



CXL: Transforming Cloud Infrastructure

Khurram Malik
Director, Product Marketing
August 2023



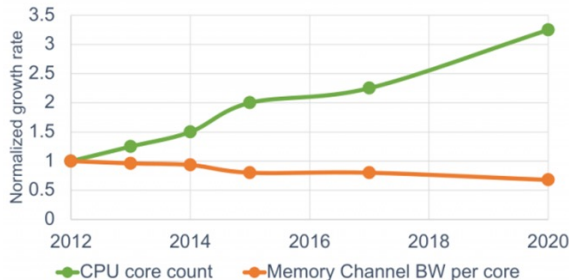
Disclaimer

This presentation contains forward-looking statements within the meaning of U.S. federal securities laws that involve risks and uncertainties. Forward-looking statements include, without limitation, any statement that may predict, forecast, indicate or imply future events or achievements. Actual events or results may differ materially from those contemplated in this presentation. Forward-looking statements are only predictions and are subject to risks, uncertainties and assumptions that are difficult to predict, including those described in the “Risk Factors” section of our Annual Reports on Form 10-K, Quarterly Reports on Form 10-Q and other documents filed by us from time to time with the SEC. Forward-looking statements speak only as of the date they are made. Readers are cautioned not to put undue reliance on forward-looking statements, and no person assumes any obligation to update or revise any such forward-looking statements, whether as a result of new information, future events or otherwise.

Why does this matter to hyperscalers

Fixes scaling challenges

DRAM BW per core historical decline
Compute efficiency ↓

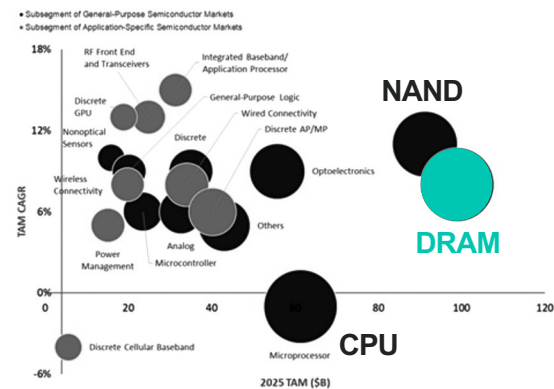


Source: Meta, OCP Summit Presentation Nov 2021

CPU core growth outpaces bandwidth

Maximizes utilization

Semiconductor Revenue Forecast by Device Type Through 2025




Source: Gartner 2021


DRAM #1 spend

Enables SKU flexibility

SSD



Memory Expander



E1.S (shown); E3.S supports more power & capacity


Pluggable, leverage & unify SSD FFs

DRAM key to unlocking TCO and Perf for emerging workloads

CXL use cases


Memory Expansion

- Scale performance and capacity
- Mix-and-match DRAM types
- Thermally optimized




Memory Acceleration

- Coherent, efficient
- Accelerate analytics, ML, etc.
- Improves efficiency & TCO




Memory Pooling

- Share memory across xPUs
- Improve utilization
- BW:cost optimized

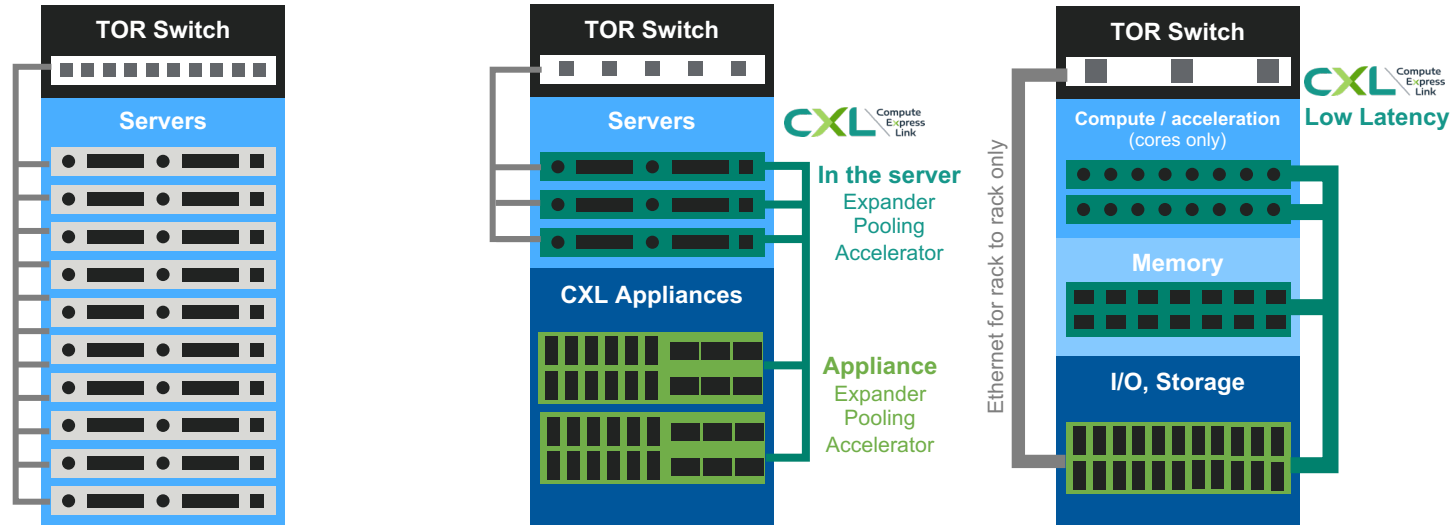


Switch

- Low-latency fabric
- Supports optics
- Enables full composability

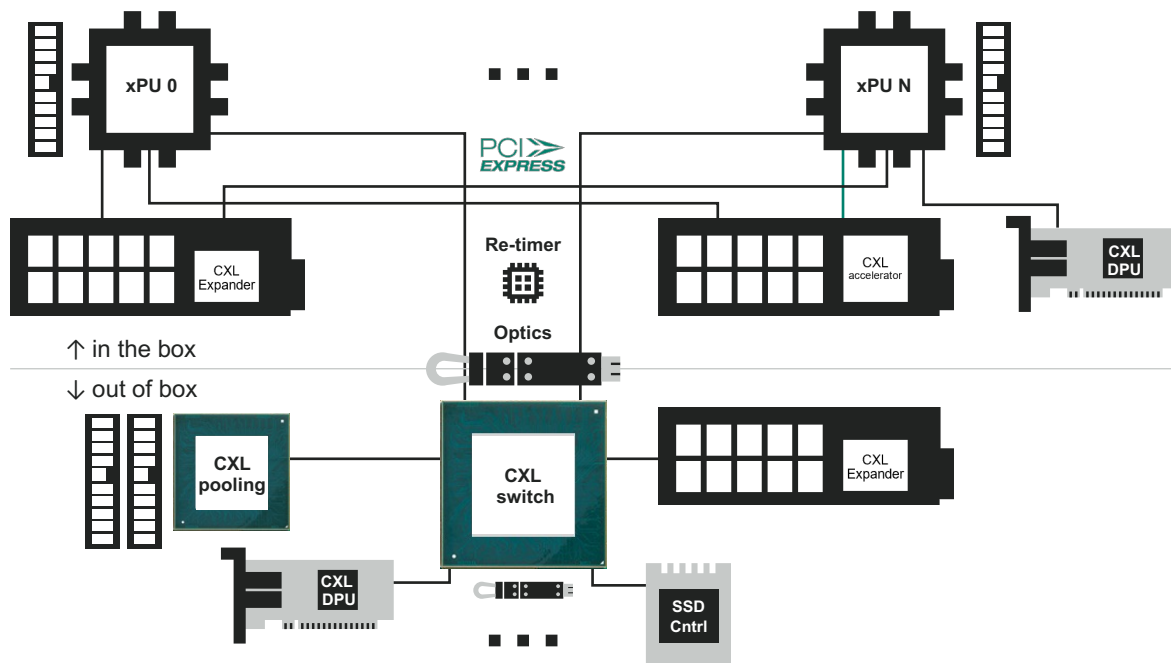


CXL vision: optimal resource utilization



Rack Fabric	Ethernet	Ethernet+CXL	CXL Only
Disaggregation	Storage and Accelerators	+Partial memory	+Full memory, compute

Comprehensive end-to-end CXL solutions



CXL use cases

- Expanders
- Pooling
- Switch
- Accelerators
- Electro-optics
- Re-timers
- Custom Compute
- DPUs / SmartNICs
- SSD Controllers

Built on industry's most comprehensive cloud portfolio



Thank You



Essential technology, done right™