fall 2022	Late 2023
Cisco	512X112G	5nm	June, 2023	Fall 2023	
Innovium (Marvell) Teralynx 10	512X112G	5nm	March, 2023	Summer 2023	Summer 2024

102.4 Tbps – Started with sampling in late 2024 and early shipment in 2025

	Form Factor	Process Node	Announcement	Original Shipment Date	Actual Shipment Date
NVIDIA Spectrum 6	512x224G	3nm			
Broadcom Tomahawk 6-100	1024X112G	3nm			
Broadcom Tomahawk 6-200	512x224G	3nm			
Cisco TBD	512x224G	3nm			
Marvell Teralynx 100	512x224G	3nm			

Merchant Silicon – Data Center Switching: ASIC Usage in the Tier 1 Cloud

650 GROUP
MARKET INTELLIGENCE RESEARCH

Merchant Silicon's product cycles accelerating in the Cloud

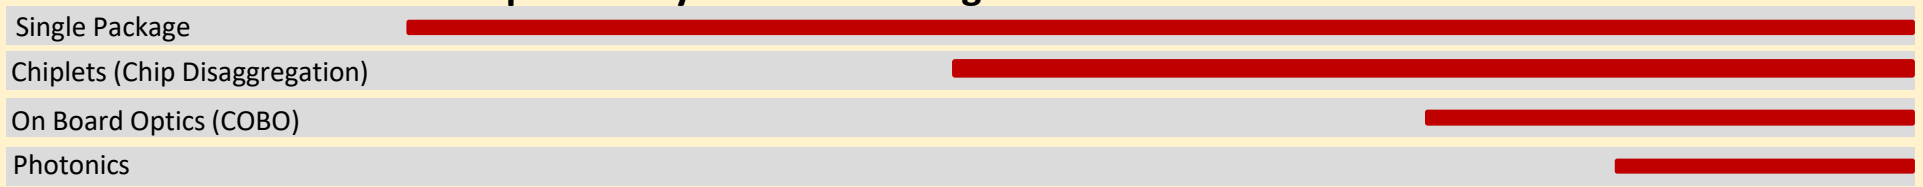

ASIC Size	Process Node	SERDES Technology	Leaf Port Speed	Spine/Core Port Speed
1.3 Tbps	65/40nm	10 Gbps	10 Gbps	10/40 Gbps
3.2 Tbps	28nm	25 Gbps	25/50 Gbps	100 Gbps
6.4 Tbps	16nm	25 Gbps	25/50 Gbps	100 Gbps
12.8 Tbps	16nm	50 Gbps	50/100 Gbps	200/400 Gbps
25.6Tbps	7nm	100 Gbps	100/200 Gbps	200/400 Gbps
51.2 Tbps	5nm	100 Gbps	400/800 Gbps	400/800 Gbps
102.4 Tbps	3nm	200 Gbps	800 G / 1.6 T	1.6 Tbps
204.8 Tbps	2nm	400 Gbps	1.6 T / 3.2 T	3.2 Tbps

- Two waves of 1.6 Tbps
  - 16 X 100 Gbps
  - 8 X 200 Gbps

- Pace of Innovation Increasing
  - Three major silicon cycles in five years
  - Some technologies will get orphaned