AI Driving Architectural Changes in the Datacenter and Beyond

Tom Williams

AI is Transforming Networking and Optics

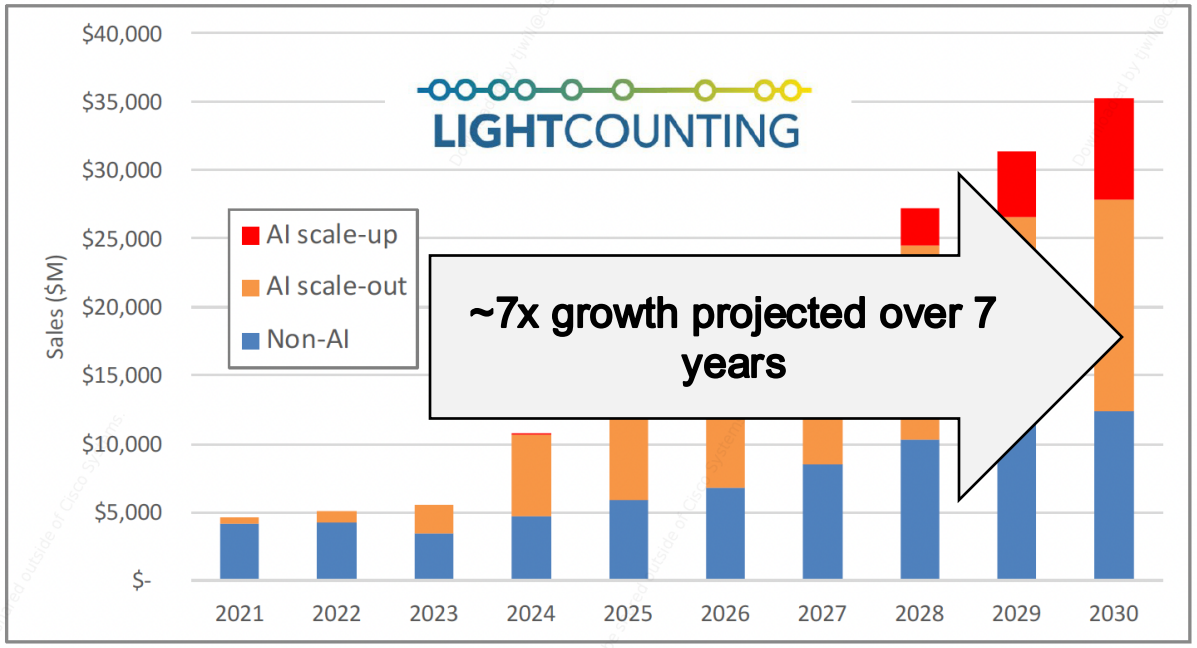

Intra-Datacenter

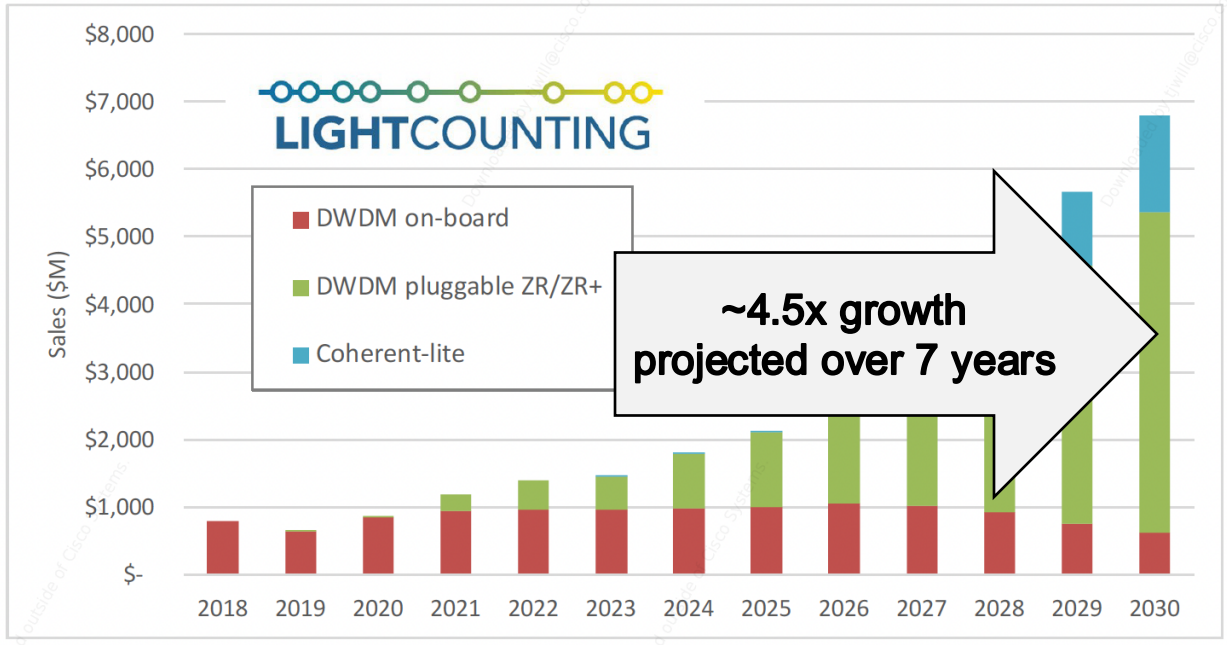

Datacenter Interconnect

Power

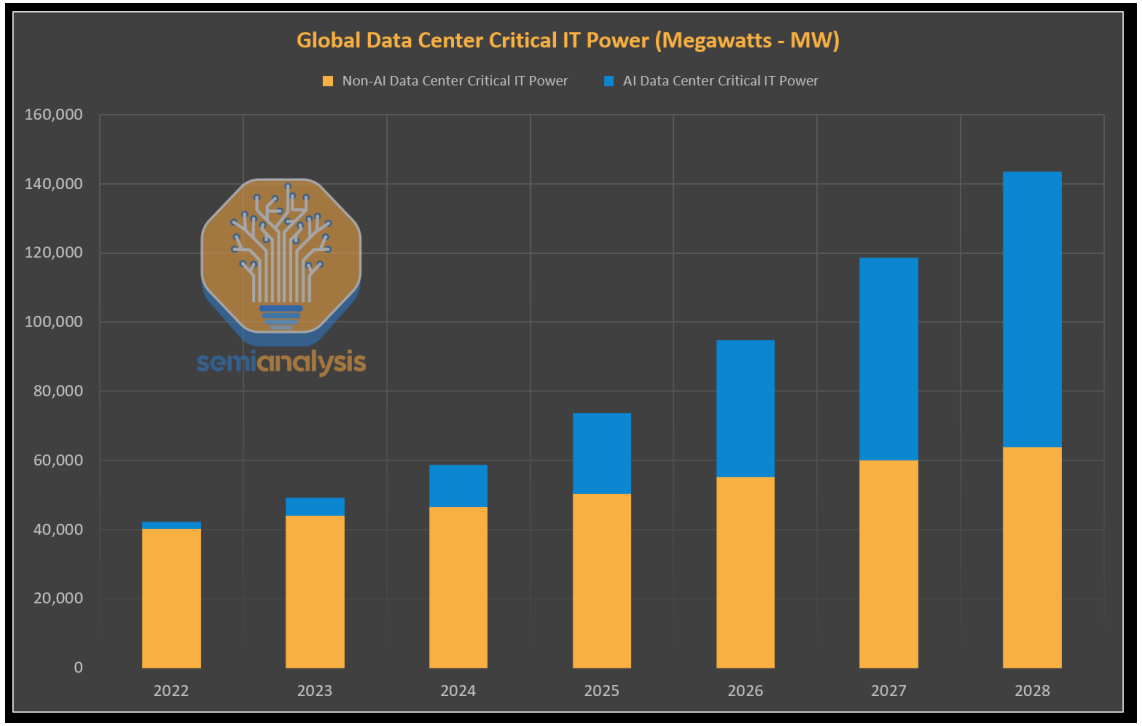

Source: SemiAnalysis

AI Training Power Requirements				
	Units	BLOOM	GPT-3	Meena
Number of Parameters		176B	175B	2.6B
Total Training Time	hours	2,820	355	720
Total Chips used	Chips	384	10,000	1,024
Chip type used		A100 SXM	V100	TPU v3
Chip TDP	W	400	300	450
Dynamic Power Used	**kWh**	**433,196**		
Network, Storage, Cooling (Infrastructure)	**kWh**	**107,305**		
Compute Idle Power	**kWh**	**254,354**		
Total AI Training Power Consumption*	**kWh**	**794,855**	**1,287,000**	**232,000**

** Net of infrastructure component power usage, thus is net of PUE and reflects actual power demand from grid.*

Source: SemiAnalysis

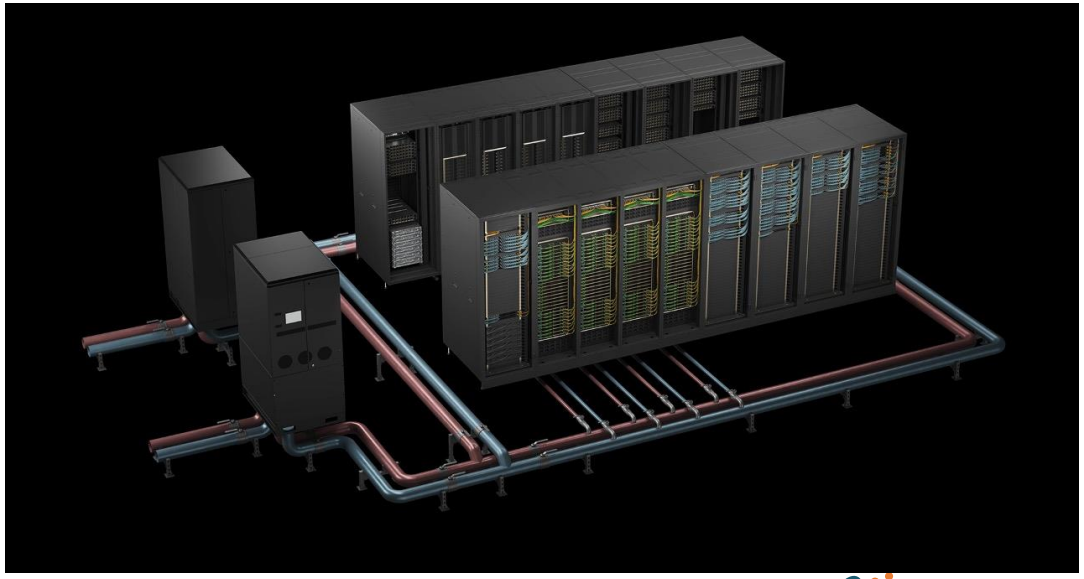

Source: Nvidia

Reducing power impacts cooling, GPU utilization and overall power availability and costs

Reliability

Why reliability impacts AI use case so severely

Scale

$$\text{MTBF} = \frac{1}{(FIT1+FIT2+FIT3.....FITn)}$$

MTBF for a 1M GPU cluster can be on the order of minutes

GPU Utilization

Even short duration link failure (i.e. link flaps) can significantly impact GPU utilization

Reliability is a key element of CPO discussions

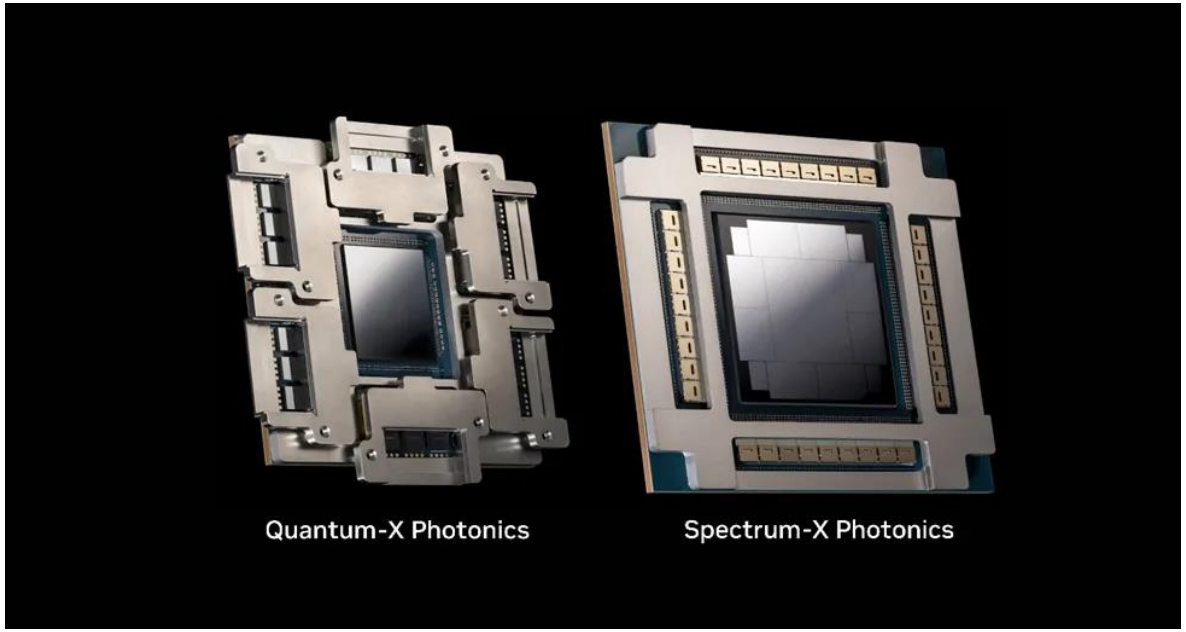

Source: Nvidia

Latency

- Latency is much more important for AI than traditional cloud networking due to the highly parallel process
- Inside the datacenter, it's critical to minimize latency variation to avoid "stragglers"
- Multi-datacenter training possible introduces increased fiber transmission latency, which must be accounted for in training algorithms

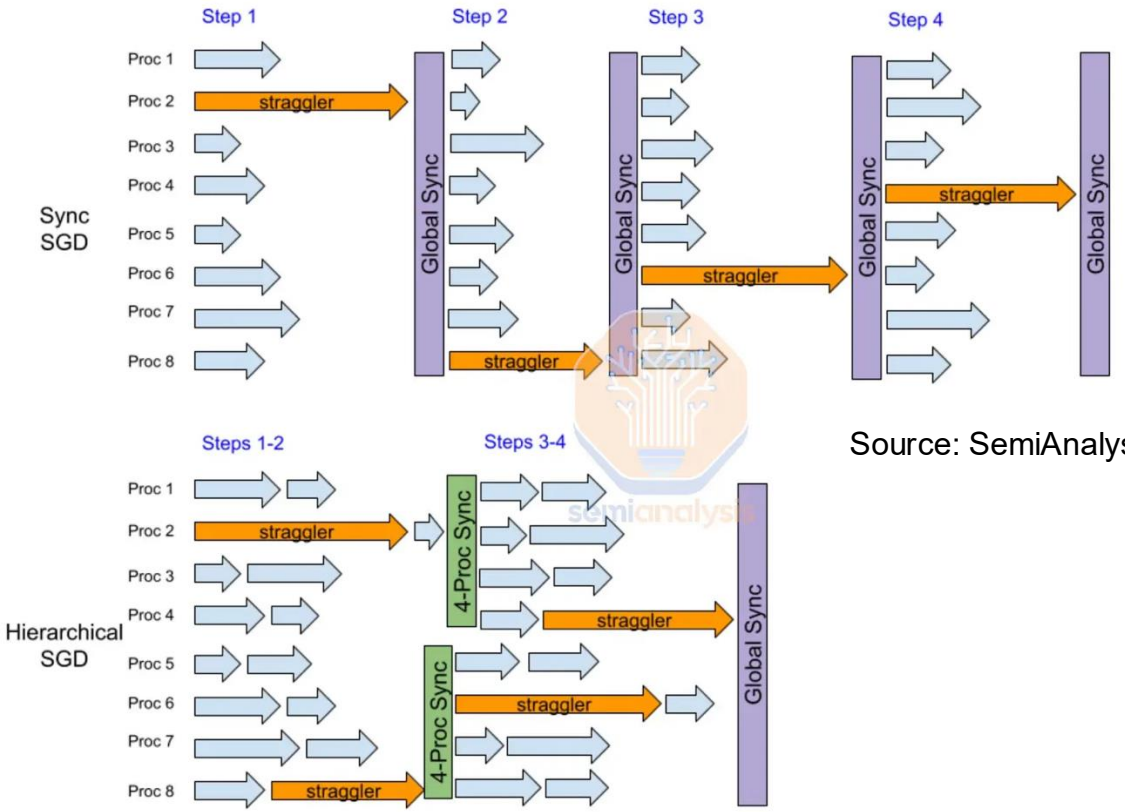

Source: SemiAnalysis

Enabling Massive DCI Connectivity – Class 3 Generation

- 800ZR/ZR+ is ramping and will play a critical role in AI buildouts over the next few years
  - 800ZR from OIF for metro DCI
  - 800G ZR+ from OpenROADM for regional DCI
  - 600G ZR+ from OpenROADM for long-haul DCI
  - 400G ULH from Acacia for maximum reach
- 400ZR continues to be widely deployed for metro DCI

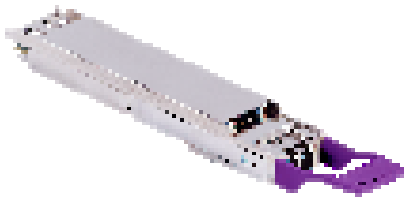

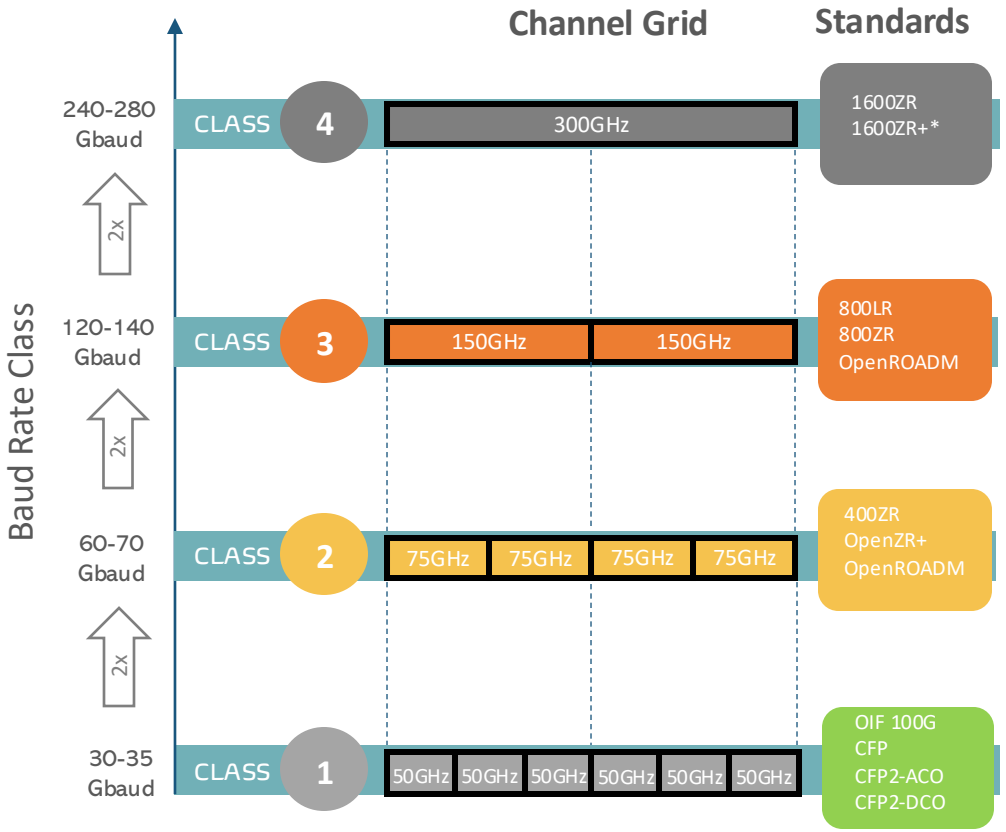

Enabling Massive DCI Connectivity – Class 4 Generation

- Higher baud rates reduce the number of wavelength per optical band
  - C+L deployments widely utilized
- Power optimized architectures, including amplification and regeneration
- Greater integration and fewer interconnections

These new architectures are expected to coexist with pluggable implementations

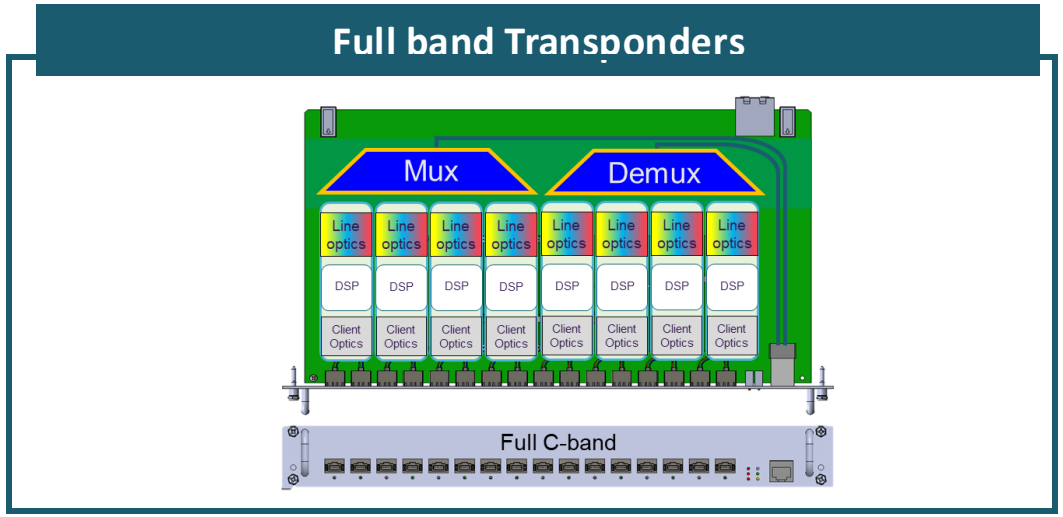

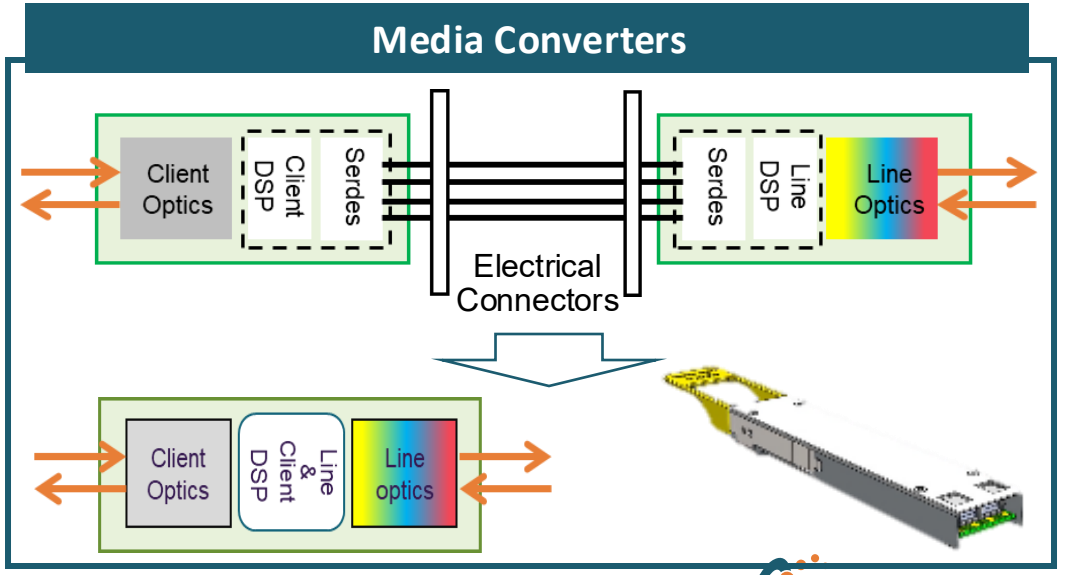

Connecting at the speed of light