Lessons learned from GW-scale AI deployments

Mark Filer

Architect, Oracle Cloud Infrastructure

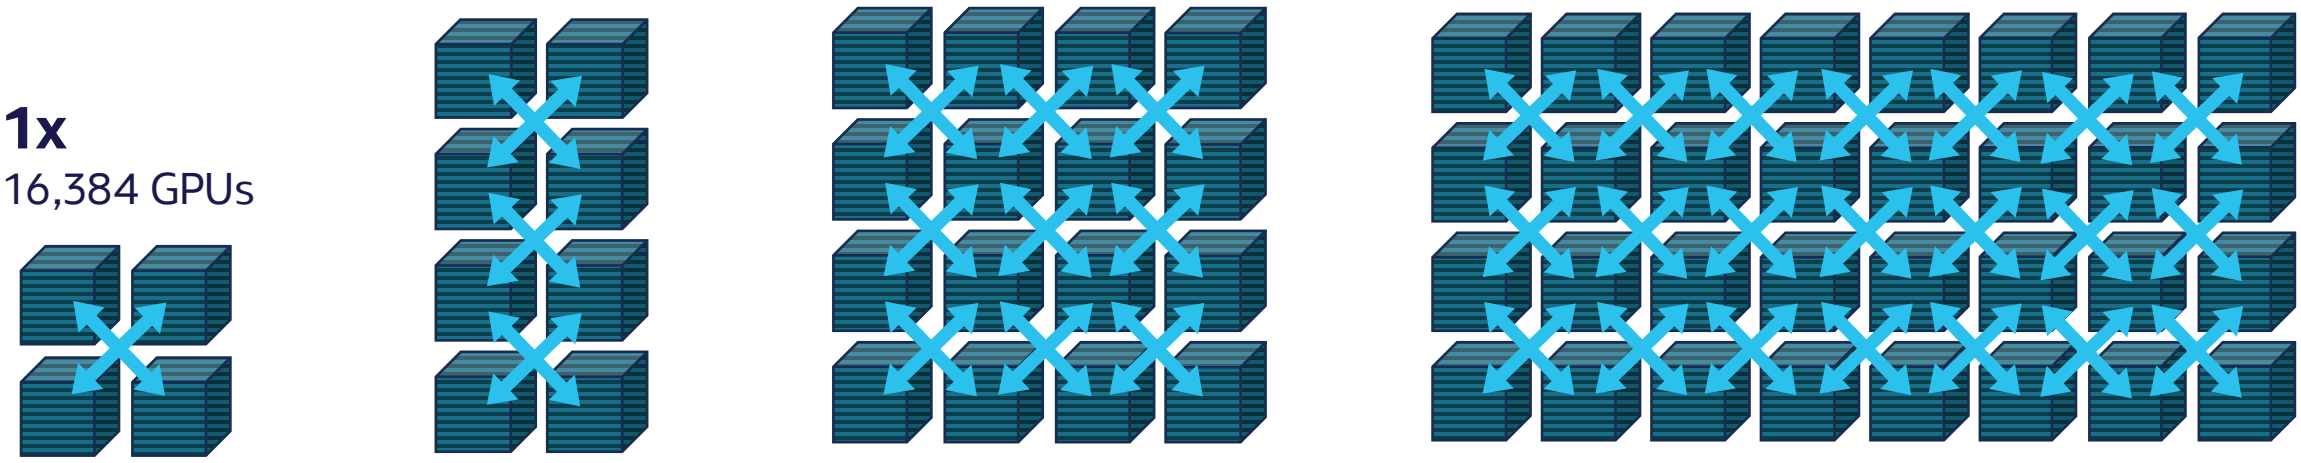
Oracle AI cluster growth

1x
16,384 GPUs

2x
32,768 GPUs

4x
65,536 GPUs

8x
131,072 GPUs



2020

2026

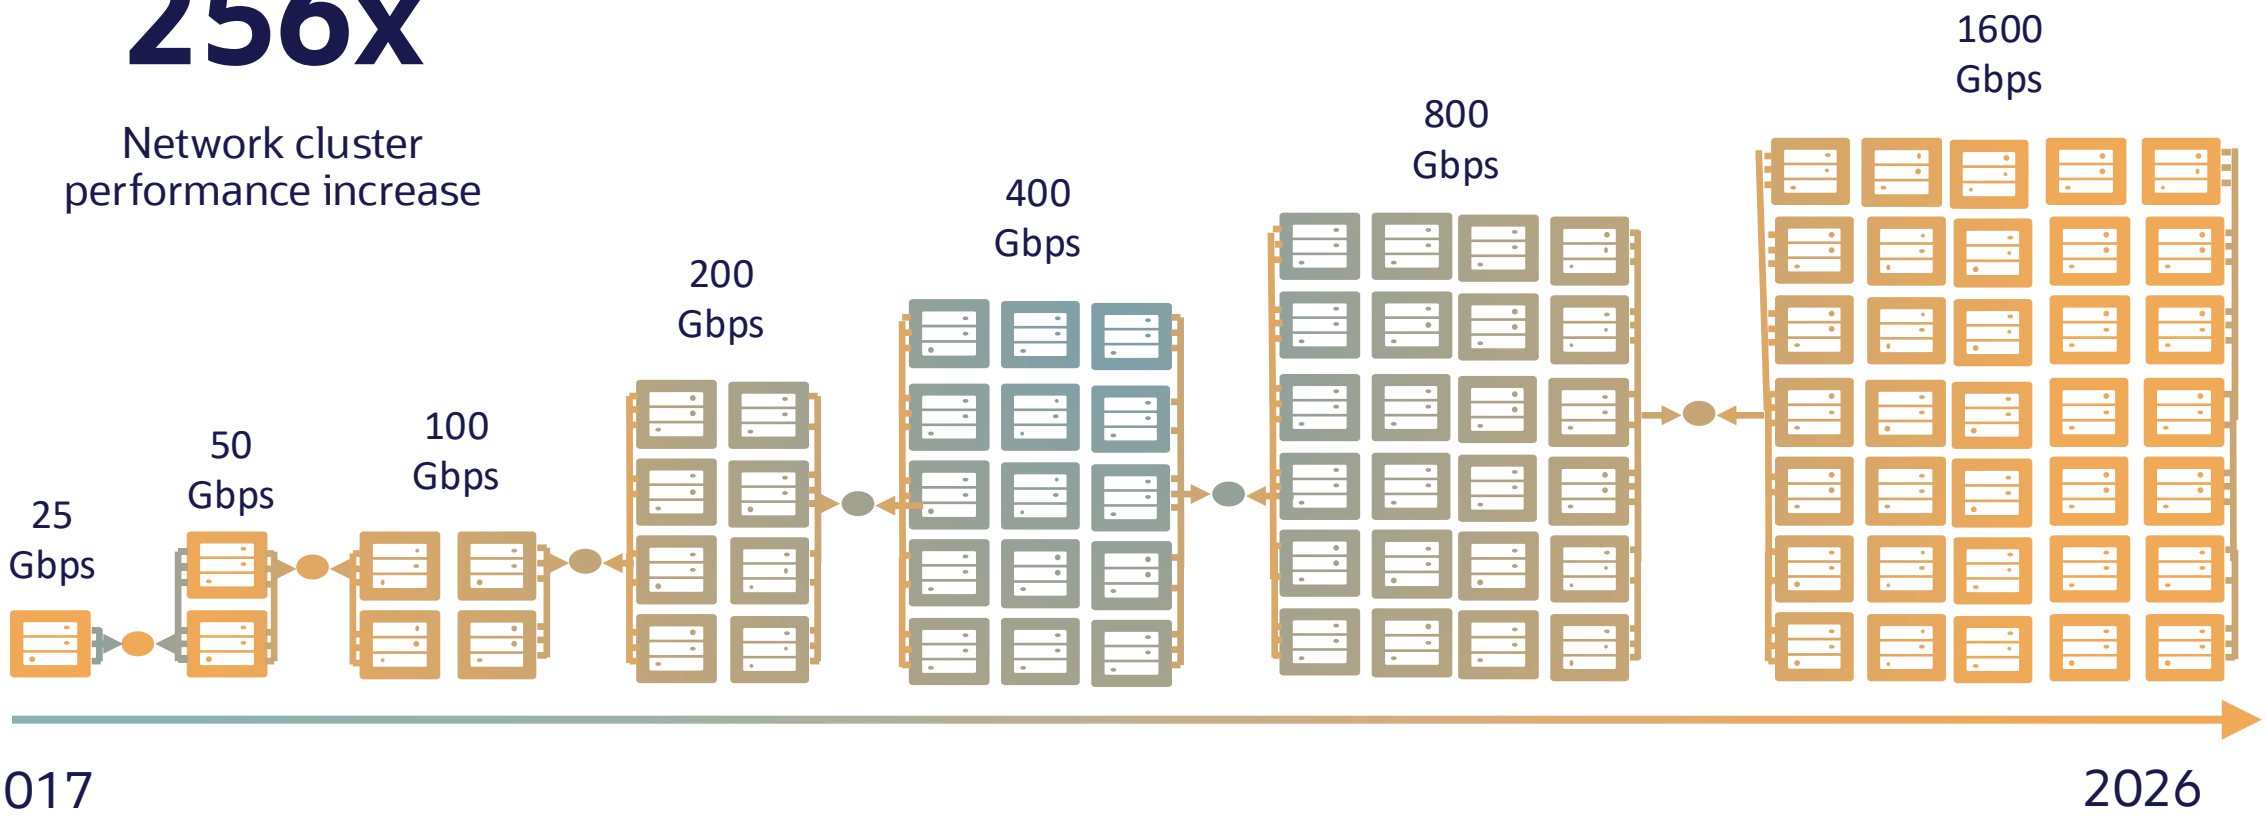
Enabled by:

1. RDMA at Zettascale with ultra-high throughput
2. Ultra-low latency
3. High network resilience
4. Advanced traffic load-balancing

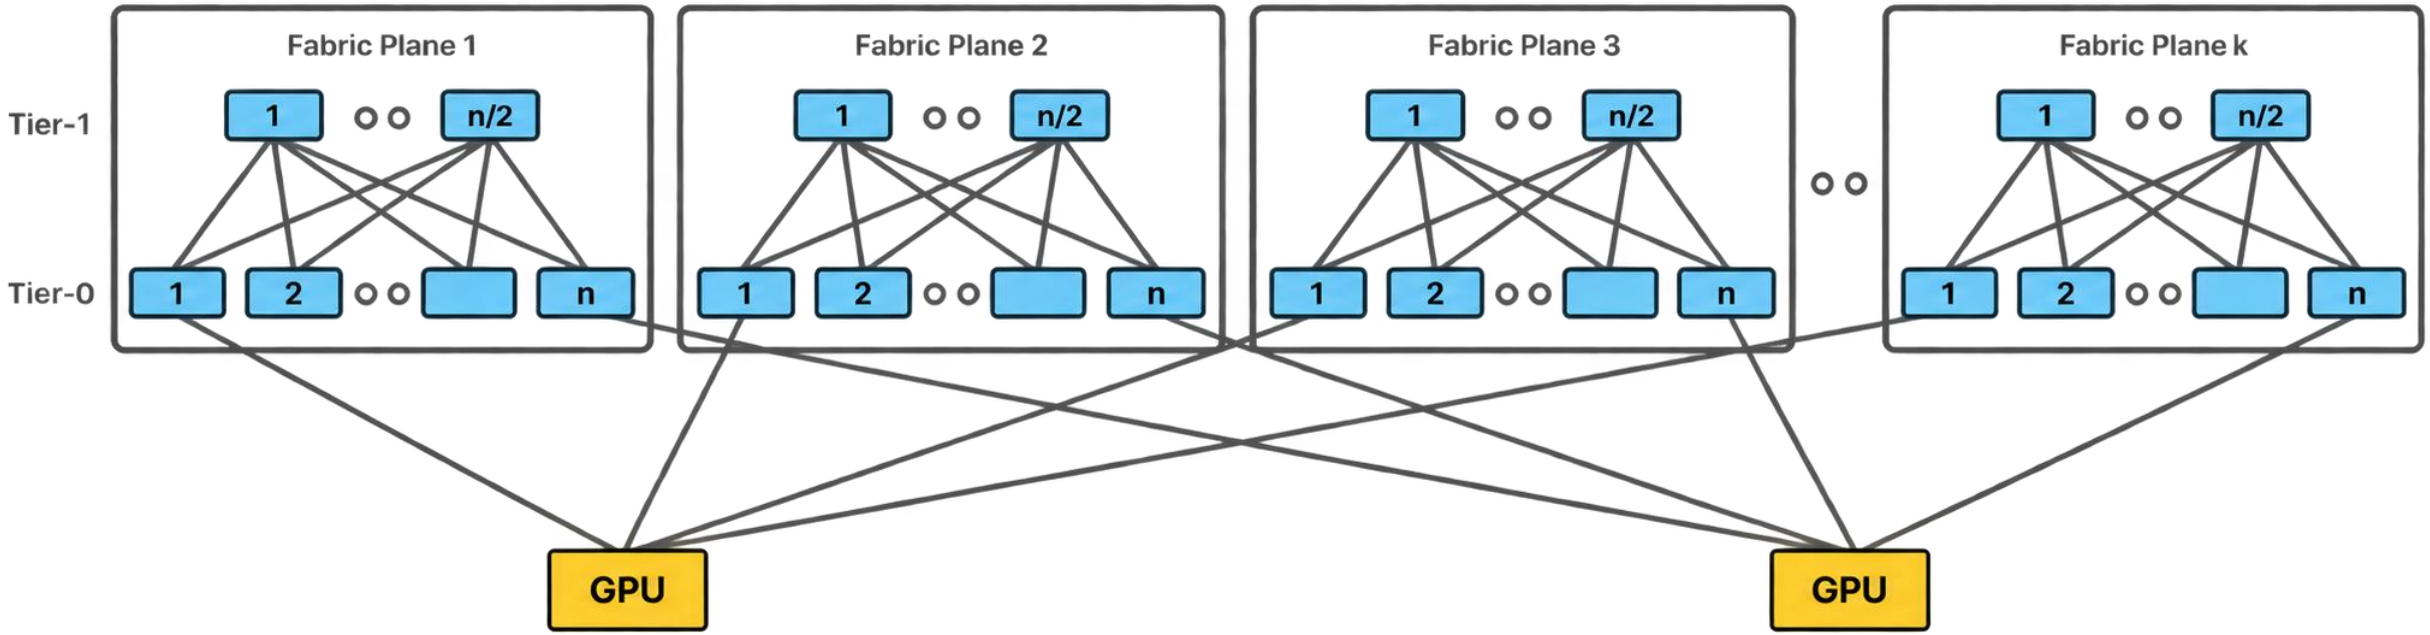
Network speed growth (NIC interfaces)

256x

Network cluster performance increase



Oracle RDMA multi-plane scale-out cluster



Features

- Standardized network interconnects, including LPO
- Large-radix devices, non-blocking Clos
- Multi-plane, rail-disjoint
- 1.6T breakouts: 8x200G, 4x400G, 2x800G, 1x1600G
- DR, FR, LR are all in scope

Challenges

- RDMA transport is lossless and sensitive to disruptions – network availability is critical
- Network power envelope is non-negligible at Zettascale
- Dust / contamination a primary source of link failure (attributed to 90% of link issues)

https://blogs.oracle.com/cloud-infrastructure/first-principles-acceleron-multiplanar-networking

What are we solving for in next-gen networks?

At the end of the day, we're solving for TCO…



Network Availability

Network up-time, higher MFU, service credit avoidance

Direct revenue impact.



Power efficiency

Less power ceded to network, more power for compute

Direct revenue impact.



Cost (CapEx)

Relative spend of network to compute is increasing YoY

Not a sustainable trajectory.

800G: Linear Pluggable Optics (LPO)

Addressing TCO at 800G, OCI has large deployment of LPO transceivers in AI scale-out networks

Plots show 24-hr max pre-FEC BER across *identical* T0<>T1 links, 156k LPO and 202k FRO in data set(359k total)

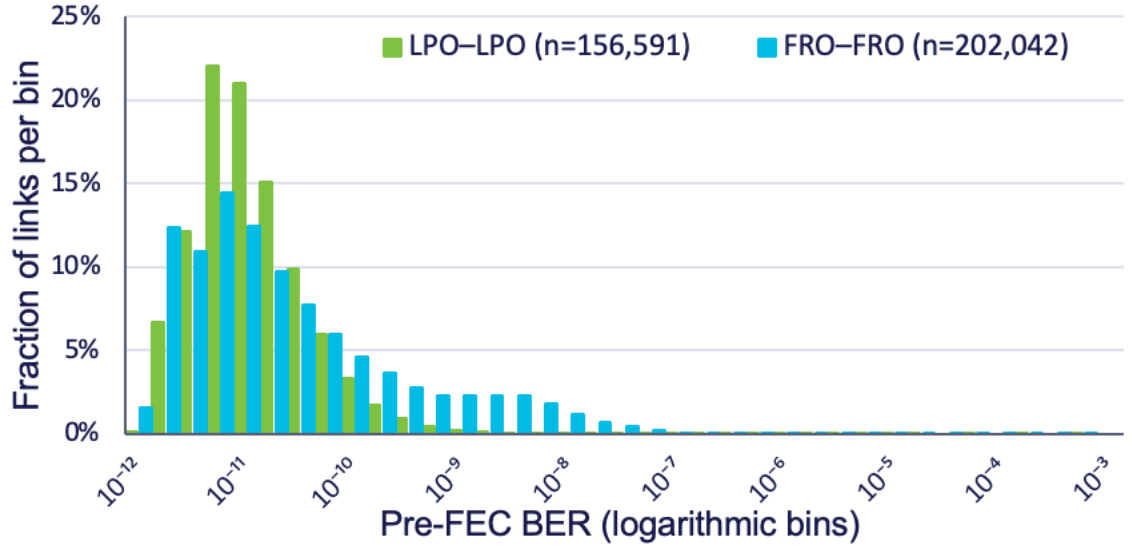
Observations

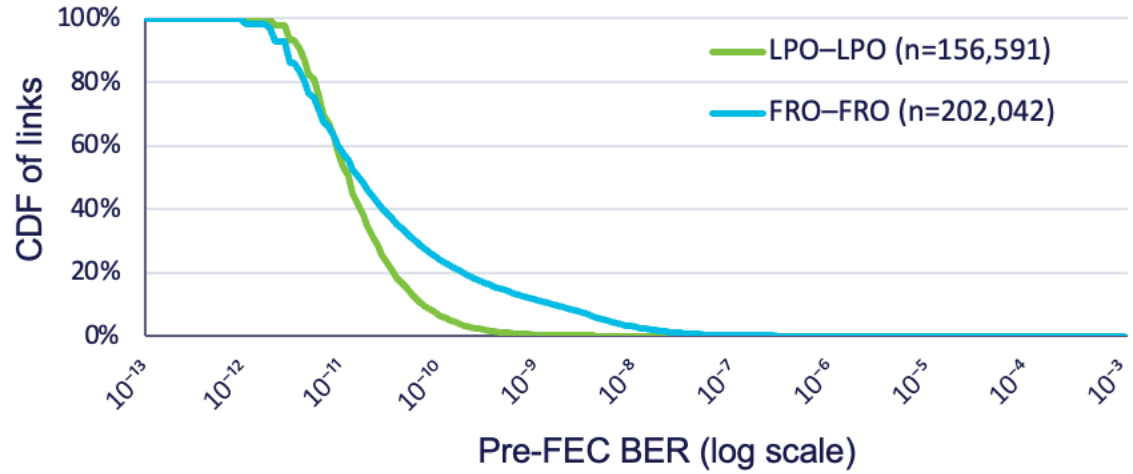
- LPO marginally outperforms FRO in median pre-FEC BER
- LPO greatly outperforms FRO in p99 pre-FEC BER

Link type	Valid links	Median BER	p99 BER
LPO–LPO	156,591	1.137×10^{-11}	5.165×10^{-10}
FRO–FRO	202,042	1.432×10^{-11}	2.888×10^{-8}

Caveats

- Short, med, and high loss ports require optimized tuning
- Same switch ASIC for all (LPO + FRO)
- Diverse optical DSPs for FRO





800G: Linear Pluggable Optics (LPO) II

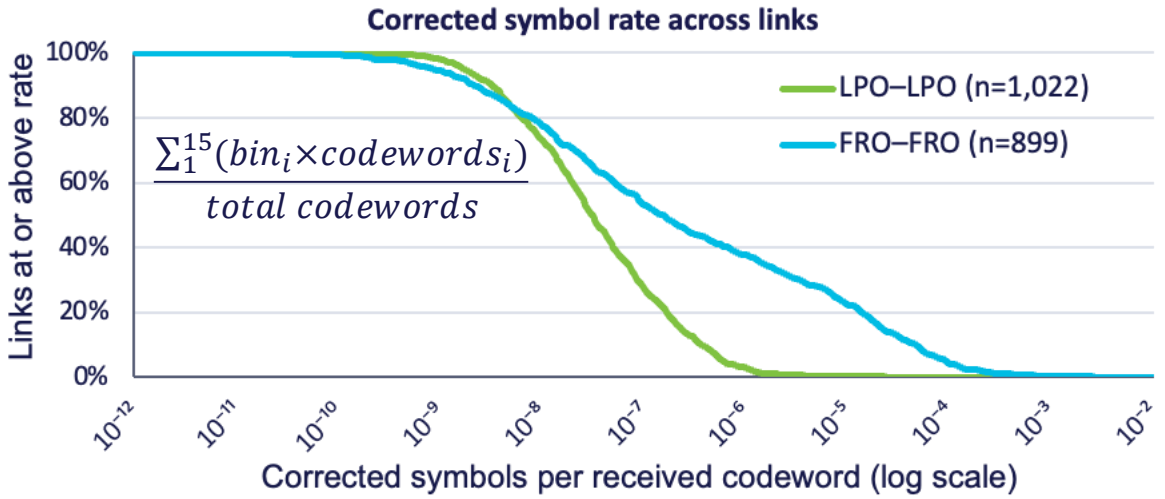
Top chart: corrected symbol rate as proxy for how "hard" the FEC is working looking at occupied FEC bins

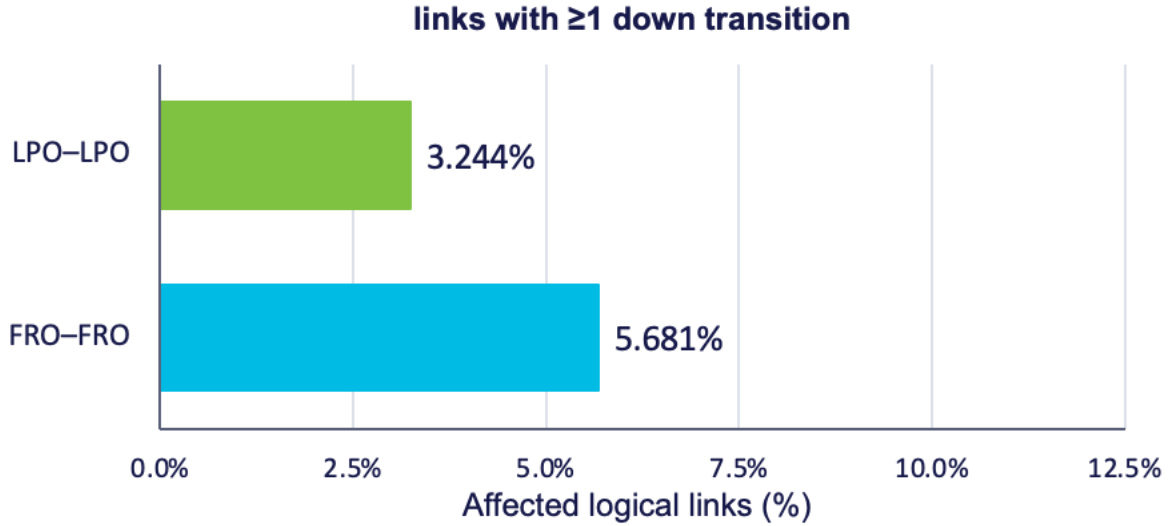
- Lower and farther left is better – fewer corrected symbols for the same amount of received traffic
- FRO sits above and to right of LPO across most of the population, i.e., in general, FRO presents the decoder with more correctable symbol errors
- At median, LPO is 3.1E-8 corrected symbols per codeword, while FRO is 1.6E-7 – roughly 5x higher correction activity

Bottom chart: fraction of links that experienced at least one transition (i.e., flap) over a 1 month period

- FRO links experienced 1.75x more transitions than LPO

Takeaway: Linear drive interfaces don't introduce penalties, and in fact seem to provide better, more stable link performance for AI systems





800G: Linear Pluggable Optics (LPO) III

64x800G switch system-level power consumption

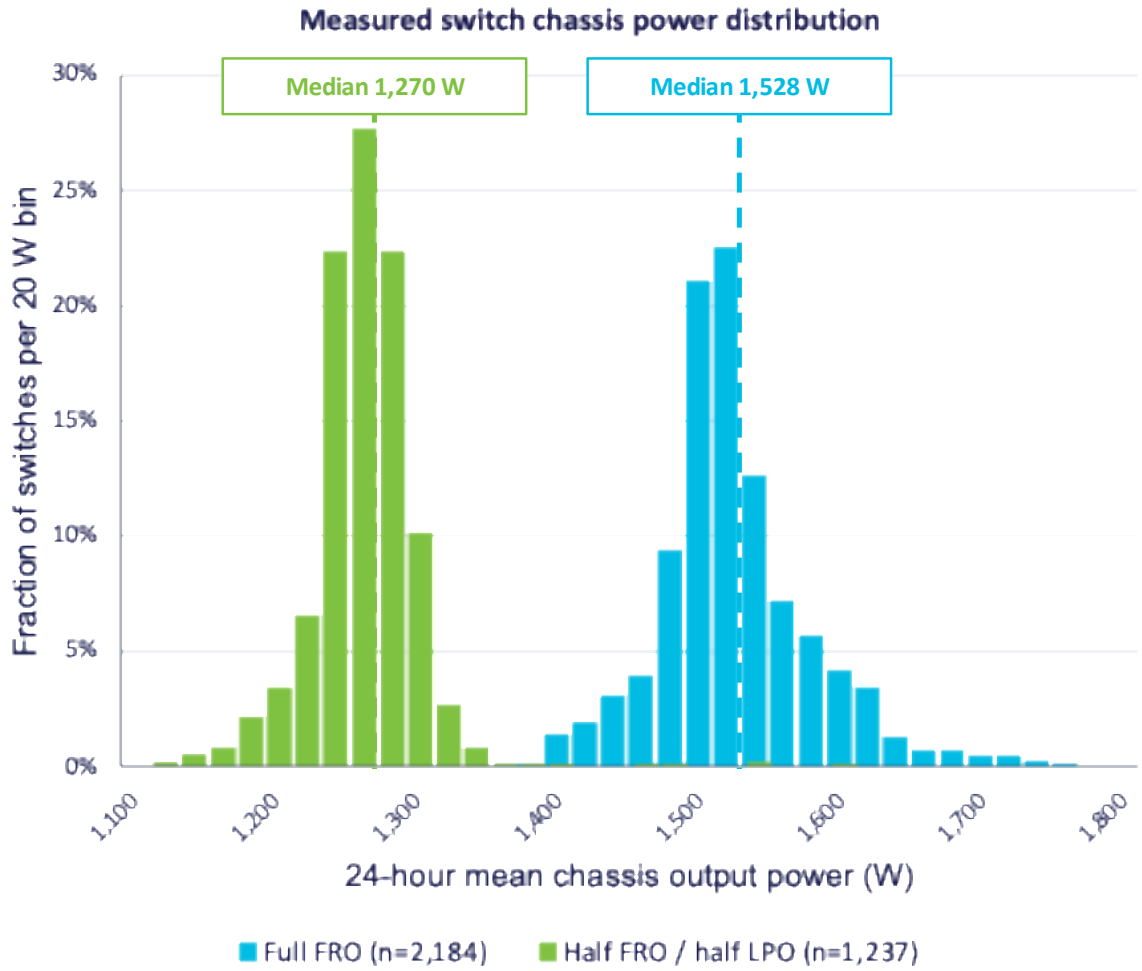
Chart compares switches deployed with 100% FRO modules vs those with 50% LPO / 50% LRO:

- Median full FRO power = 1528W
- Median half LPO/FRO power = 1270W
- Net savings = 258W

800G LPO modules save 5-7W compared to FRO. Expected savings due to modules alone: 32x ~6W = 192W

Observed net savings (258W) vs savings due to modules (192W) indicate 66W additional due to system (fan) power reduction.

Takeaway: Confirmation that system-level power savings are compounded for LPO due reduced thermal burden of linear interfaces

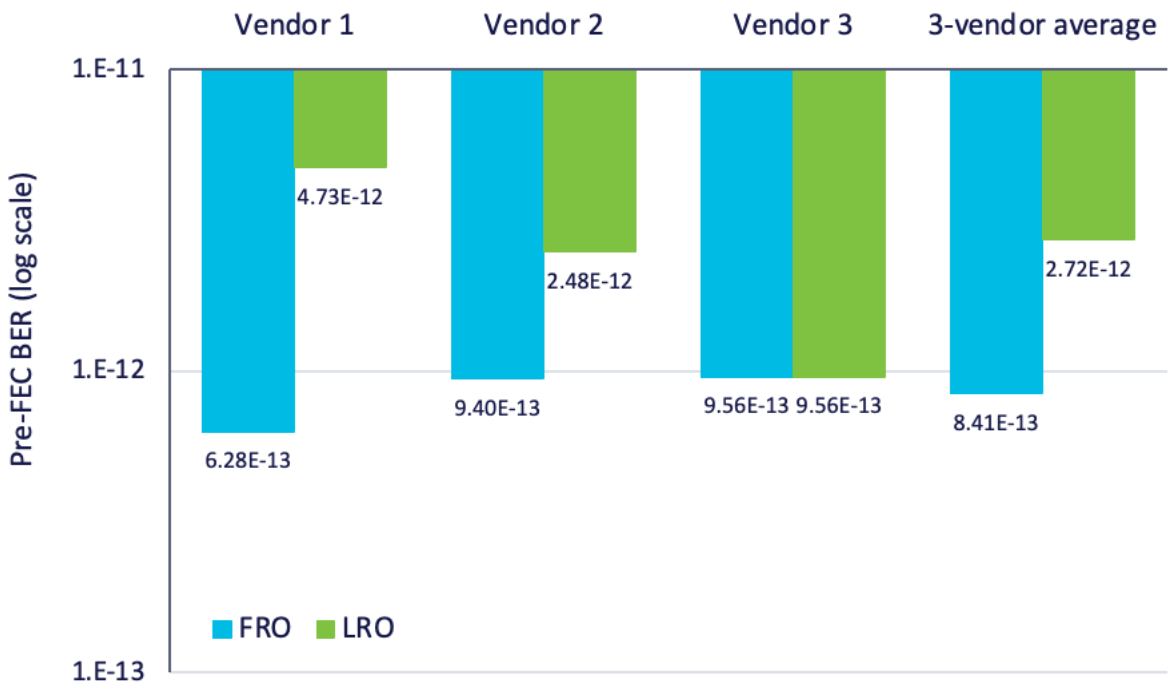


1.6T: Linear Receive Optics (LRO)

LPO cannot comfortably meet SI and interop performance requirements at 1.6T; LRO provides the best low power solution currently available:

- 26dB channels at 16W power consumption (10 pj/bit)
- Performs equally well over short, medium, long reach channels
- Supports DR4+ (4dB), and FR4 links in the upper tiers
- Interop has been demonstrated from different suppliers and multi-vendor 3nm DSPs
- More aggressive power reduction expected as LRO DSP technology matures in 2nm.

Takeaway: LRO is viable at 1.6T and markedly improves network TCO for AI scale-out applications



configuration	FEC bin p50	FEC bin p95
LRO	1	2
FRO	1	1

Beyond pluggables – CPO/NPO systems

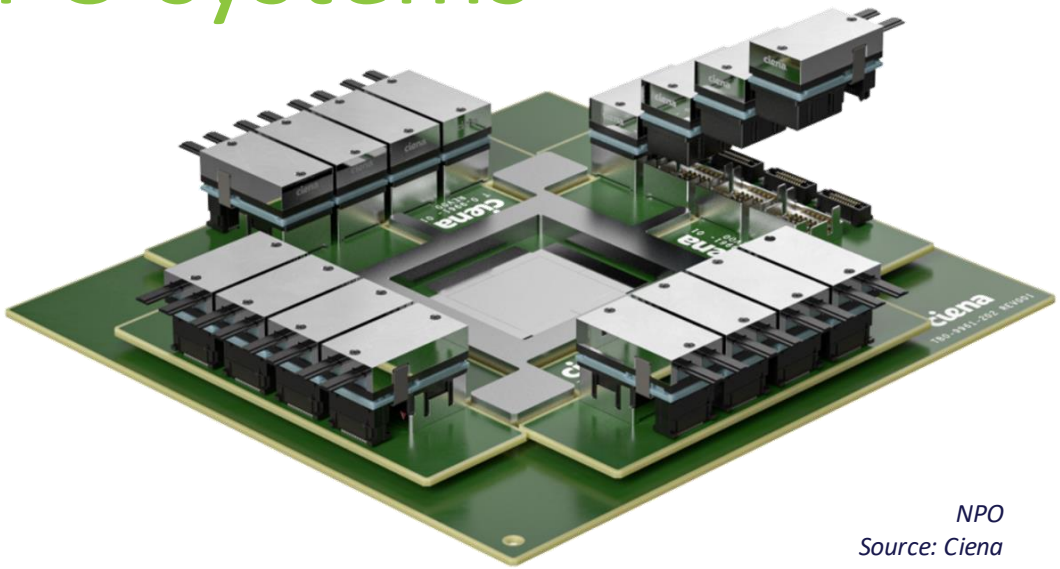
Reliability / Availability

- (Mis)handling of transceivers a primary failure mode – CPO removes humans from the deploy and ops process
- System FIT expected *lower* than equivalent pluggable-based system[1]

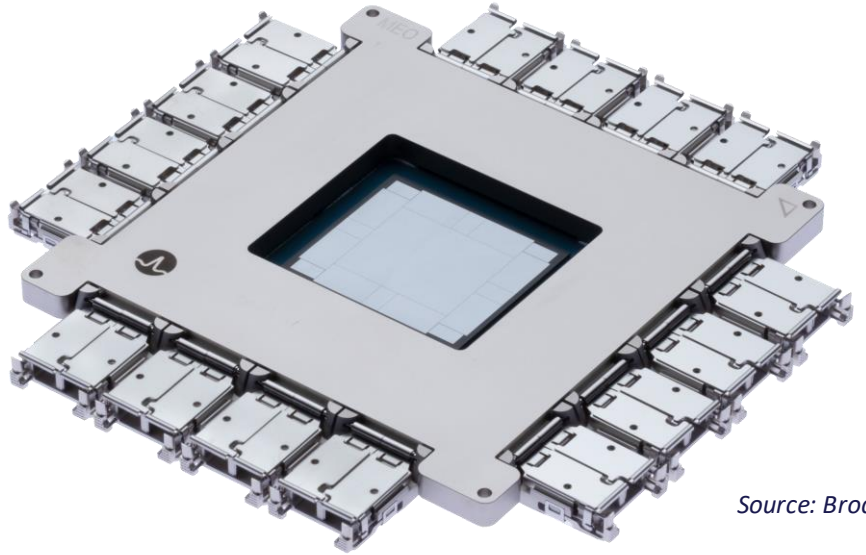
Power efficiency: expected power savings of up to 50% over the equivalent fully-retimed pluggable-based system (200G/lane generation)[2]

Cost reduction: theoretical 30%+ cost efficiencies possible (assuming volume deployments)[3]

NPO offers some attractive features for end users and switch manufacturers, but must be proven to deliver reliability and power efficiency on par with soldered CPO



NPO
Source: Ciena



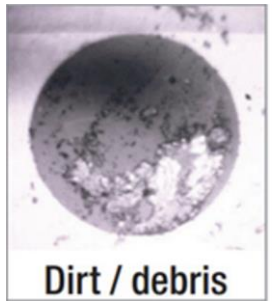
CPO
Source: Broadcom

1. *Evaluation of CPO performance and pluggable optics health for reliable AI infrastructure, S. Amiralizadeh (Meta), OCP 2025*
2. *https://developer.nvidia.com/blog/scaling-ai-factories-with-co-packaged-optics-for-better-power-efficiency*
3. *https://newsletter.semianalysis.com/p/co-packaged-optics-cpo-book-scaling*

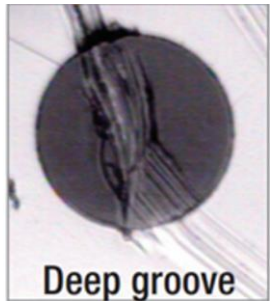
Dust, dirt, and other contaminants

Common failure modes

- Dirt, dust, debris – airborne contaminants and mishandling, particularly in areas with high clay content in soil
- Finger oils (sebum) – installer / technician mishandling
- Grooves / etching – permanent damage requiring replacement of cable or transceiver



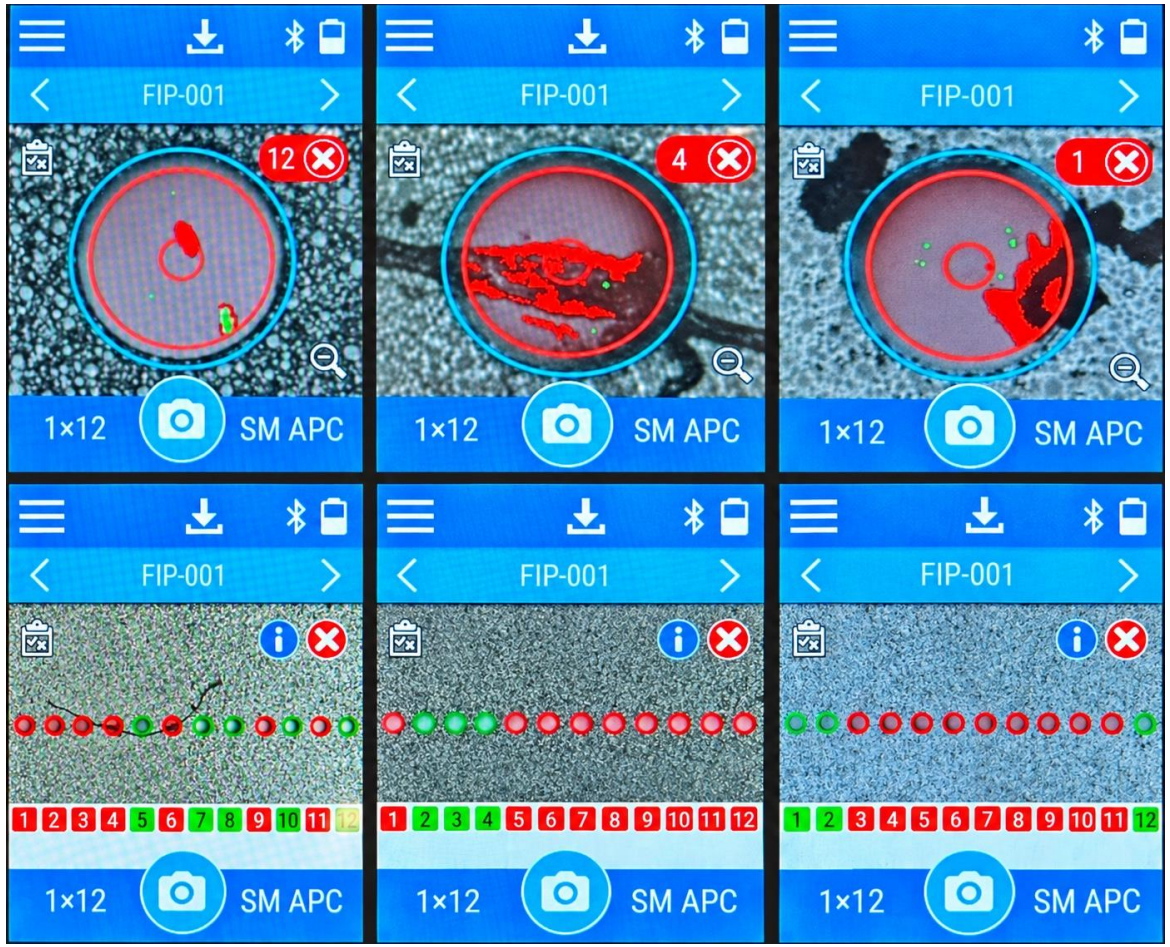




Some interesting stats from Oracle AI deployments

- 90% of observed L1 link failures are contamination related
- 97% of transceiver replacements are contamination

Sampling of observed contaminants in actual deployments

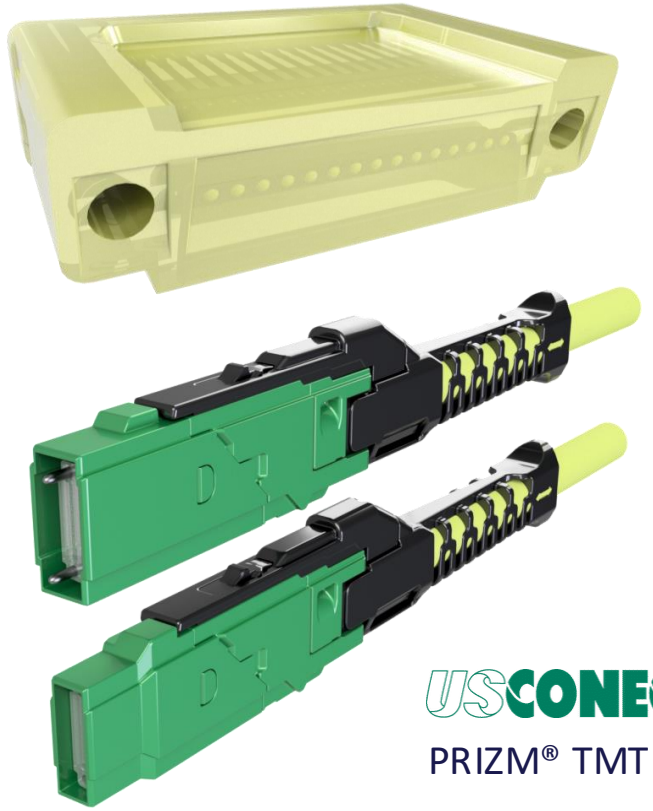


Expanded beam optical technology

Resilient connector design as an enabler of scale

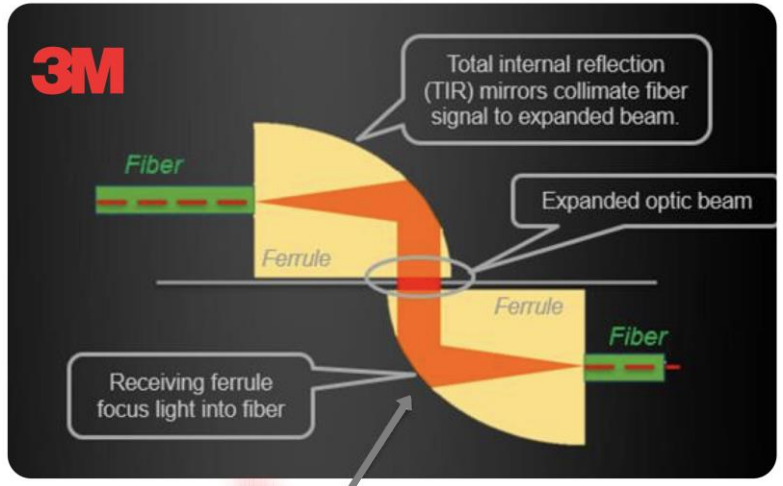
Expanded beam technology all but eliminates contamination as source of link down / flaps, and connector damage as a high-blast radius failure mode:

- optical beam expanded for dust, debris, oil resilience
- ferrule recessed to avoid direct contact with contaminants
- some designs feature 90° turn to contain optical path within housing
- no physical contact enables ~20x reduction in mating force, allowing large-format ganged insertion suitable for optical backplane applications

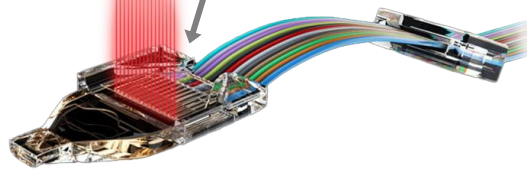


USCONEC

PRIZM® TMT



EBO ferrule



EBO MPO-style connectors



molex



SENKO® Advanced Components

EBOMSA

https://www.ebomsa.org/

Call to Action

- It will take a multi-pronged approach to tackle all aspects of TCO for AI infrastructure – network availability, power efficiency, and capital expenditure
- Linear drive interfaces are shown to improve on all of these metrics, paving the way forward for photonic integration (CPO/NPO)
- Innovations like expanded beam technology will contribute further but like CPO/NPO, require multi-vendor, interoperable ecosystems to become a reality
- Get involved!

  https://www.opencompute.org/

  https://www.oiforum.com/

  https://ebomsa.org

  https://opencpxmsa.org

Thank You!

mark.filer@oracle.com

