400G, 800G, and Terabit Pluggable Optics:

What You Need to Know

Mark Nowell
Cisco Fellow

Errol Roberts
Distinguished Architect

BRKOPT-2699

Webex App

Questions?

Use Webex App to chat with the speaker after the session

How

1. Find this session in the Cisco Events App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 27, 2026.

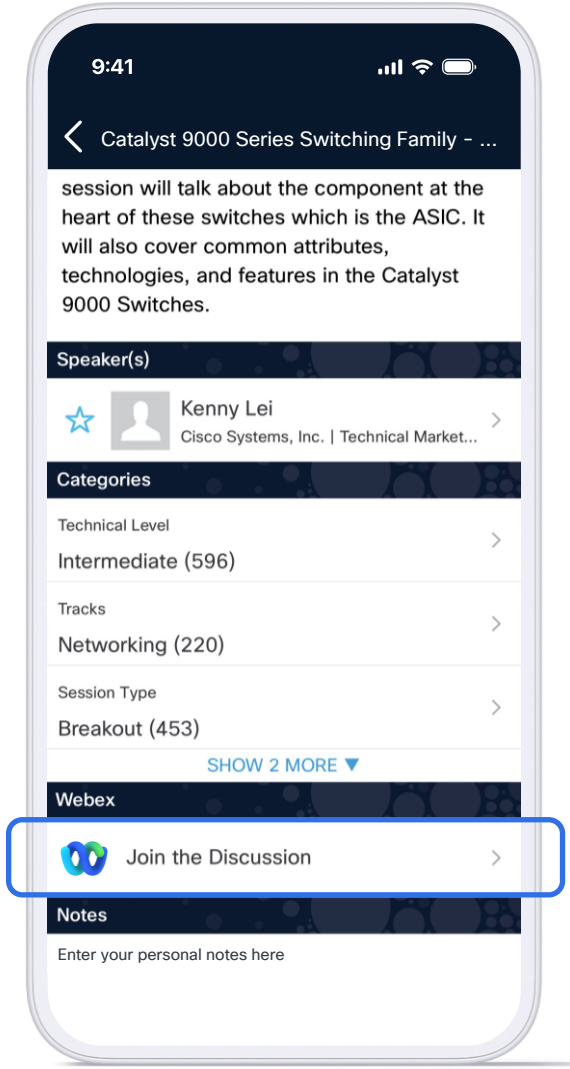

Are pluggable optics dead or alive for the AI era?

Pluggable optics at the dawn of the AI era

- The AI market is accelerating the growth of networking and compute technologies to keep up with the deployments
- Lots of announcements and news stories about the challenges of pluggables:
  - Too much power
  - Copper cables can’t keep up
  - Hope that integration of optics with the GPU or Switch ASICs (aka co-packaged optics) can solve all the problems
- Are pluggables relevant in the AI era?

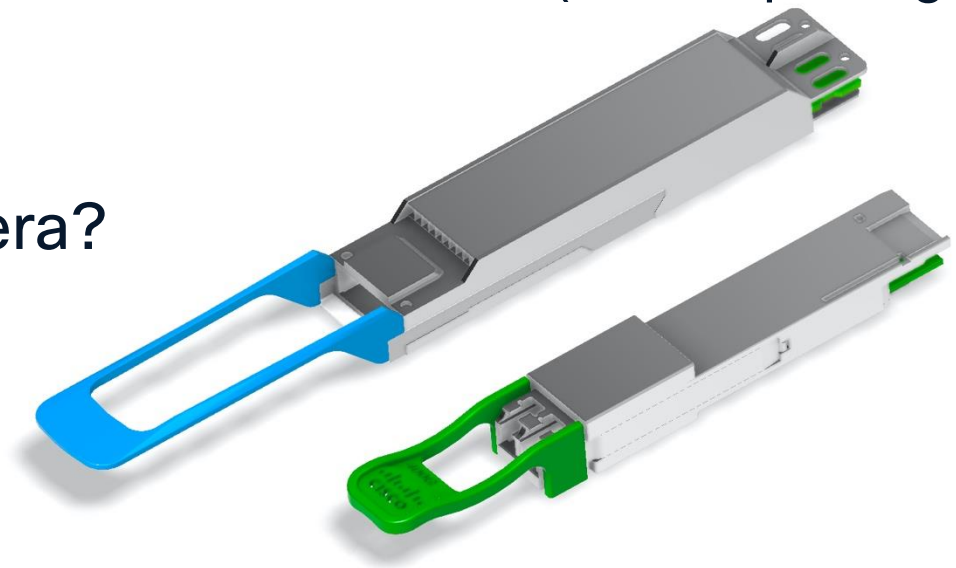

AI Infrastructure & Pluggable Optics

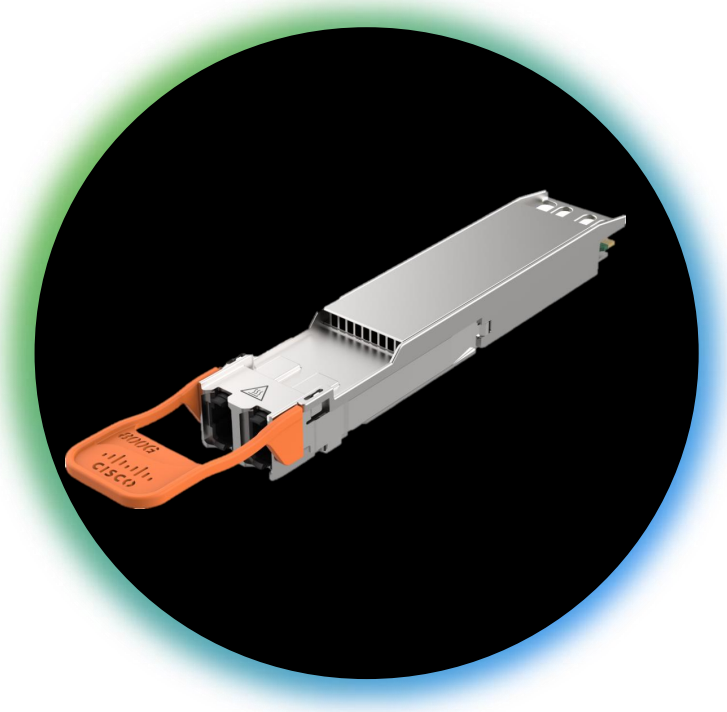

- AI infrastructure demands are a disruptive transformation
- AI is driving innovations in pluggable optics
- Demands higher – reliability, performance, integration
- Cisco leadership in SiPhotonics addressing AI infrastructure demands

Agenda

01 Market Dynamics

02 400G & 800G Pluggable IMDD and Coherent Status

03 Optics Advancements for AI Infrastructure

04 Wrap Up

Market dynamics

Ethernet Worldwide Switch Port Speed Transitions

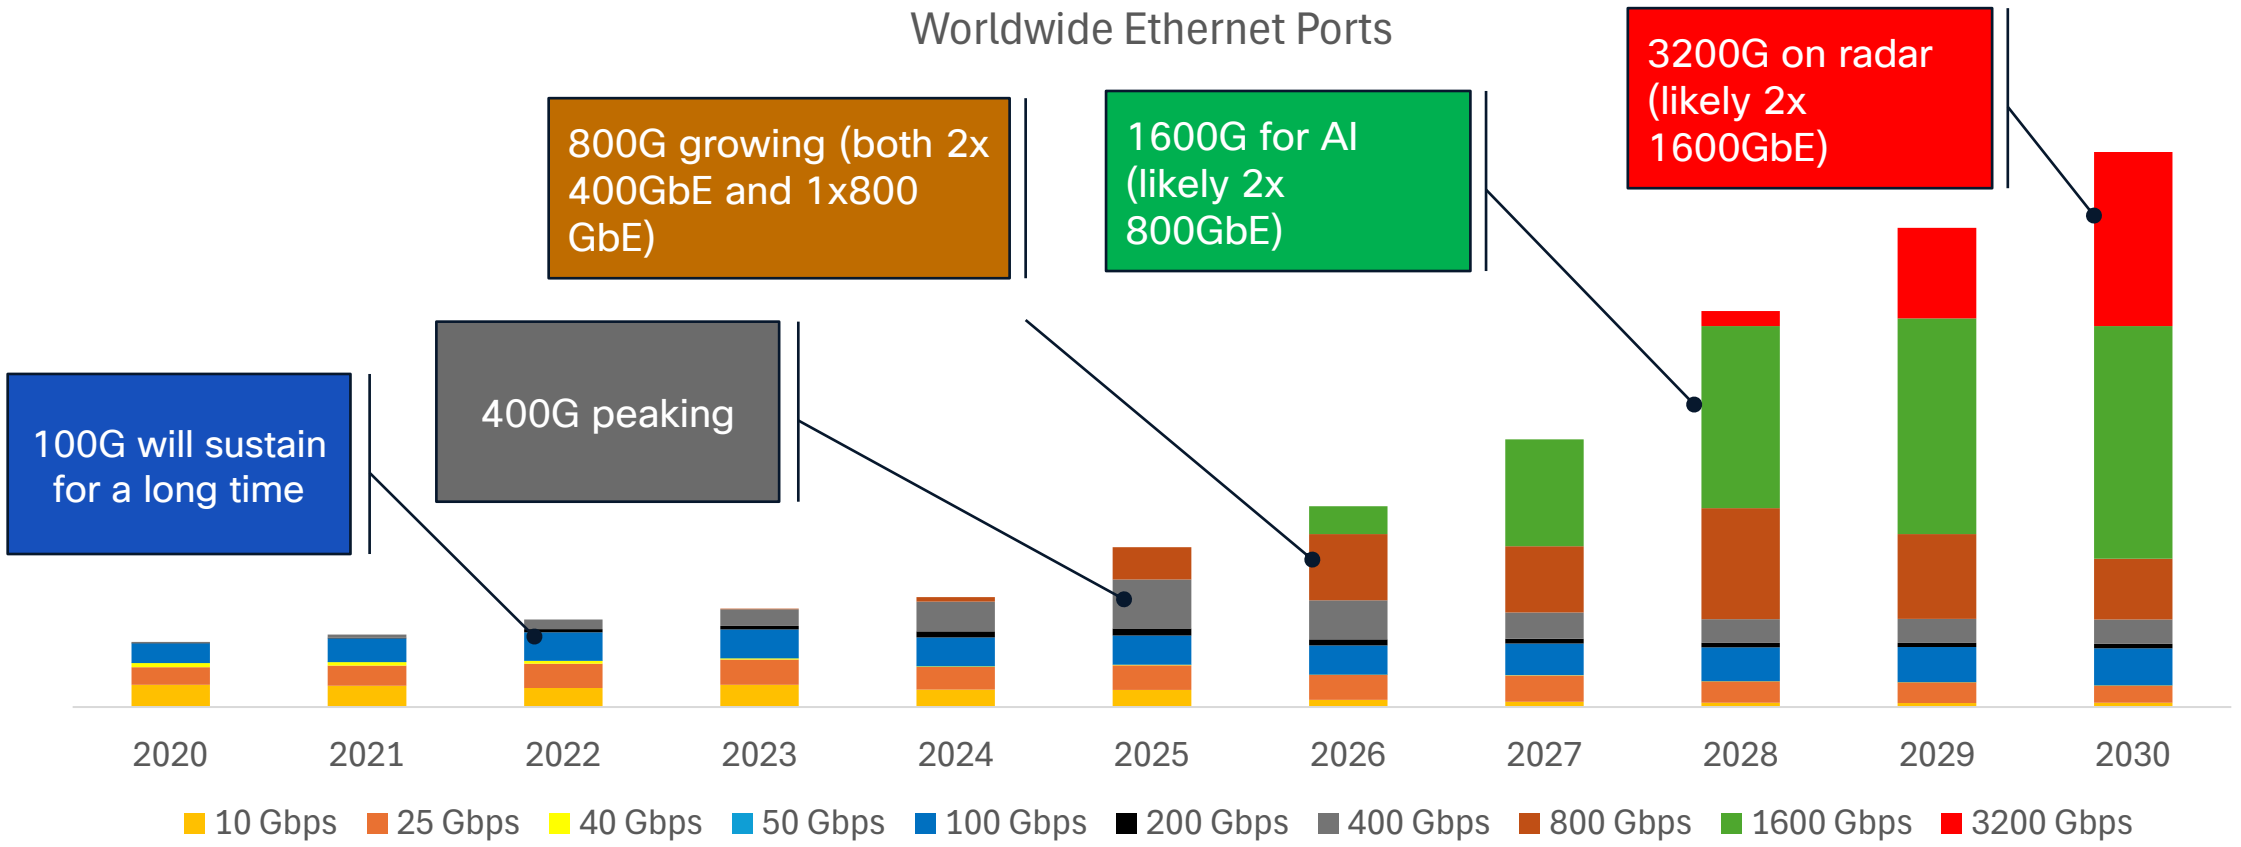

Source: Source: Dell'Oro, "Data Center Switch– Total 2026 – 2030"

- AI Infrastructure is dominating the market (Backend Fabric)
- Hyperscalers will drive these speed transitions quickly to scale
- Rest of the market will leverage that scale in their own time frame

Speed evolution in a data center (Front-end)

100% pluggable optics

Inter-Data Center

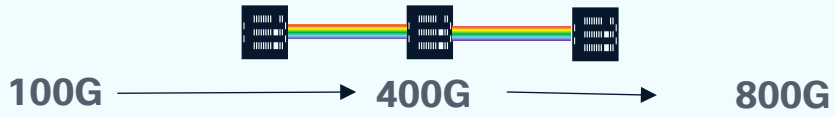

- Significant pluggable DCI (DWDM coherent)
- Open Line System

Switch fabric

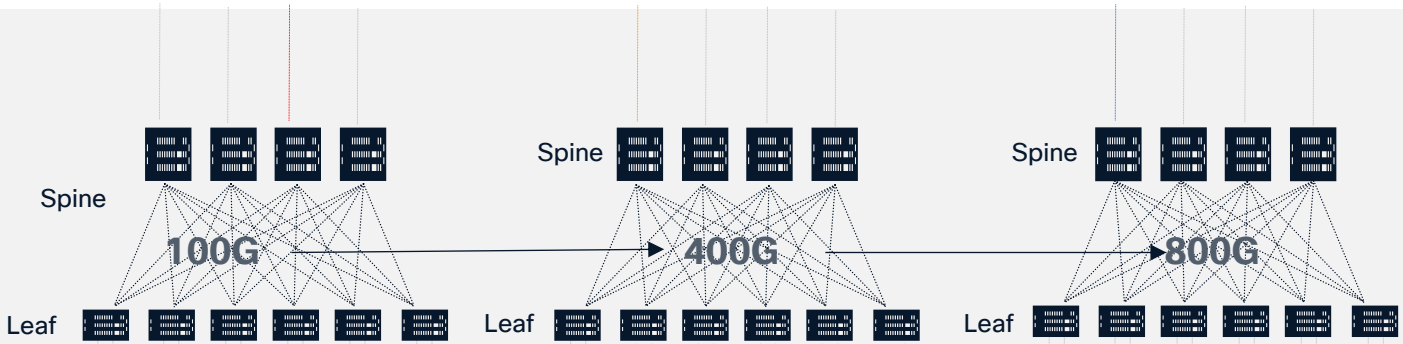

- Switch silicon bandwidth growing due to higher Radix and faster Serdes speeds
- Switch ASIC throughput growing: 6.4 Tbps to 12.8 Tbps to 25.6 Tbps to 51.2 to 102.4 Tbps
- Optics increasing from 100 Gbps to 400 Gbps to 800 Gbps

Access

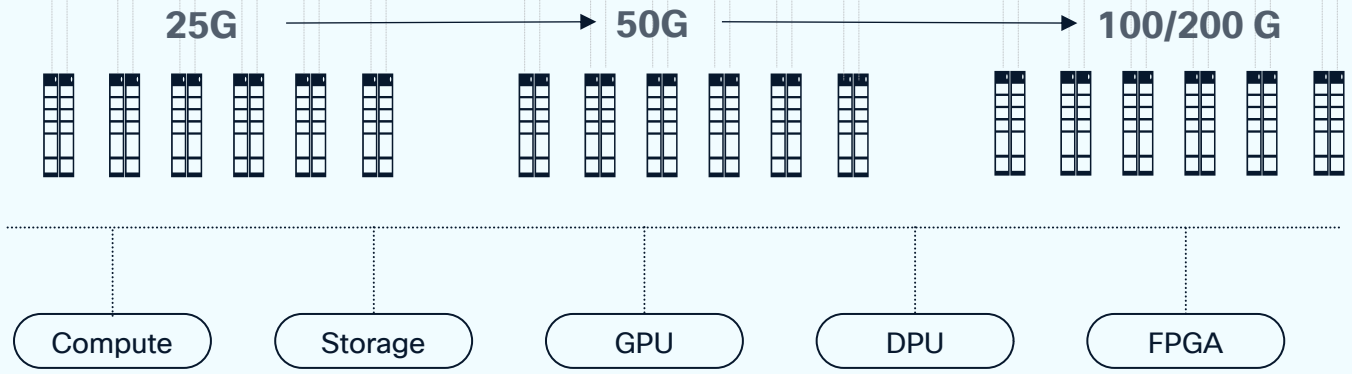

- Server network connectivity evolves with server processor upgrade cycles as data center traffic grows
- Server port speed is transitioning from 25 Gbps to 200 Gbps
- Storage, GPU, DPU, FPGA driving connectivity bandwidth, PCIe speed increase

Ethernet Speed Transitions in AI Networks

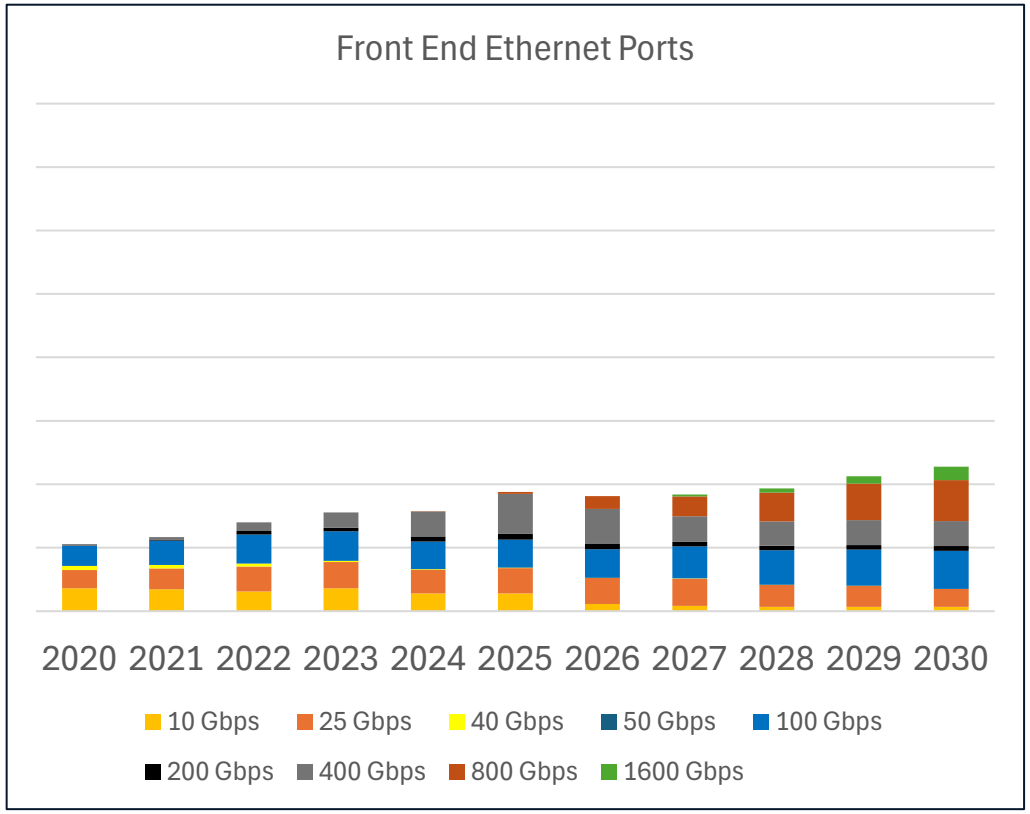

Source: Dell'Oro, "Data Center Switch – Front-end networks 2026 – 2030"

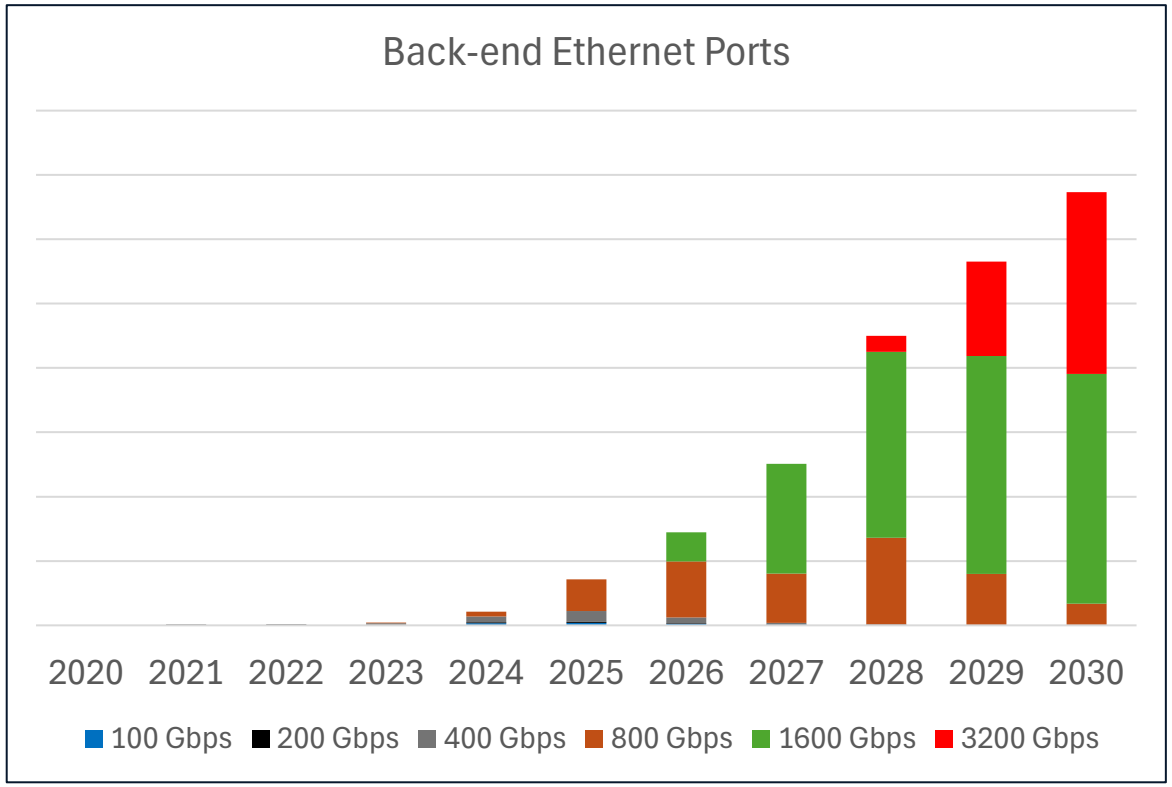

Source: Dell'Oro, "Data Center Switch – AI Back-end Networks 2026 – 2030"

Majority of the switch ports in AI back-end Networks will be 1600 Gbps in 2027, showing a very fast migration to the highest speeds available in the market.

AI Network Type Fundamentals

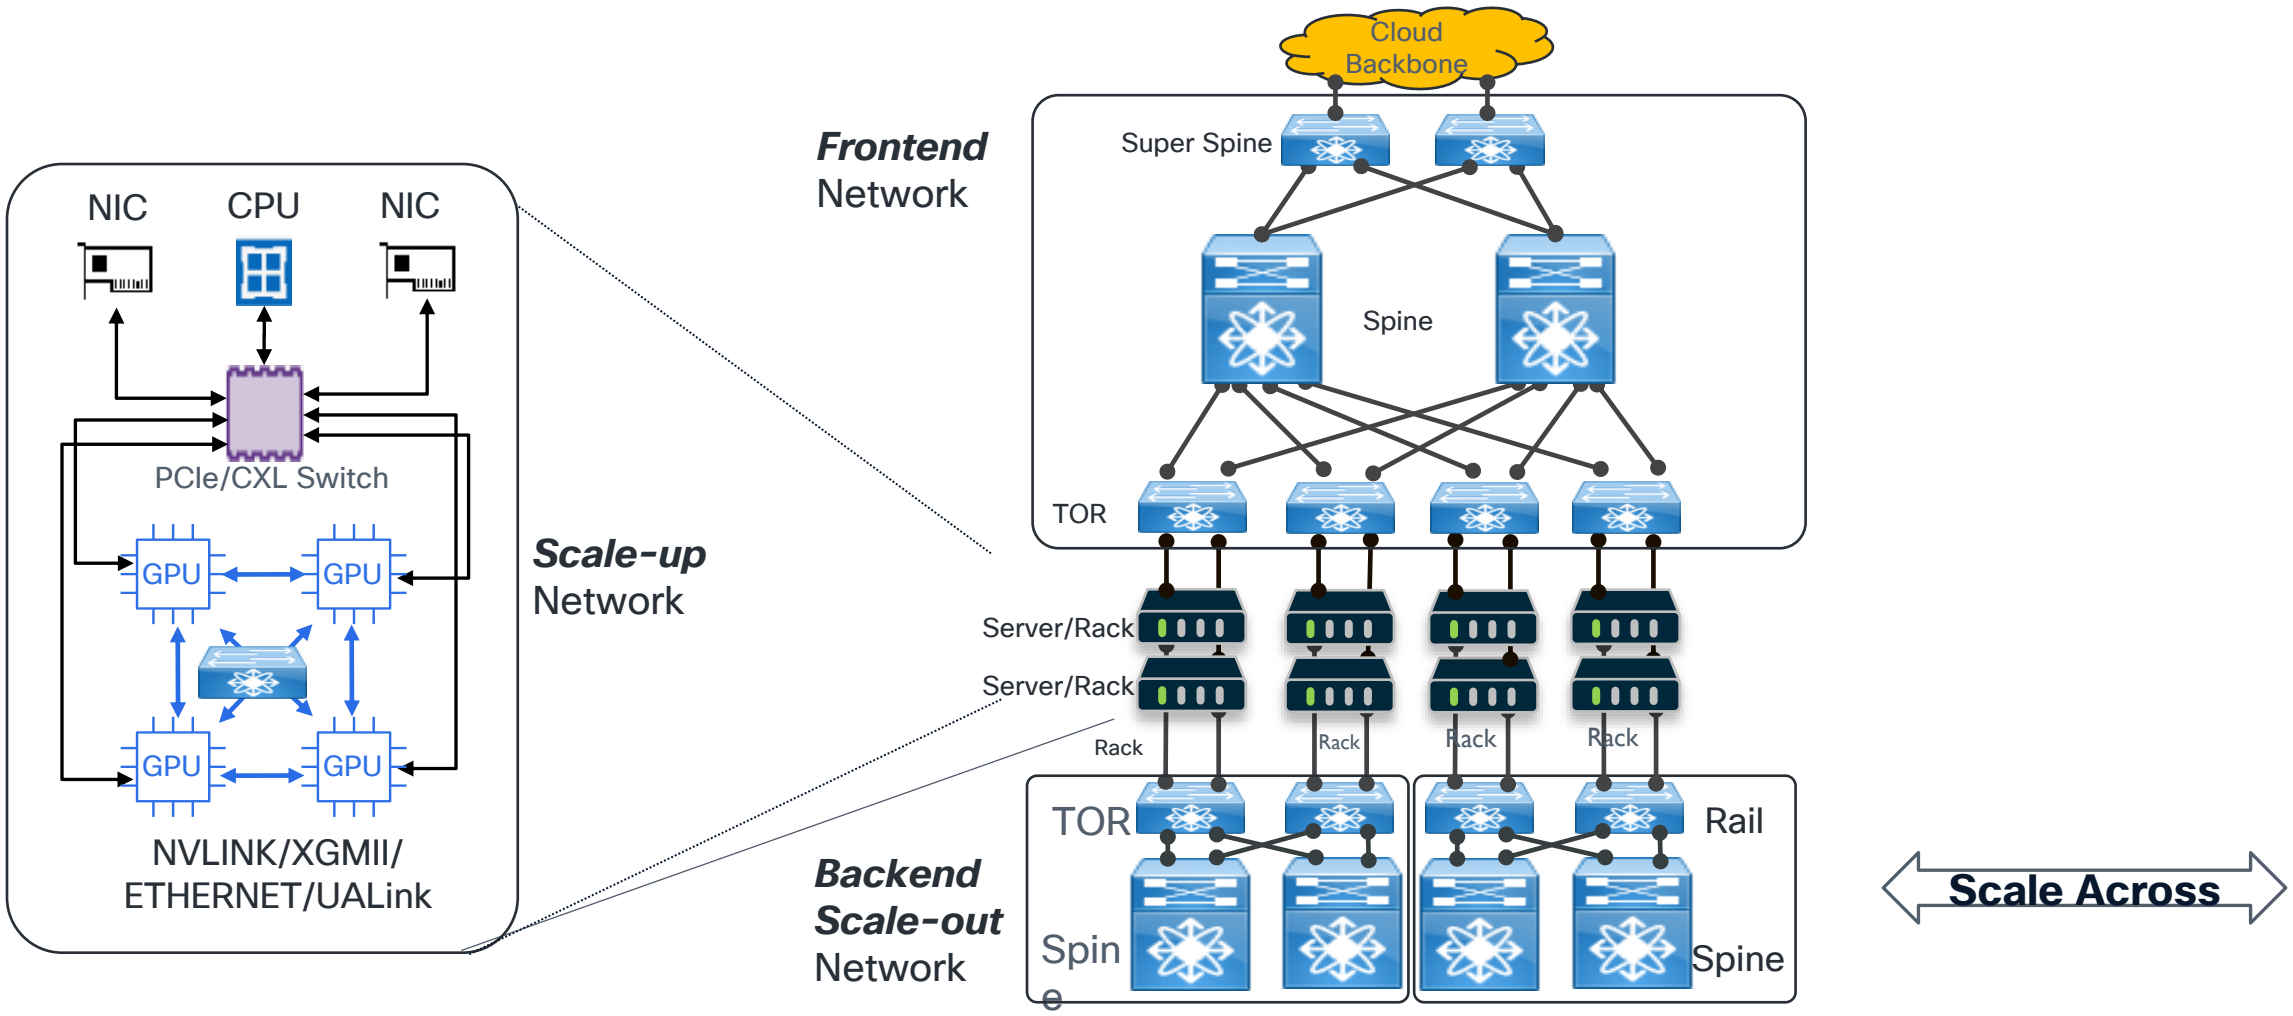

AI/ML is a disruptive event for equipment design

		Traditional Front-end DC	AI/ML Back-end DC
Rack Bandwidth (ToR/MoR)		3.2T-12.8T	>> 100T
Rack power		~10 kW	100 kW+
Packet Loss impact (reliability)		Low importance	Critical importance
Latency importance	Absolute	Low	Low
	Tail	-	High

These challenges are forcing innovation to happen at all levels, including pluggable modules. But pluggable modules still offer:

- Flexibility of media (cables, optics)
- Flexibility of supplier
- Ease of replacement
- Commonality of designs → robust module and switch ecosystem

400G/800G/1.6T use cases

Market Drivers

Cloud & NeoCloud

Earliest adopters on next speeds and variants.

High volume drives economies of scale and optimization

Telco service providers

Leverage early adopter solutions

Enterprise

Later adoption of speeds and feeds.

Strong leverage of previous solutions

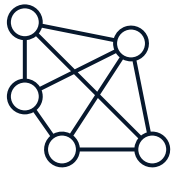

Media networks

Strong leverage of existing solutions

Continued Cisco Optics Innovation

Fiber investment protection
with 400G BiDi Optics

800G Silicon Photonics Optics Maximum
reliability and lowest power in AI infrastructure

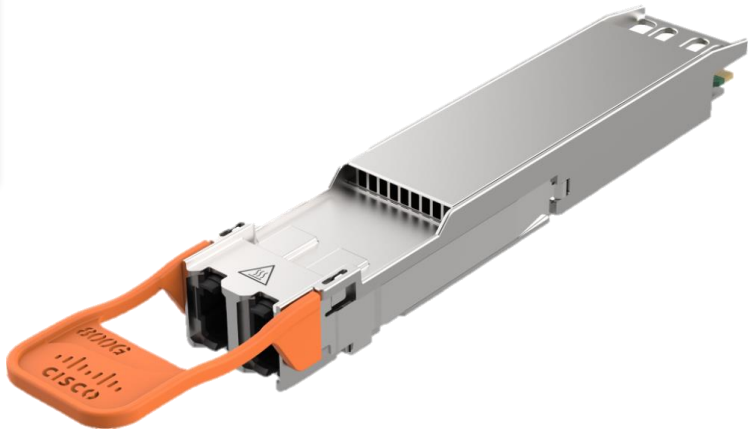

High speed and ultra-long haul
coherent pluggable optics

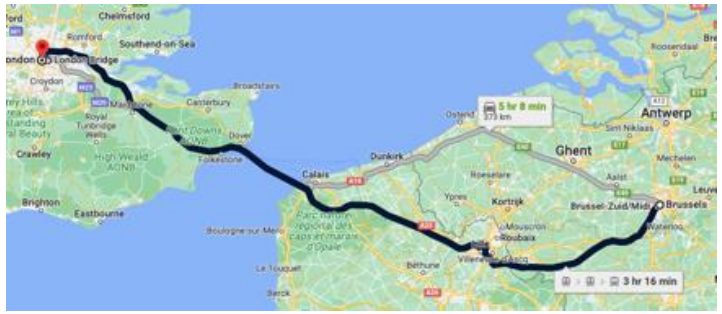

Optics firmware updates onsite
Leadership in standards development

Resilient Optics for AI → Reliability/Performance, Links always on, Telemetry

400G & 800G Pluggable IMDD and Coherent Status

400G & 800G: Where are we today?

Client IMDD (Shorter Reach)		Coherent (Long Reach)
400 GbE, 800 GbE	Mature standards	400G (OIF, Open ZR+)
100 Gb/s	Mature technology	400ZR: 16QAM @ 60 Gbd 800ZR: 16QAM @ 118 Gbd
1.6 TbE & 800 GbE @ 200 Gb/s	Next Gen technology	1600ZR: 16QAM @ 240 Gbd
Incl. Breakout	Mature Deployment	400ZR, Bright 400ZR+
CMIS	Optics Management	CMIS

Full range of 400G / 800G pluggable modules

- Copper cables
- Multimode Fiber – 100m
- Single Mode Fiber inside DC – 500m & 2km
- Single Mode Fiber Campus – 10 km
- Outside plant, DCI – 100 km → 3000+ km

Most 400G/800G Standards are done

✅ Complete

Cisco-led

Standards	IEEE 802.3bs ✓		400 GbE & 200 GbE MAC & Initial Interfaces
	IEEE 802.3cd ✓	cisco	50 GbE MAC & Interfaces (also 100 GbE & 200 GbE PMDs)
	IEEE 802.3cm✓		400 GbE MMF (BiDi and SR8)
	IEEE 802.3cn ✓		Extended reach (40km) 50 GbE, 200 GbE, 400 GbE
	IEEE 802.3ct ✓	cisco	100GbE Coherent 80km
	IEEE 802.3cu ✓	cisco	100G-FR, 100G-LR, 400G-FR4, 400G-LR4-6
	IEEE 802.3ck ✓	cisco	100GE serdes
	IEEE 802.3db ✓		100/200/400GE MMF (100Gb/s short wavelength)
	IEEE 802.3df✓	cisco	200G/400G/, 800G Ethernet Task Force @ 100Gb/s per lane
	OIF400ZR ✓		400 GbE Coherent 120km
	OIF800ZR		800 GbE Coherent 120km
	OIF1600ZR		800 GbE Coherent 120km
	Open ZR+ ✓	cisco	Long Distance 400G
	CMIS	cisco	Common Management Interface Specifications
	802.3ds		200 Gb/s per Wavelength MMF PHYs Task Force
	802.3dj	cisco	200G/400G/800G/1.6T Ethernet Task Force @ 200Gb/s per lane
	802.3dk		Greater than 50 Gb/s Bidirectional Optical Access PHYs Task Force.
MSAs*	LPO MSA ✓	cisco	100G Linear pluggable
	100G Lambda MSA ✓	cisco	100G-FR, 100G-LR, 400G-FR4, 400G-LR4
	QSFP-DD MSA ✓	cisco	400G/800G/1.6T Form factor, started CMIS
	OSFP MSA ✓	cisco	400G/800G/1.6T Form factor
	400G-BiDi MSA ✓		400 GbE MMF BiDi

Key elements of 400G & 800G optics

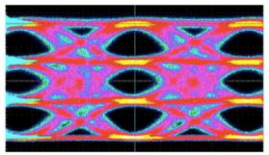

Pluggables with 8-wide connector to support 400G and 800G ports

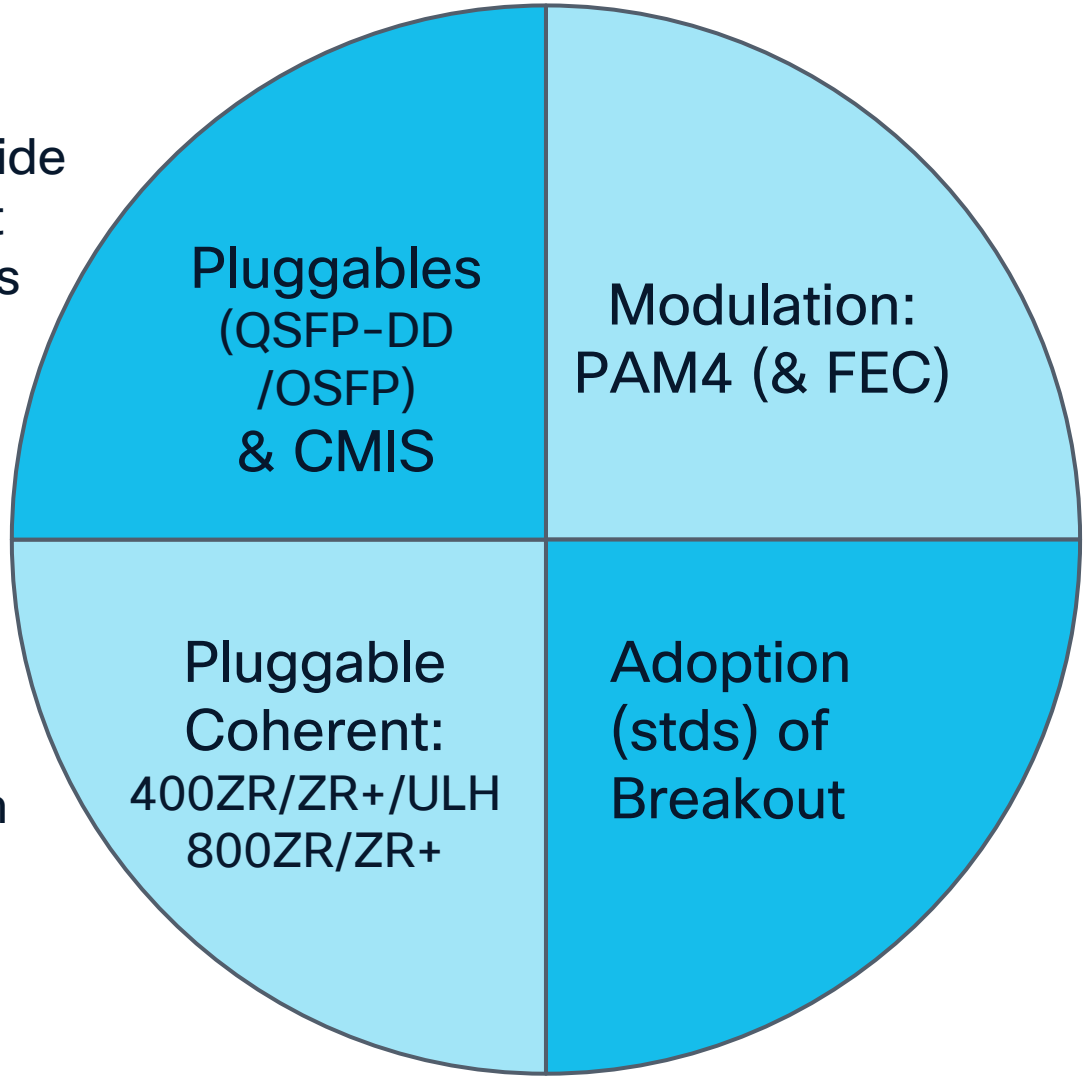

Higher speed interfaces PAM4 modulation.

Ubiquitous use of FEC.

Long reach coherent without any system port density reduction → Routed Optical Networking

Pluggable modules supporting multiple lower speed interfaces

400 GbE / 800 GbE – Uses mature 100 Gb/s technology

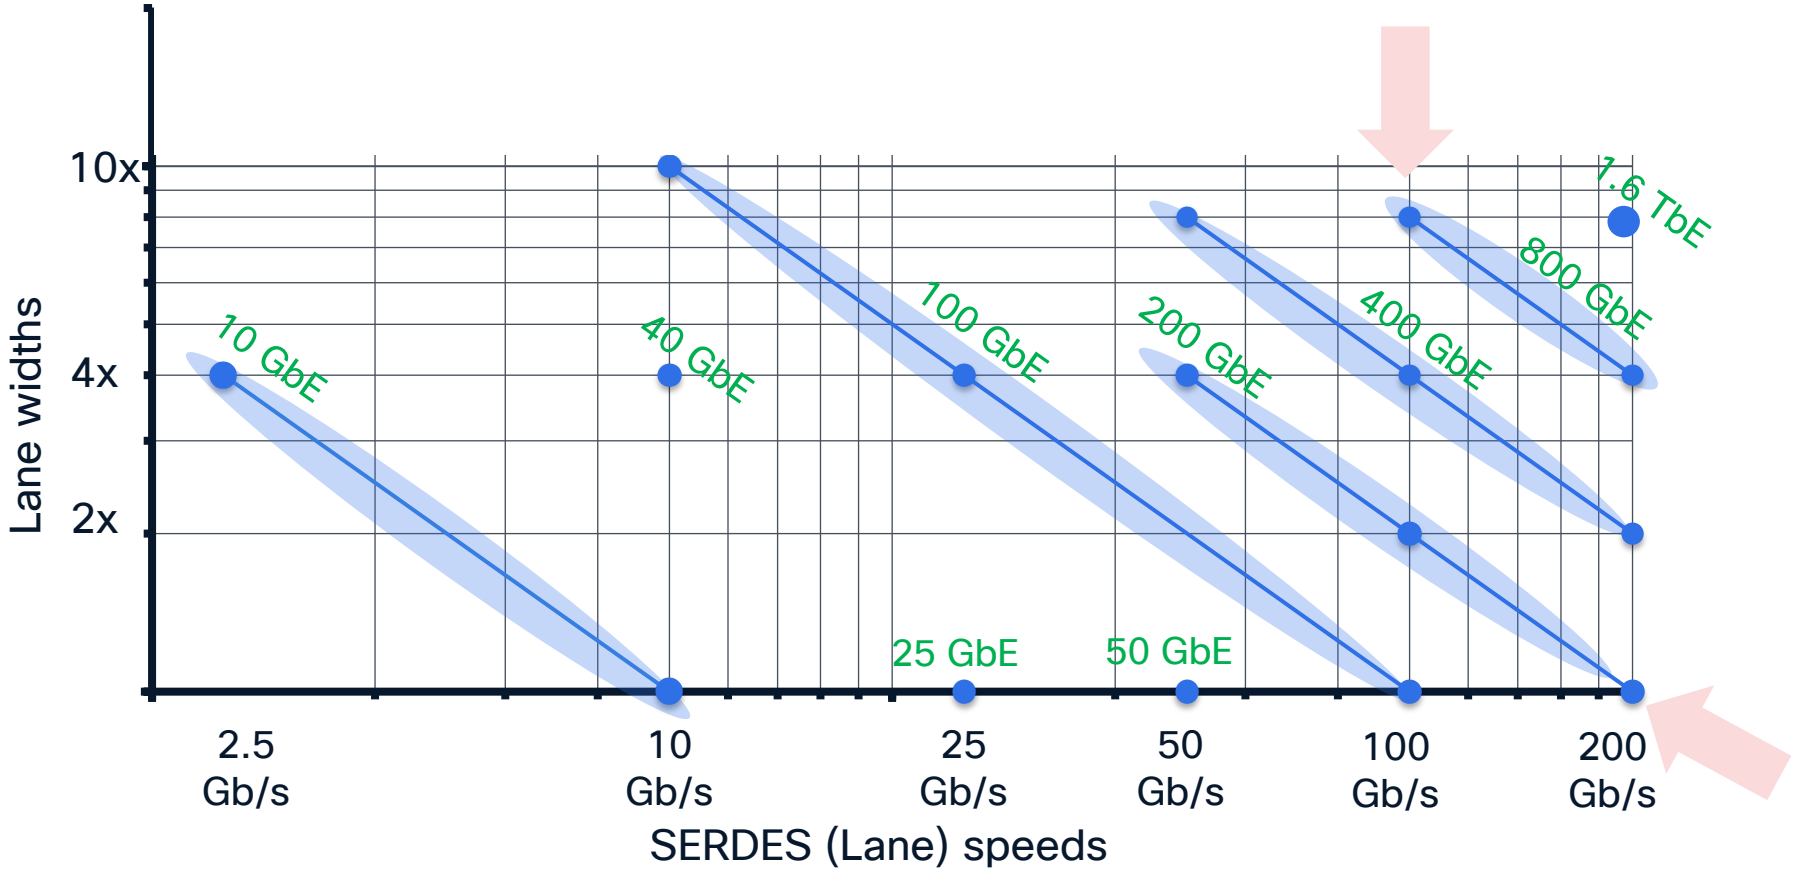

- (almost) nothing is new or only used once
- 8x100 “today”
- nx200 “tomorrow”
- Requires (yet more) integration
- (yet) higher mechanical and thermal densities

Newer 200Gb/s technology in development now.

Technology maturity of 100 Gb/s technology enables:

- Excellent performance
- Excellent interoperability
- Broad adoption

Technologies used for 400G, 800G optics

Advanced Modulation

Client optics (IMDD)

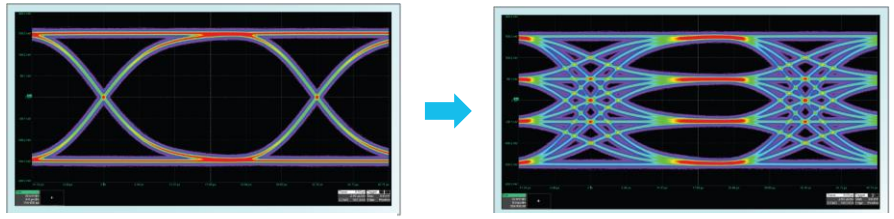

NRZ
1 bit/sec/symbol

PAM4
2 bit/sec/symbol

Coherent optics

Coherent QPSK
2 bit/sec/symbol

Coherent 16-QAM
4 bit/sec/symbol

Using more complex modulation allows us to increase the data rate (Gb/s) without increasing the signaling speed (Gbaud)

Forward Error Correction (FEC)

Allows correction of errors at receiver

Enables use of relaxed specs (saves $) to get same performance or enables much higher performance.

Usually embedded in Ethernet switch or module DSP ASIC

Some overhead (faster data rate needed)

Common Management Interface Spec (CMIS)

Standardized now in OIF → 100G modules, ubiquitous adoption for 400G modules and above

Host (switch/router)

High speed electrical interface (OIF CEI, IEEE, etc)

Pluggable Module

Optical/Copper (IEEE, MSAs, OIF)

Management

- configure
- initialize
- monitor
- manage

Often overlooked, CMIS standardizes management between modules and hosts
Consistent definition for configuration, initialization, monitoring, telemetry, firmware update etc.
Critical for monitoring and telemetry – more hooks into the DSP for quality and performance, critical for AI
Versatile Diagnostics Monitoring – observables for link performance e.g. pre FEC BER, SNR...

Virtual Diagnostic Monitoring

Including optical module firmware upgrade in the field

Virtual diagnostic monitoring (VDM)

Real-time insights into performance and health across thousands of deployed pluggable optics

Measure pre-FEC and post-FEC performance for every transceiver

Troubleshoot and root cause analysis

Optimize network performance and efficiency

Proactive fault detection and diagnosis

Versatile Diagnostics Monitoring (VDM) features

OIF

Implementation Agreement (IA)

Common Management Interface Specification (CMIS)
Revision 5.3

- Observables include real time samples and threshold

VDM observables supported by Cisco	Benefit of VDM monitor
Laser Age	Track module lifetime and help with estimate/plan field replacement
Laser Temperature	Monitor laser temperature along with module temperature for optimizing system fan speed
SNR Media Input	aid in telemetry for degraded optical signal from remote Tx
SNR Host Input	aid in telemetry for degraded electrical input from host
PAM4 LTP Media Input	aid in telemetry for noise characterization at optical input
PAM4 LTP Host Input	aid in telemetry for noise characterization at electrical input
Pre FEC BER Media Input	**pre-FEC BER** monitor to diagnose faulty/degraded remote transmitter
Pre FEC BER Host Input	**pre-FEC BER** monitor to diagnose faulty/degraded host electrical input signal
FERC Media Input	uncorrectable **post-FEC errors** at optical input – link health telemetry
FERC Host Input	uncorrectable **post-FEC errors** at electrical input – link health telemetry
SEWmax Media Input	FEC statistics monitor (optical input) for maximum FEC codeword error bin count
SEWmax Host Input	FEC statistics monitor (electrical input) for maximum FEC codeword error bin count

Pluggable Optical Modules Options

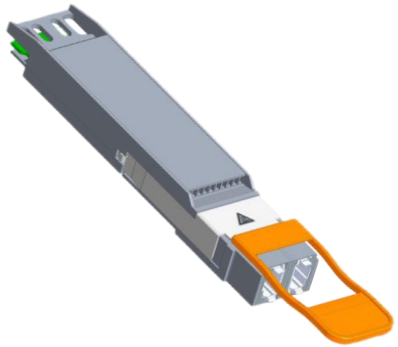
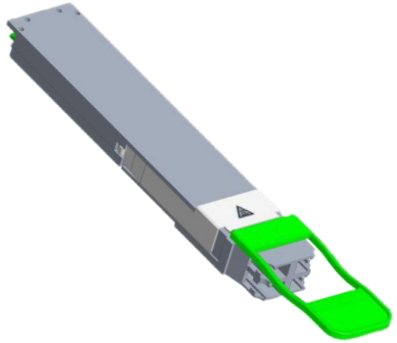
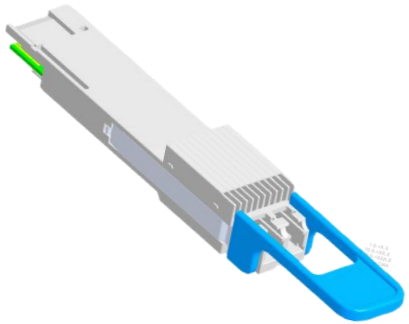
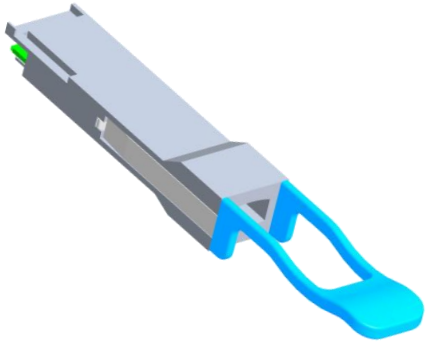

	OSFP	OSFP-RHS	QSFP-DD	QSFP
Host high-speed pins	8	8	8	4
Module capacity (Gbit)	400, 800, 1600	400, 800, 1600	400, 800, 1600	40, 100, 200, 400
Cooling	Integrated heatsink	Riding heatsink	Riding heatsink	Riding heatsink
High-speed IO Compatibility (Gbps)	50-100-200	50-100-200	50-100-200	1–25-50-100
Port compatibility	OSFP	OSFP-RHS	QSFP-DD & QSFP	
Management	CMIS			

System and ASIC design dictate which modules are supported and speed backwards compatibility

Pluggable form factors – a wealth of options

OSFP

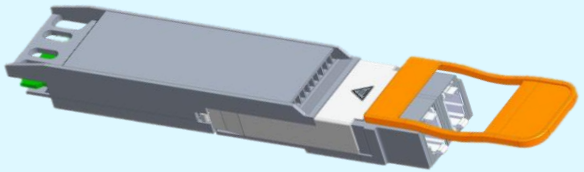

400G/800G/1.6T

OSFP-RHS

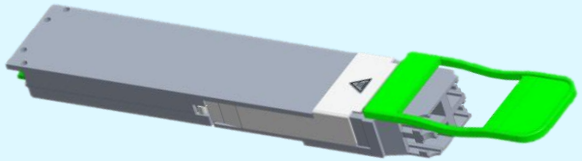

400G/800G/1.6T

Switches

Compute NICs

QSFP-DD

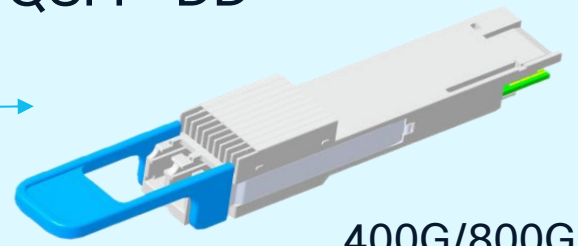

400G/800G/1.6T

QSFP

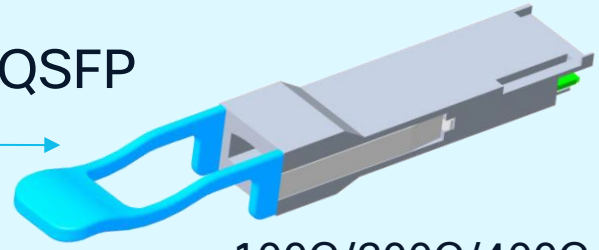

100G/200G/400G
QSFP28, QSFP56, QSFP112

Port compatibility

Everything within the square is compatible with each other

Pluggable form factors – interface interoperability

OSFP

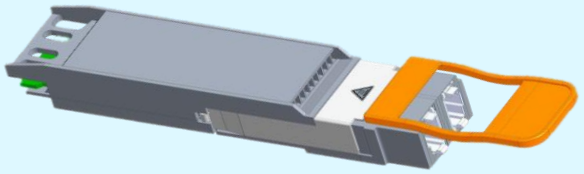

400G/800G/1.6T

OSFP-RHS

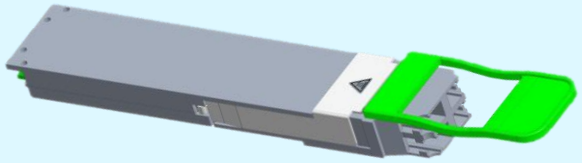

400G/800G/1.6T

Ethernet plug and play as normal.

Breakout as well!

QSFP-DD

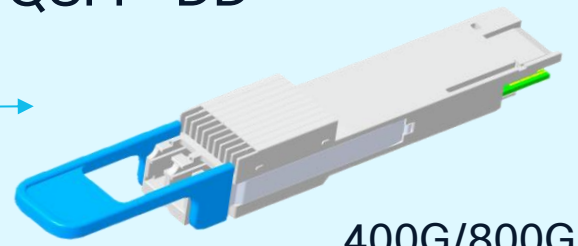

400G/800G/1.6T

QSFP

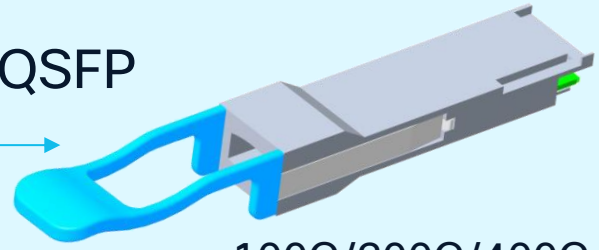

100G/200G/400G
QSFP28, QSFP56, QSFP112

Ethernet Interface compatibility

Port compatibility
Everything within the square is compatible with each other

Current trend: 800G Pluggables supporting dense 400 GbE

Both 400G & 800G form factor enables an economical way to implement breakout to lower speed Ethernet interfaces. This maximizes

- Cost and density on one end of link
- Compatibility and return on investment with existing equipment

Example: 800G module

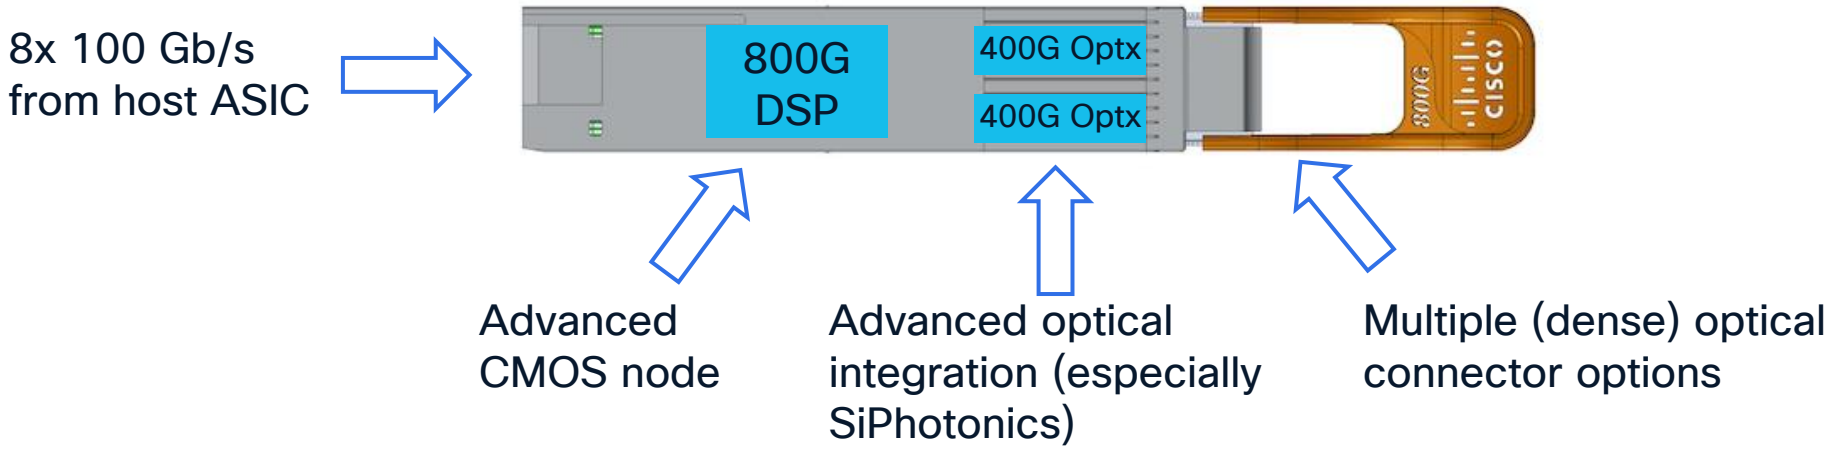

Implementing Dense 400 / 800 GbE

Pluggable modules allow wide range of deployment options between switch and compute equipment with compatible modules

Examples based on QSFP/QSFP-DD

100 Gb/s
25 Gb/s SerDes
QSFP28

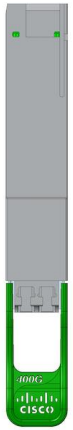

400 Gb/s
50 Gb/s SerDes
QSFP56-DD

400 Gb/s
100 Gb/s SerDes
QSFP112

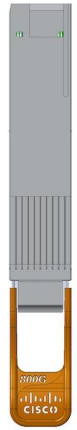

800 Gb/s
100 Gb/s SerDes
QSFP-DD800

100 GbE	1x	4x	4x	8x
200 GbE	-	2x	2x	4x
400 GbE	-	1x	1x	2x
800 GbE	-	-	-	1x

Broad Cisco availability for 200/400/800G

Distance	1–3+ m*	50/100 m	500 m–2 km	10 km	100+ km
200G QSFP56	QSFP-200-CUxM QDD-2Q200-CUxM	QSFP-200G-SL4 QSFP-200G-SR4-S	QSFP-200G-FR4-S		
400G QSFP-DD	QDD-400-CUxM QDD-2Q200-CUxM QDD-4ZQ100-CUxM QDD-400-AOCxM	QDD-400G-SR8-S QDD-400G-SR4.2-BD QDD-400G-VR4 QDD-400G-BD	QDD-400G-DR4-S QDD-4X100G-FR-S QDD-400G-FR4-S	QDD-400G-LR4-S QDD-2X100-LR4-S	QDD-400G-ZR QDD-400G-ZRP
400G QSFP112		QSFP-400G-VR4	QSFP-400G-DR4 QSFP-400G-DR4-2 QSFP-400G-FR4		
800G QSFP-DD	QDD-800-CUxM QDD-2Q400-CUxM QDD-4Q200-CUxM	QDD-800G-VR8	QDD-800G-DR8 QDD-8X100G-FR QDD-2X400G-FR4	QDD-800G-DR8-10 QDD-2X400G-LR4	
800G OSFP/R	OSFP-800-CUxM OSFP-2Q400-CUxM OSFP-4Q200-CUxM	OSFP-800G-VR8 OSFP-800G-VR8P OSFPR-800G-VR8	OSFP-800G-DR8 OSFP-800G-DR8P OSFP-L800G-DR8P OSFP-800G-DR8-2 OSFP-2X400G-FR4 OSFPR-800G-DR8	OSFP-800G-DR8-10 OSFP-2X400G-LR4	

New Form Factors

https://tmgmatrix.cisco.com/iop

https://optsel.cisco.com

Introducing the 400G BiDi Pluggable Optic

What's special about it?

It's the next generation of the popular 40G and 100G BiDi optics deployed widely in data centers.

- Breakthrough optical technology enabling 400 Gigabit link data rates over duplex multimode fiber (MMF).
- Cisco proprietary DSP design with advance equalization.
- The industry's only solution for 400G over duplex MMF today.

Increase the capacity of installed multimode fiber

Introducing Cisco 400G BiDi for duplex fiber

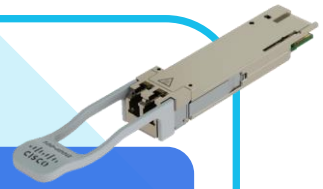

1000BASE-SX

10GBASE-SR

40G-BiDi

100G-BiDi

400G-BiDi

IEEE Standard

Cisco duplex MMF technology

Install more fiber

With 400G BiDi, you don't have to install more fiber

40G-SR4

100G-SR4

400G-SR4.2

400G-SR8

IEEE Standard (parallel MMF)

Increase the capacity of your trunk cable

8x more 400G links with the same number of trunk cable fiber pairs

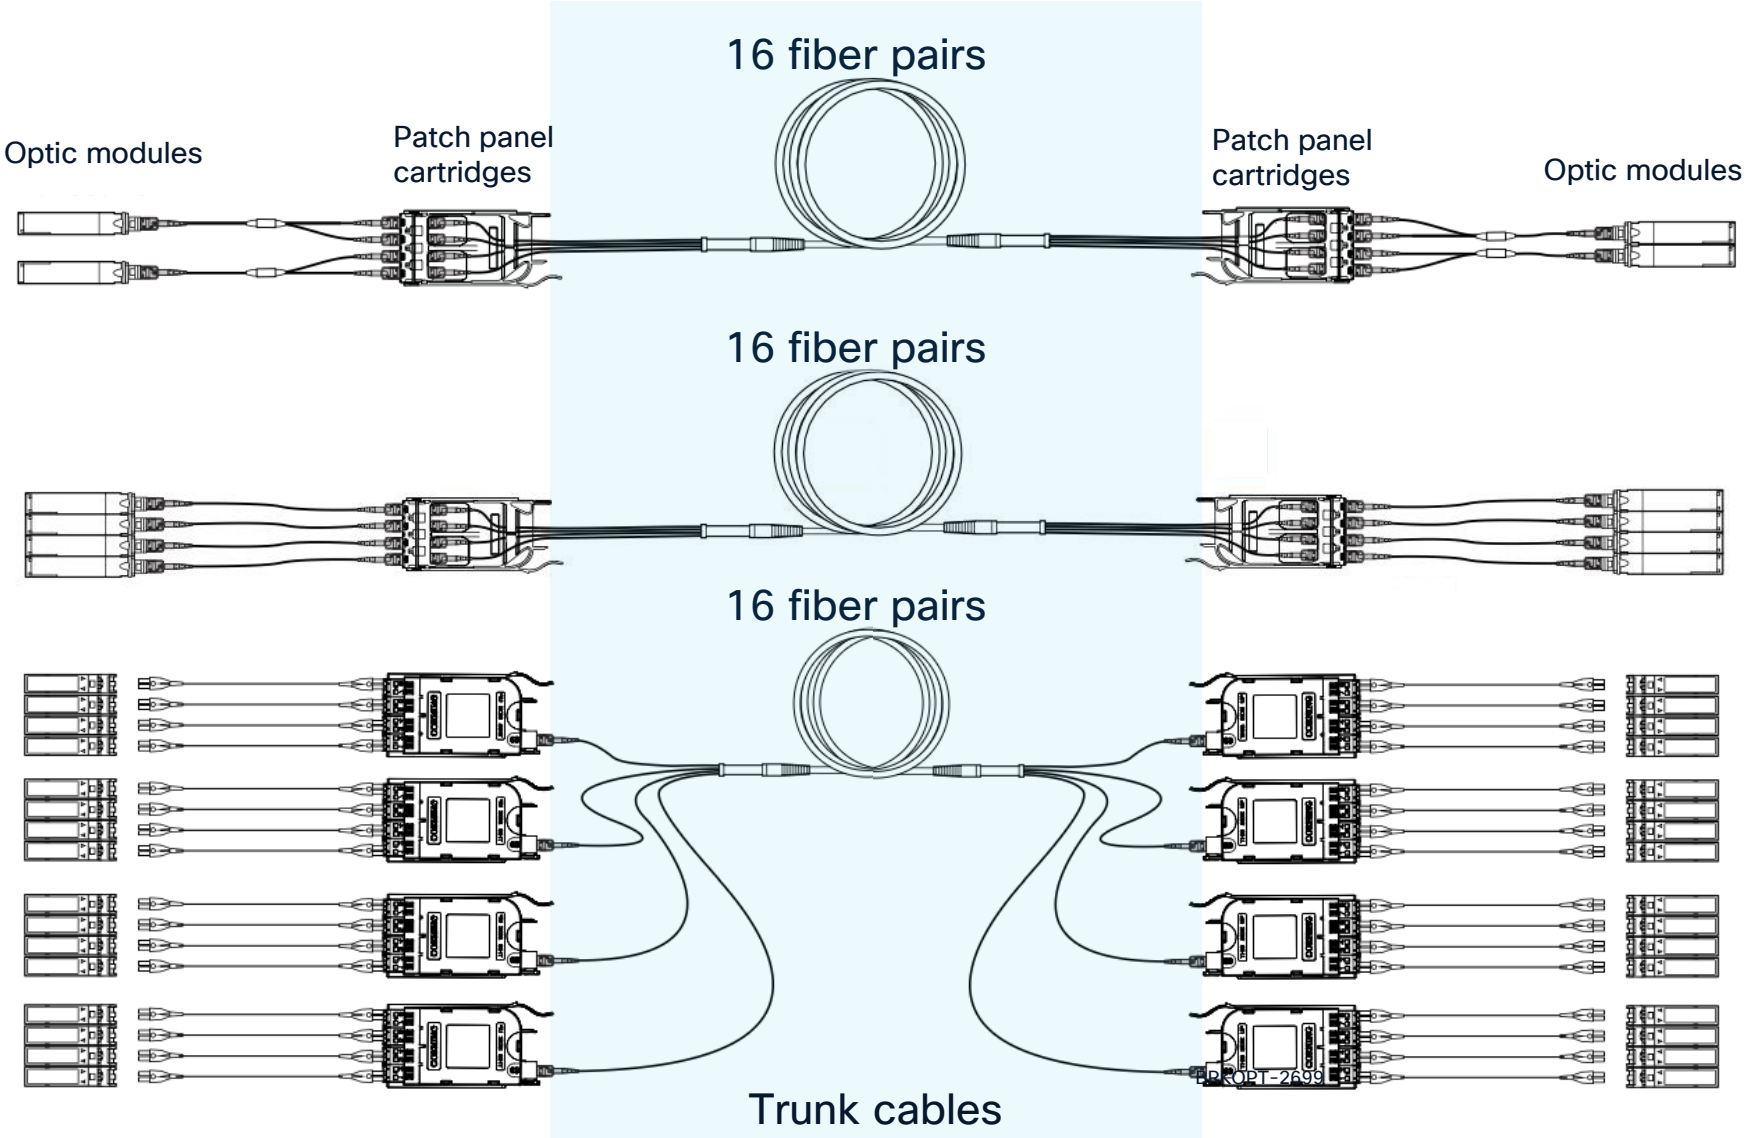

	Bandwidth increase
400G SR8: 2 links (total bandwidth 800Gbps)	1x
400G SR4.2: 4 links (total bandwidth 1.6Tbps)	2x
400G BiDi: 16 links (total bandwidth 6.4Tbps)	8x

Dual-rate and interoperability with 100G SR1.2

400G Link

- Only 1 fiber pair
- 4 wavelengths
- Duplex LC connectors

100G Link

Useful for partial upgrades of network

- Only 1 fiber pair
- 2 wavelengths
- Duplex LC connectors

Investment protection with 400G BiDi optics

Brownfield deployments: save 65% per link by re-using fiber infrastructure

- No new fiber on trunks
- Re-use patch panels
- Re-use patch cords
- No network disruption or new fiber install

Green-field deployments: use 8 times less fiber compared to IEEE 400G SR8 optics

- Less expensive patch panel cassettes
- Less expensive patch cords

Fiber investments for duplex fiber and parallel fiber

Fiber infrastructure	Duplex MMF	Parallel MMF
400G optic	Cisco 400G BiDi	IEEE 400G SR8
Fiber trunks	2 fibers per link	16 fibers per link
Patch panels	Duplex LC connectors	MPO16 connectors
Patch cords	2 fibers and duplex LC connectors	16 fibers and MPO16 connectors

Key takeaways – client (IMDD) reaches

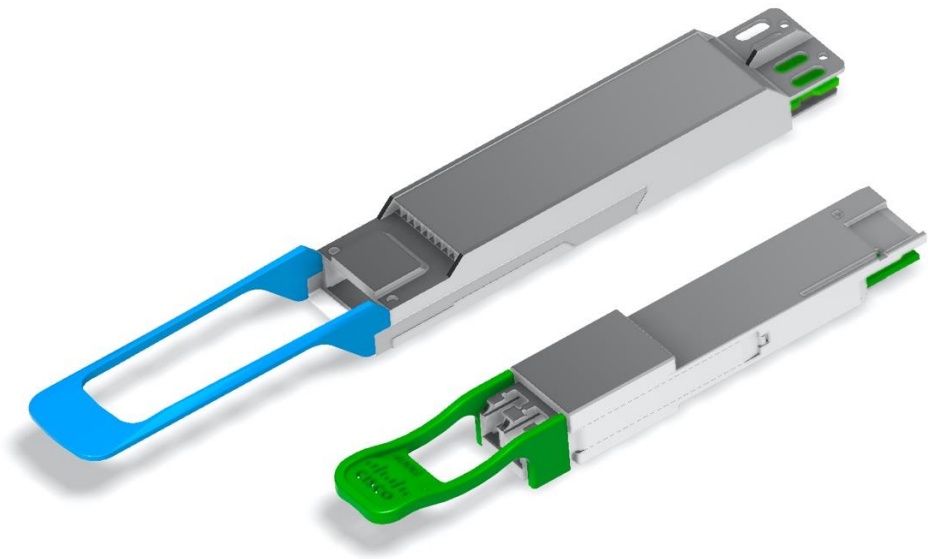

400G & 800G Standards are mature. Ethernet interop is guaranteed regardless of form factor

Optics and copper interconnect available in all reaches. **New** 400G MMF BiDi.

Cisco has a wide variety of high-density transceiver and cabling breakout solutions

Simplified deployment options are available

Coherent optics addresses higher reaches (and speeds)

- IMDD and Coherent technology will continue to be used
  - 800G, 1.6T and beyond
  - Coherent pushing towards shorter reaches. Not only in DWDM but also Grey applications
  - Focus shifting from performance enhancements ($$$) to interoperable interfaces and pluggables
- Coherent multivendor Interop more and more prevalent
  - 100ZR, 400ZR/ZR+, 800ZR/ZR+, 1600ZR

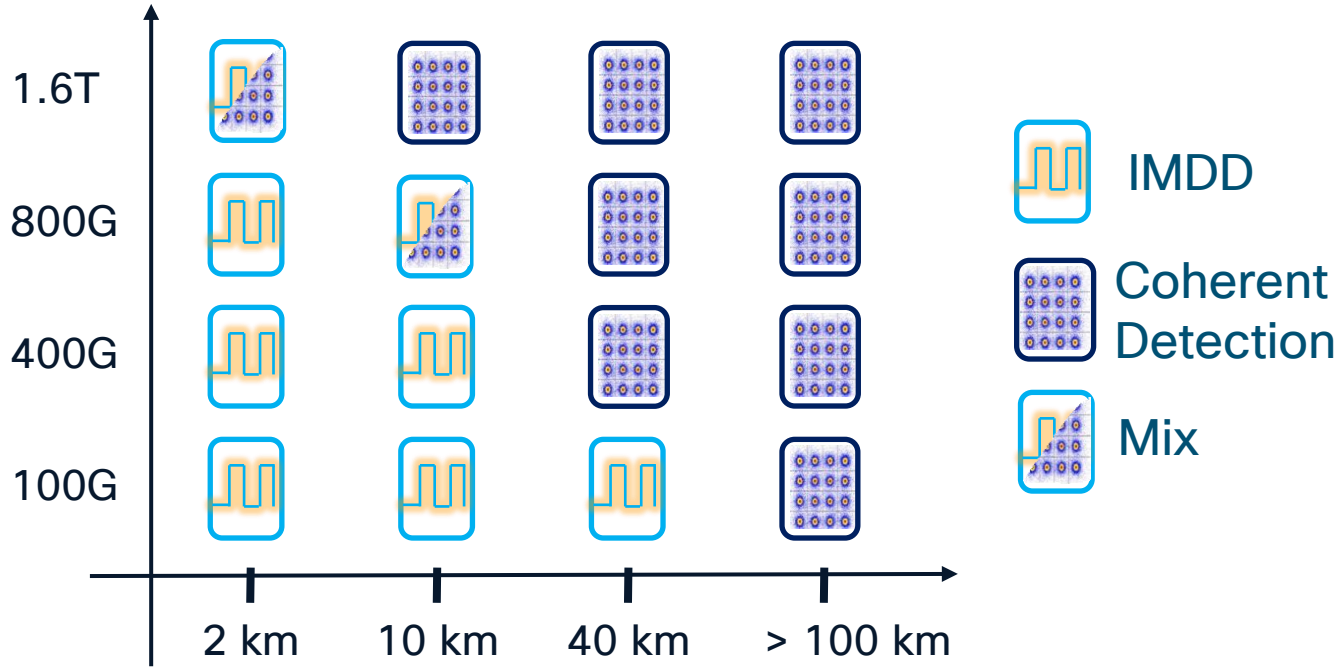

As speed increases, the crossover distance between IMDD & coherent decreases

Market/Technology Trends – Coherent Pluggables

Coherent Pluggables % Continue Increasing

- Today coherent ports are dominated by pluggables
- IP-over-DWDM gaining widespread adoption

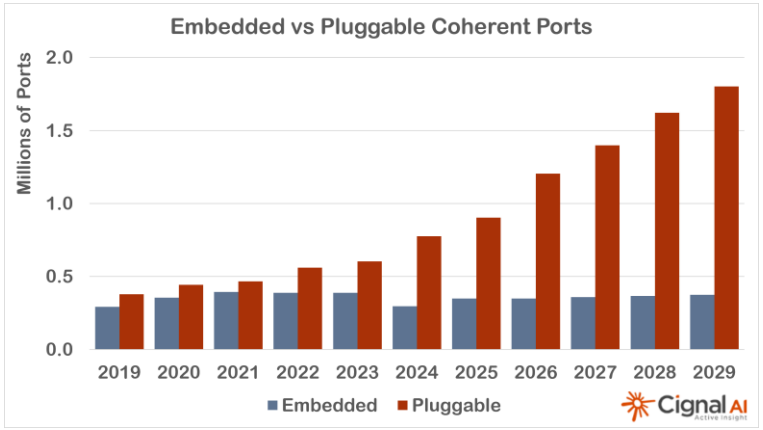

Coherent Pluggables Expanding Applications

Farther: 400G ULH

Wider: L-Band

Shorter: Coherent Lite

800G Coherent Market Takes Shape

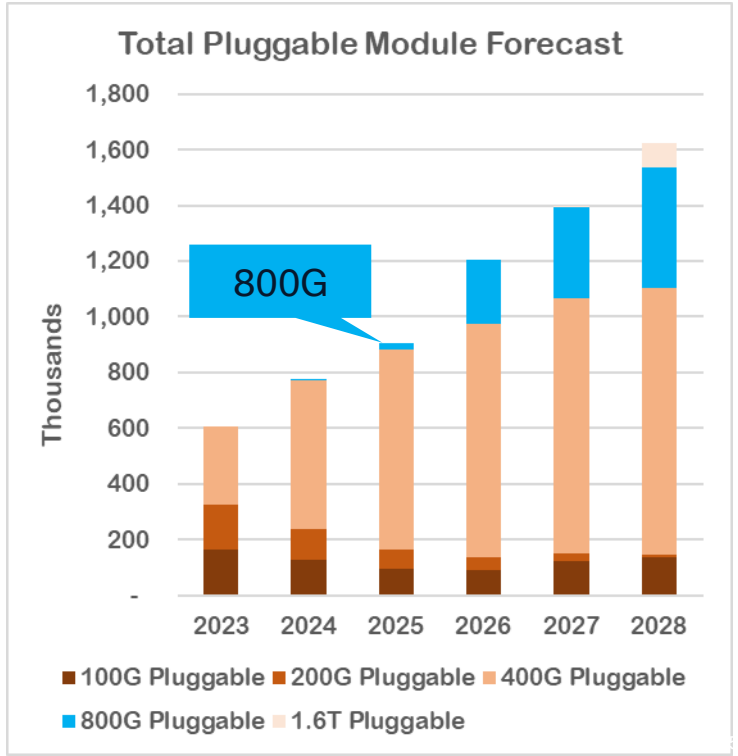

Acacia/Cisco continues to be the pioneer in coherent pluggables: 800G ZR+ (interop PCS), 400G ULH

Coherent Pluggable Brings Some Important Levers

01

Network costs and efficiencies

CapEx cost breakdown, $/100G, PAYG models, etc.

02

Power, thermal, and space

Compact devices, Watt/100G, DC power and rack space optimization

03

Network architectures

Disaggregation, optimize topologies, spectrum efficiency

04

Drives interoperability

MSAs/standards, best in breed adoption

Pluggable Coherent Scales Across the Network

- Coherent interfaces in routers save power
  - Designs optimized for power and performance
  - Eliminate extra client optics interfaces between routers and transponders
- Based on standardized interfaces with proprietary options for challenging applications
- Use cases expanding from metro to regional and long haul

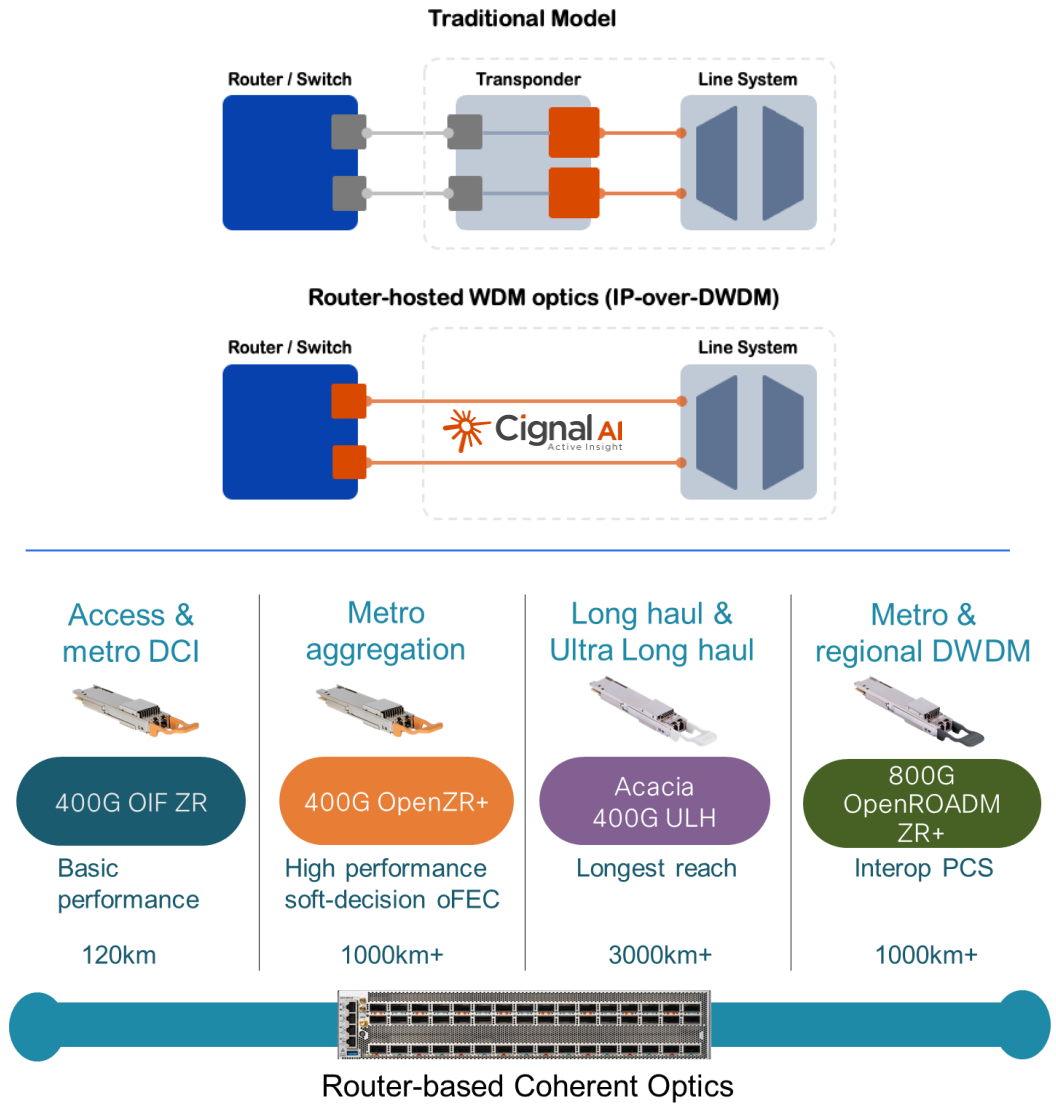

Current 400G MSA Pluggables Portfolio

Industry leading 400G pluggable deployments

Best in class quality

Standards compliant

OIF 400ZR · OpenROADM · OpenZR+ MULTI-SOURCE AGREEMENT · ITU · CableLabs · IEEE

QSFP-DD

ULH 400GZR+x → > Bright + Multiple Baud-Rate Support

Bright 400ZR+x → >+1dBm, Internal OA, TOF

OpenZR+ → Enhanced OSNR (over amplified link)

400ZR → 120km (over amplified link)

400G ER1 → 45km (unamplified link)

100ZR → Scale up in QSFP-DD

CFP2

400G OpenROADM → Enhanced OSNR, OTN

Coherent BiDi → Dual laser, Single fiber Bidi, LEO

cisco

Expanding Pluggables to Access & Long-Haul networks ..

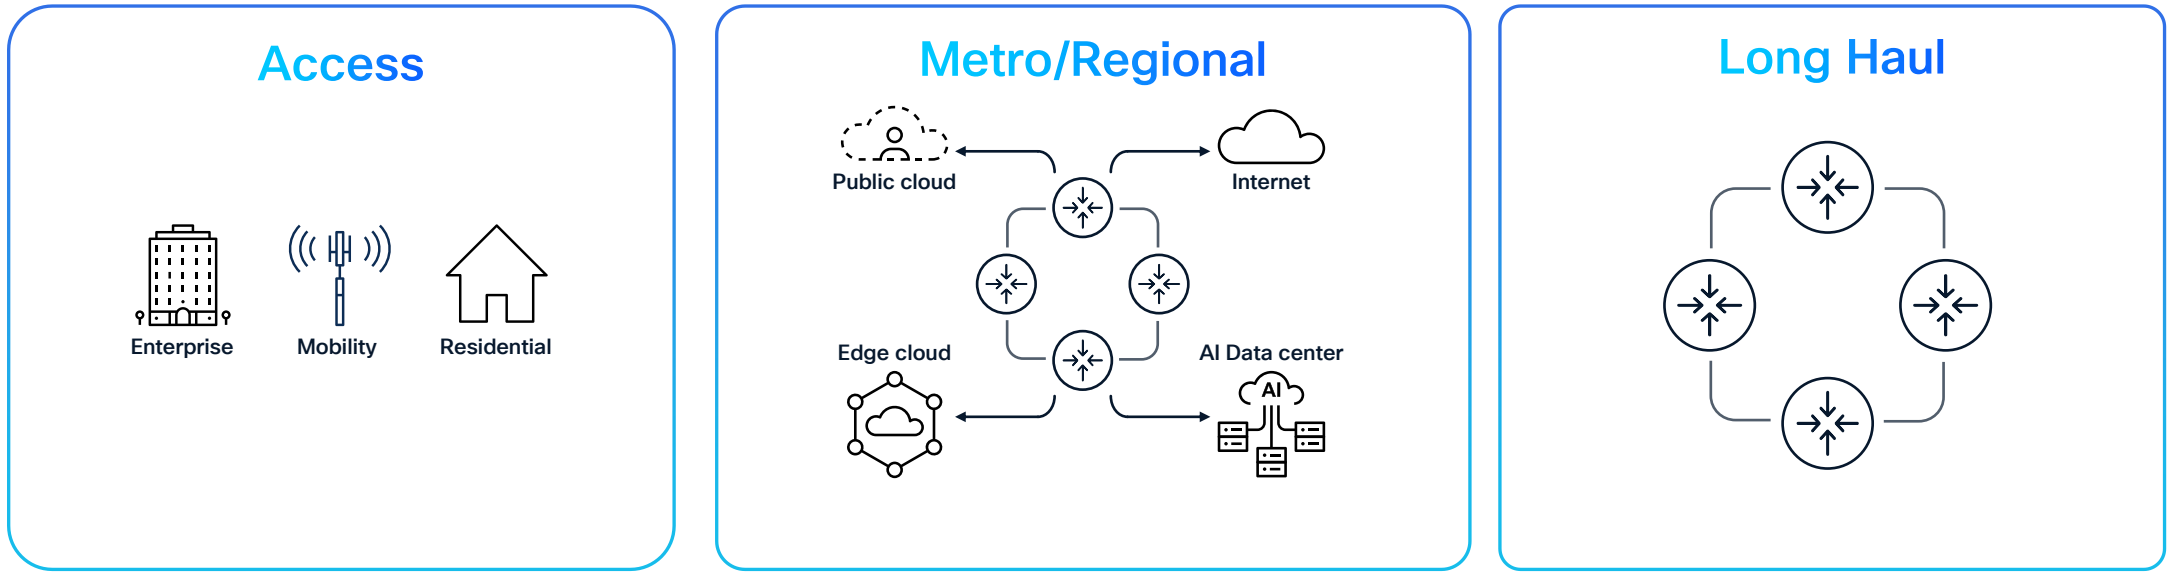

100ZR QSFP28: interop with Bright

400ULH – interop with Bright

Shipped 600K+ 400G pluggable ports (400ZR/ZR+/Bright)

800ZR/ZR+: Improved W/G, $/G

Vertical Integration is Key to Lowest Cost and Best Performance

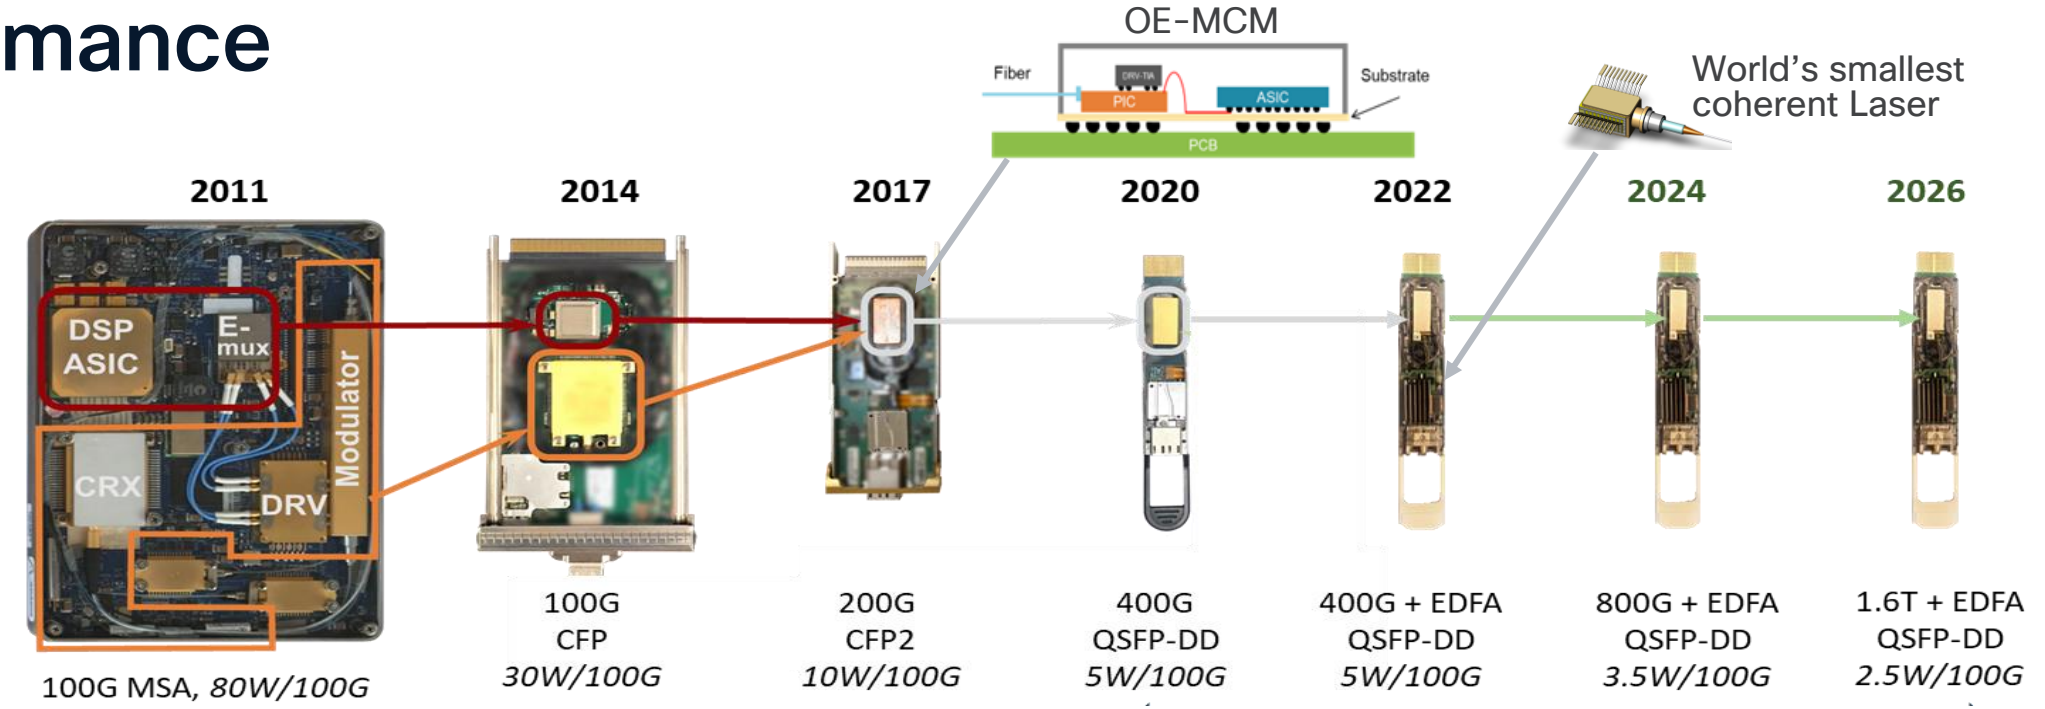

In-House:	DSP	+ SiPh Optics	+ OE-MCM	+ ASIC	+ Laser	100% Vertical
QAM:	QPSK	QPSK	16QAM	16QAM	PCS	PCS
CMOS:	40nm	28nm	16nm	7nm	4/3nm	2nm
Optics:	Discrete, InP, LiNbO3	SiPh	SiPh	SiPh	SiPh	SiPh/TFLN

Cisco is the only company that is 100% Vertically Integrated today

800G Coherent Pluggables

Industry's first DSP compliant to OIF and OpenROADM (incl. PCS – Probabilistic Constellation Shaping)

MSA

Feature Rich

QSDD & OSFP form factors
C/L Band Variants,
Built in EDFA & TOF

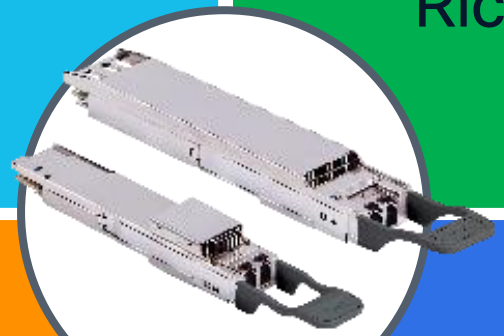

4nm CMOS,
SiPhotonics architecture
<28W (800ZR), <30W (800ZR+)

Low Power

Applications

Metro DCI,
Metro/Regional DWDM
P2P router interconnect

Welcome to the Next-Gen Coherent Pluggables !

- Latest 4nm CMOS process
- Up to 131 Gbaud Optics
- 1st to support interop PCS 800G ZR+ (OpenROADM) and OIF 800ZR modes
- Supports CMIS 5.3
- C & L Band

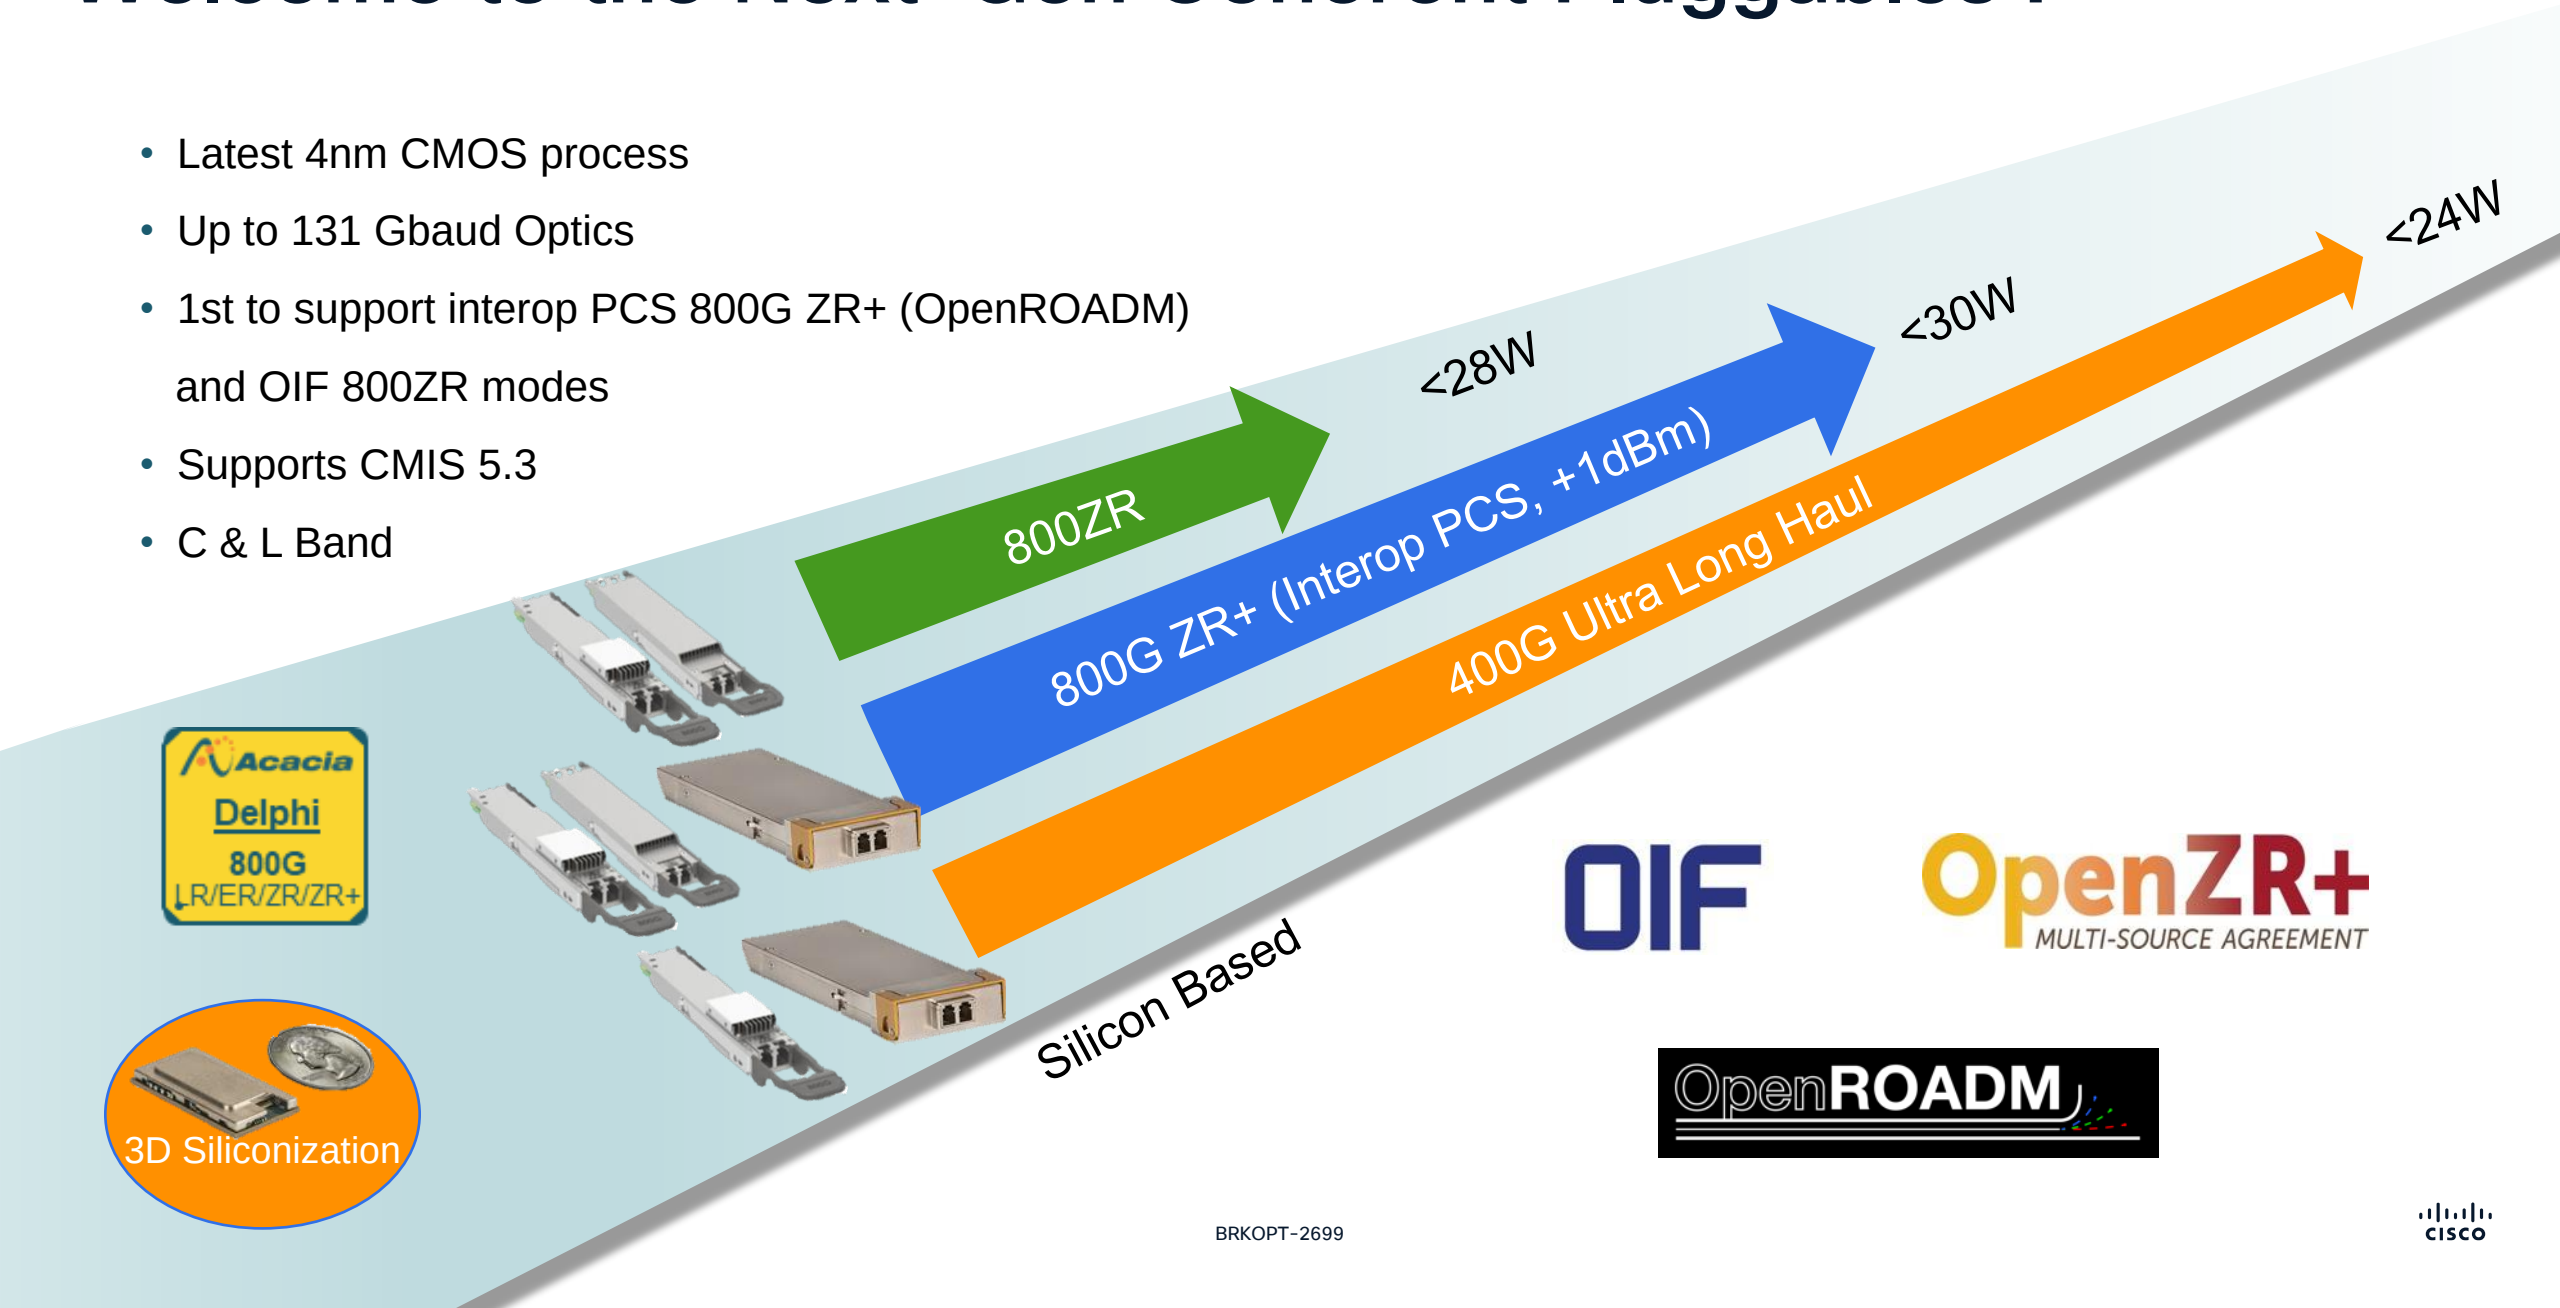

Coherent Optics enables IP-Optical Convergence

Traditional Architecture

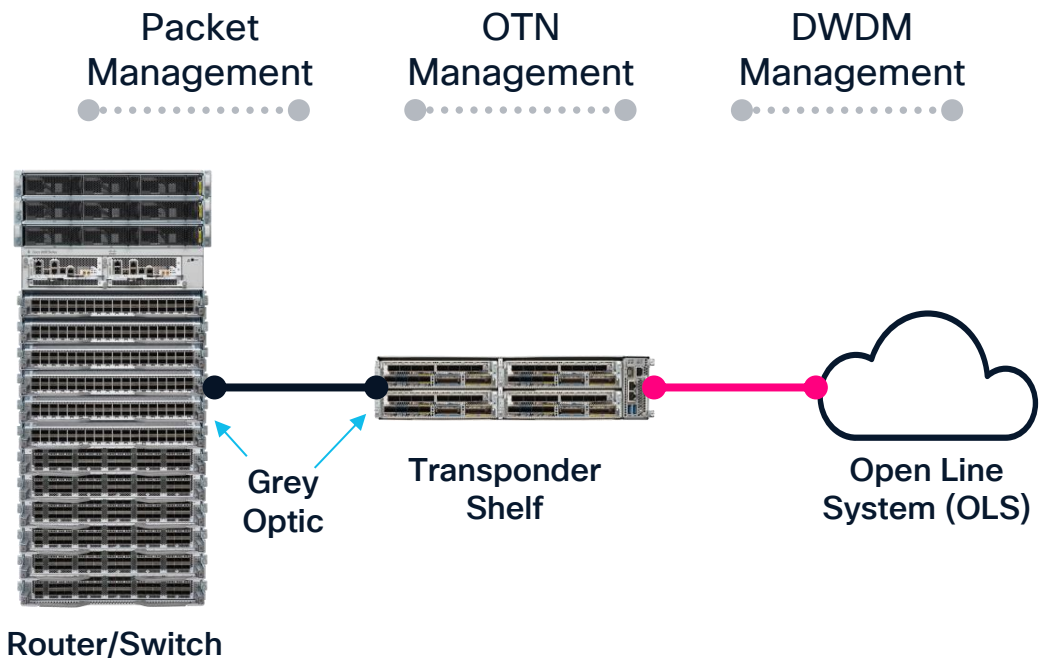

Demarcation between IP and Optical

RON Architecture

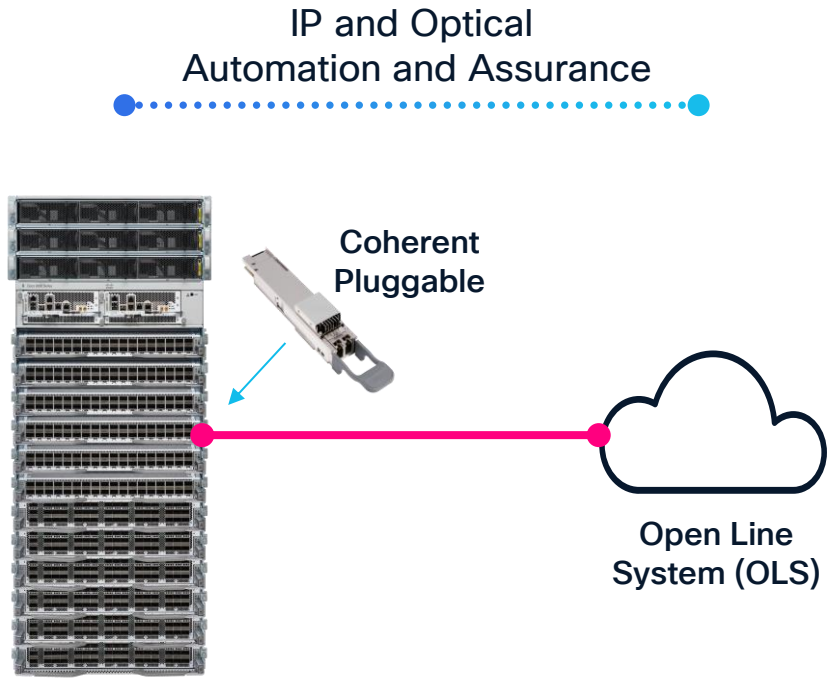

Much Lower Total Cost of Ownership (TCO)

Significantly reduce capital and operating expenses **by 50% or more**

Routed Optical Network (RON)

TCO*	46%
CapEx*	35%
OpEx*	57%

Supporting different use cases (access to ULH) and different deployment models

1 QDD-ZR/ZR+ over third-party OLS

2 QDD-ZR over dark fiber or passive DWDM

3 QDD-ZR+ over Cisco OLS

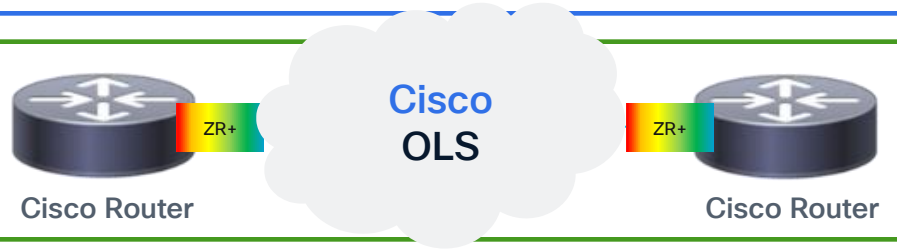

>75% of the RON deployments are over third-party line system!

OLS: Open Line System

* ACG Research 2023

Key takeaways – coherent interfaces

Pluggable solutions are leading option for coherent interfaces – including 100G now

400G coherent solutions (and standards) are mature with lots of deployment options

800G coherent solutions (and standards) available – interoperable PCS standardization improves performance

Routed optical networking: improved performance, operational efficiency & cost efficiency

Increasingly relevant for AI Infrastructure

Optics Advancements for AI Infrastructure

Accessing AI Potential | Scaling AI Cluster

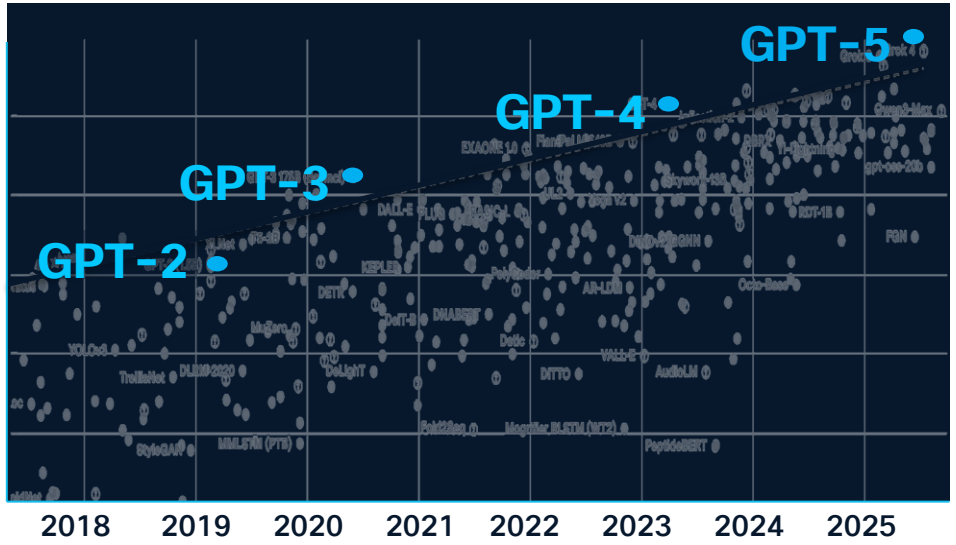

Baseline graph from Epoch AI, September 2025, added GPT5 point manually

Other major models like Llama, Gemini, Grok, others have excellent performance. Only using GPT to simplify trends.

3,000x	Parameters	2019 – 2025
100,000x	FLOPs	
5,000x	GPUs	
2.5x	Knowledge & Reasoning MMLU	

FLOPs | Floating point operations
MMLU | Massive Multitask Language Understanding (measures general knowledge & reasoning)

GPT-2 | 2019
~1.5B parameters
~1e21 FLOPs
~10s-100s GPUs
~36% MMLU performance

GPT-3 | 2020
~175B parameters
~3e23 FLOPs
~1,000s GPUs
~43% MMLU performance

GPT-4 | 2023
~1T parameters
~1e25 FLOPs
~10,000s-25,000 GPUs
~86% MMLU performance

GPT-5 | 2025
~2-10T parameters
~1e26 FLOPs
~50,000s-100,000+ GPUs
~92% MMLU performance

Statistics generated by ChatGPT4.1

Cluster size unlocks intelligence

Data Centers | Limited by Power

Meta Hyperion Data Center

Up to 5GW

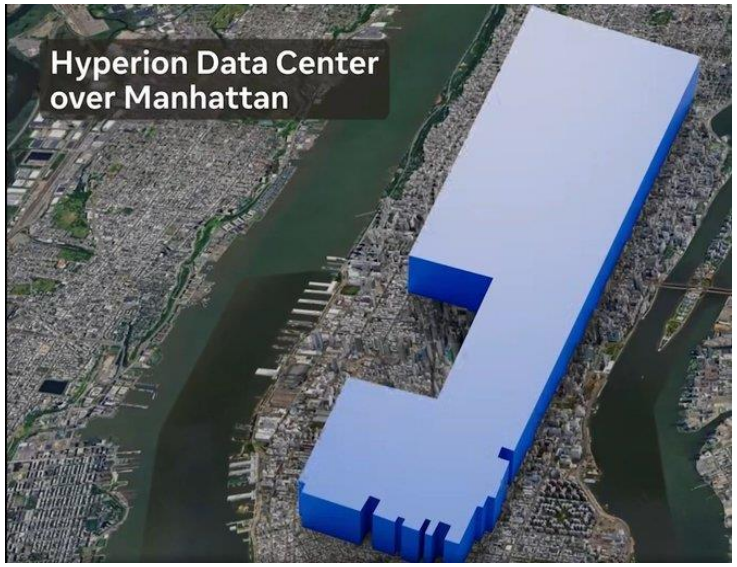

Meta's planned Hyperion data center (Credit:Meta)

Average kWh Rates

March 2025

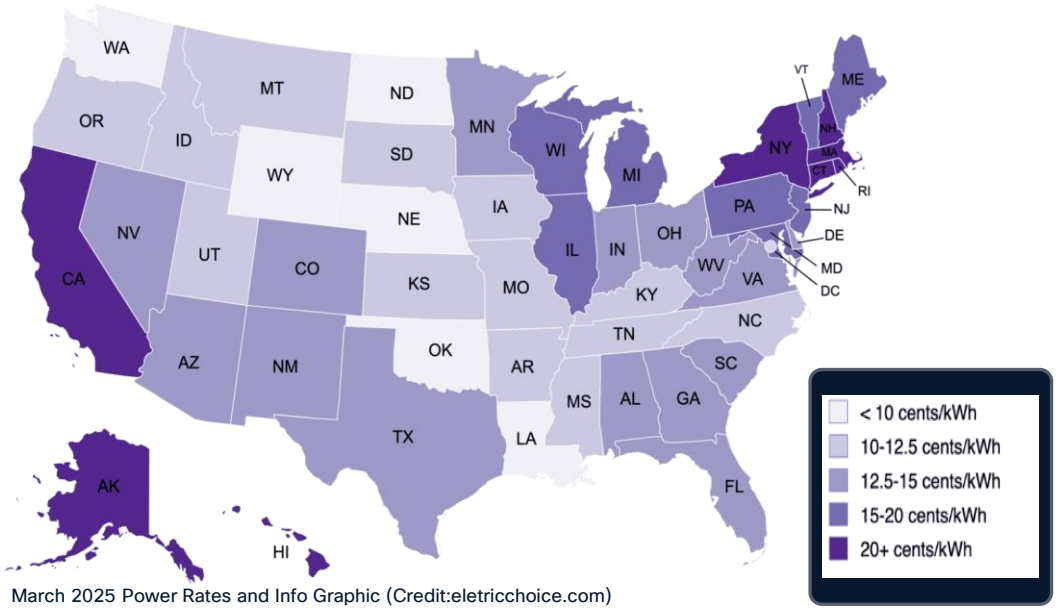

March 2025 Power Rates and Info Graphic (Credit:eletricchoice.com)

Limited real estate & high electricity costs

Data centers must "Scale Across" to other sites

Backhaul the traffic to consumers

Explosive growth in WAN traffic

Power limits the ability to scale

Priorities for Interconnect in an AI world

Minimize workload impacts
- Physical reliability
- Link Performance

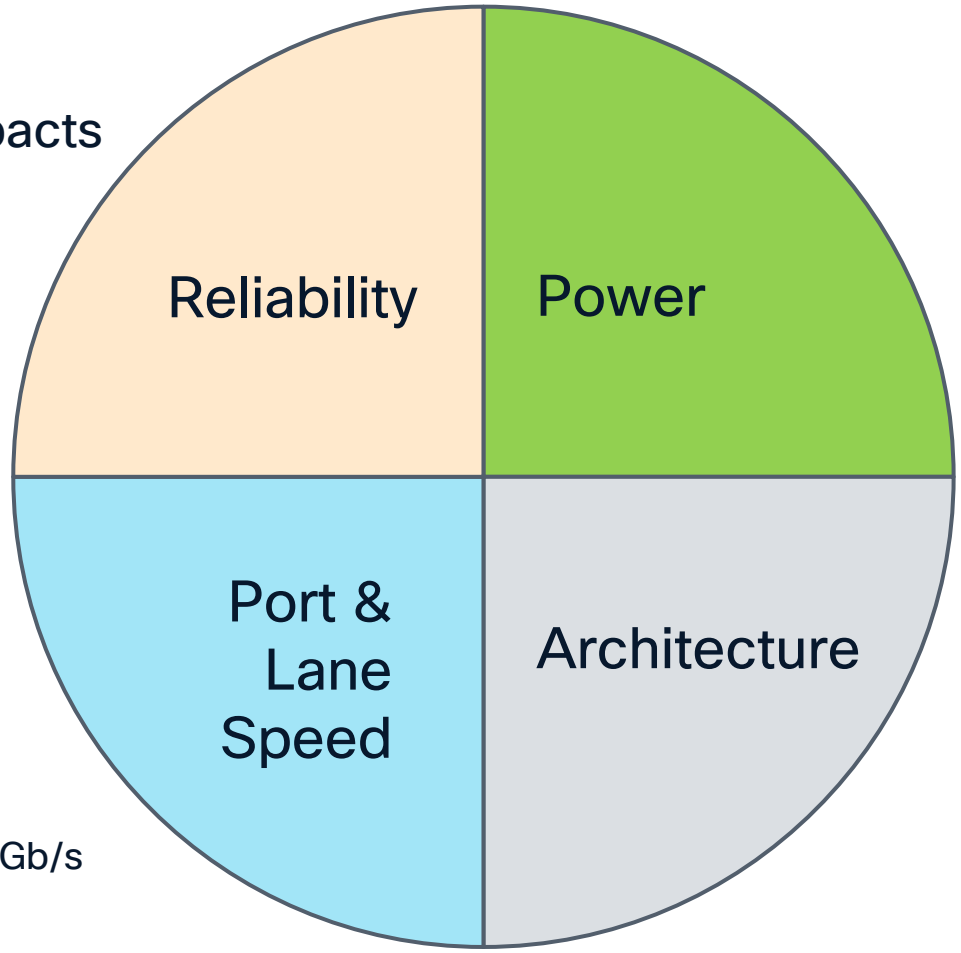

Lower Watts per Gigabit
- Copper cables
- Linear Optics

Interfaces
800 GbE→1.6 TbE → 3.2 TbE

SERDES and λ
100 Gb/s → 200 Gb/s → 400 Gb/s

New architectures to support the necessary scale

Changing Architectures| supporting AI infrastructure

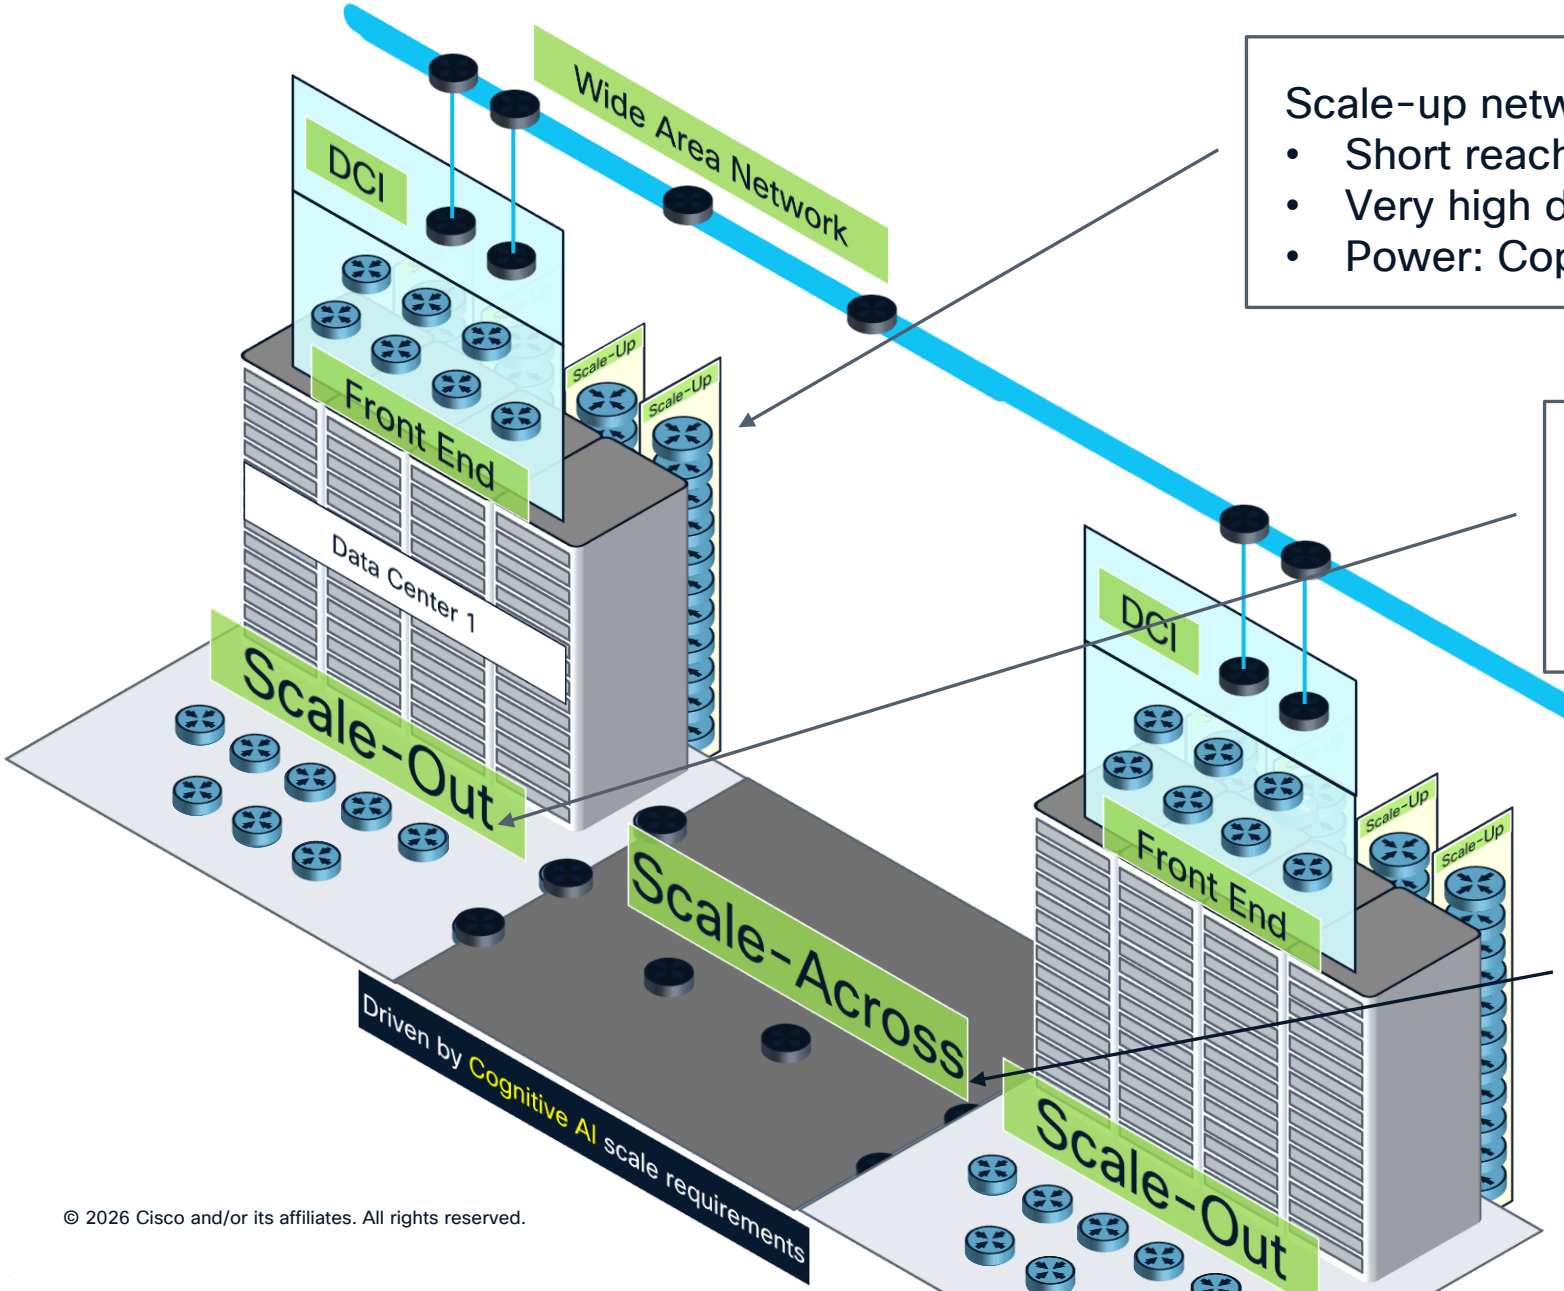

Scale-up network
- Short reach (latency)
- Very high density
- Power: Copper preferred, optics coming

Scale-out network
- Longer reach (optics)
- Pluggable

Scale-Across network
- Extending AI Cluster across data centers
- Pluggable long reach (coherent optics)
- Deep buffer switches
- Will drive exponential growth in WAN traffic

Considerations for Next-Gen Scale Across Architectures

AI Driving New Traffic Granularity

Traditional Transport
How many wavelengths?

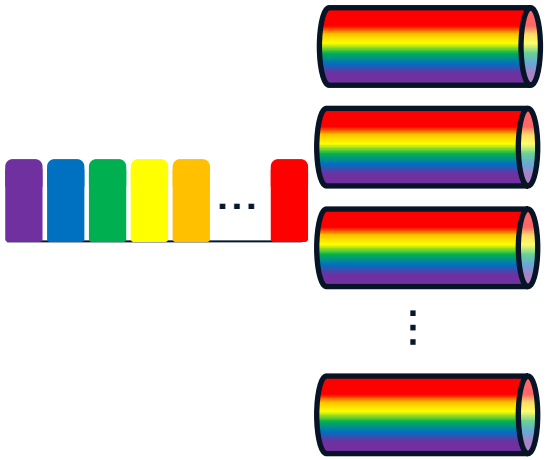

Scale-across AI
How many fibers?

Alignment of C-band and Router Capacity

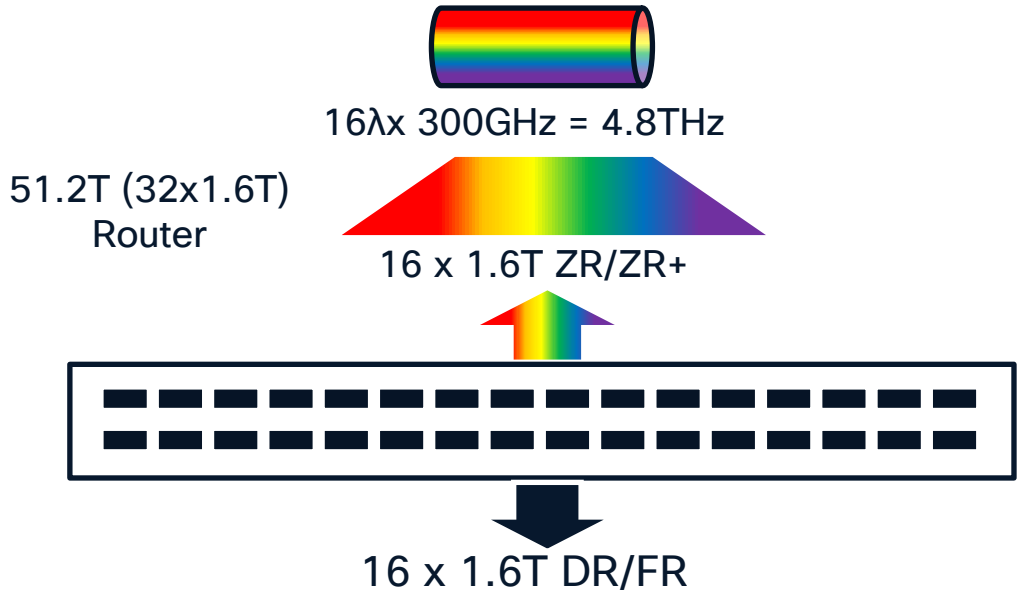

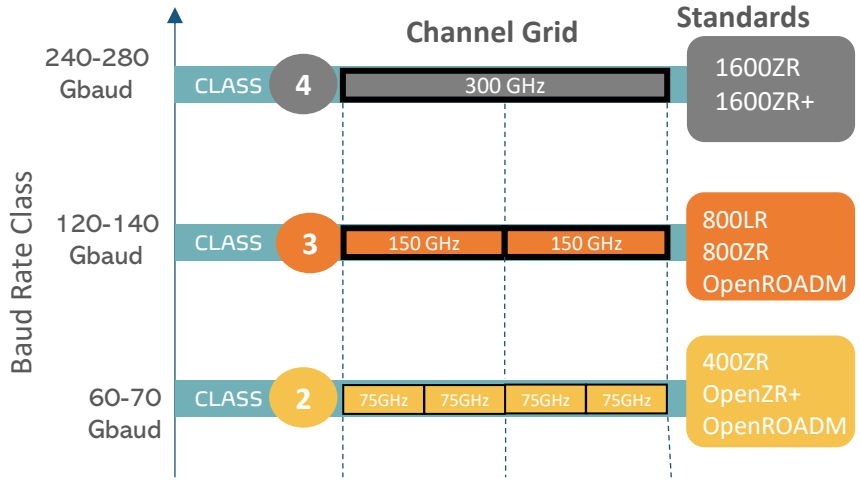

Reliability

Why Optics Reliability is critical for AI datacenter

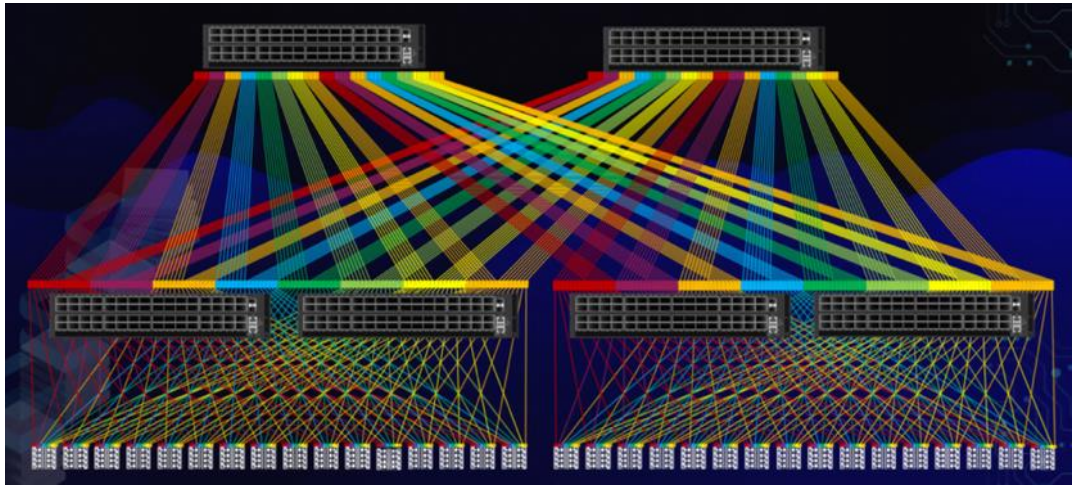

- Distributed AI have non-traditional E-W workloads
- Sometimes, a single transceiver failure can cause measurable impact, even if *intermittent*:
  - can affect data synchronization across nodes
  - can slow down parallel computation
  - can reduce training accuracy
  - can reduce system stability
  - adds operational/ financial losses

Causes and Consequences of Optics Failures in AI Datacenters

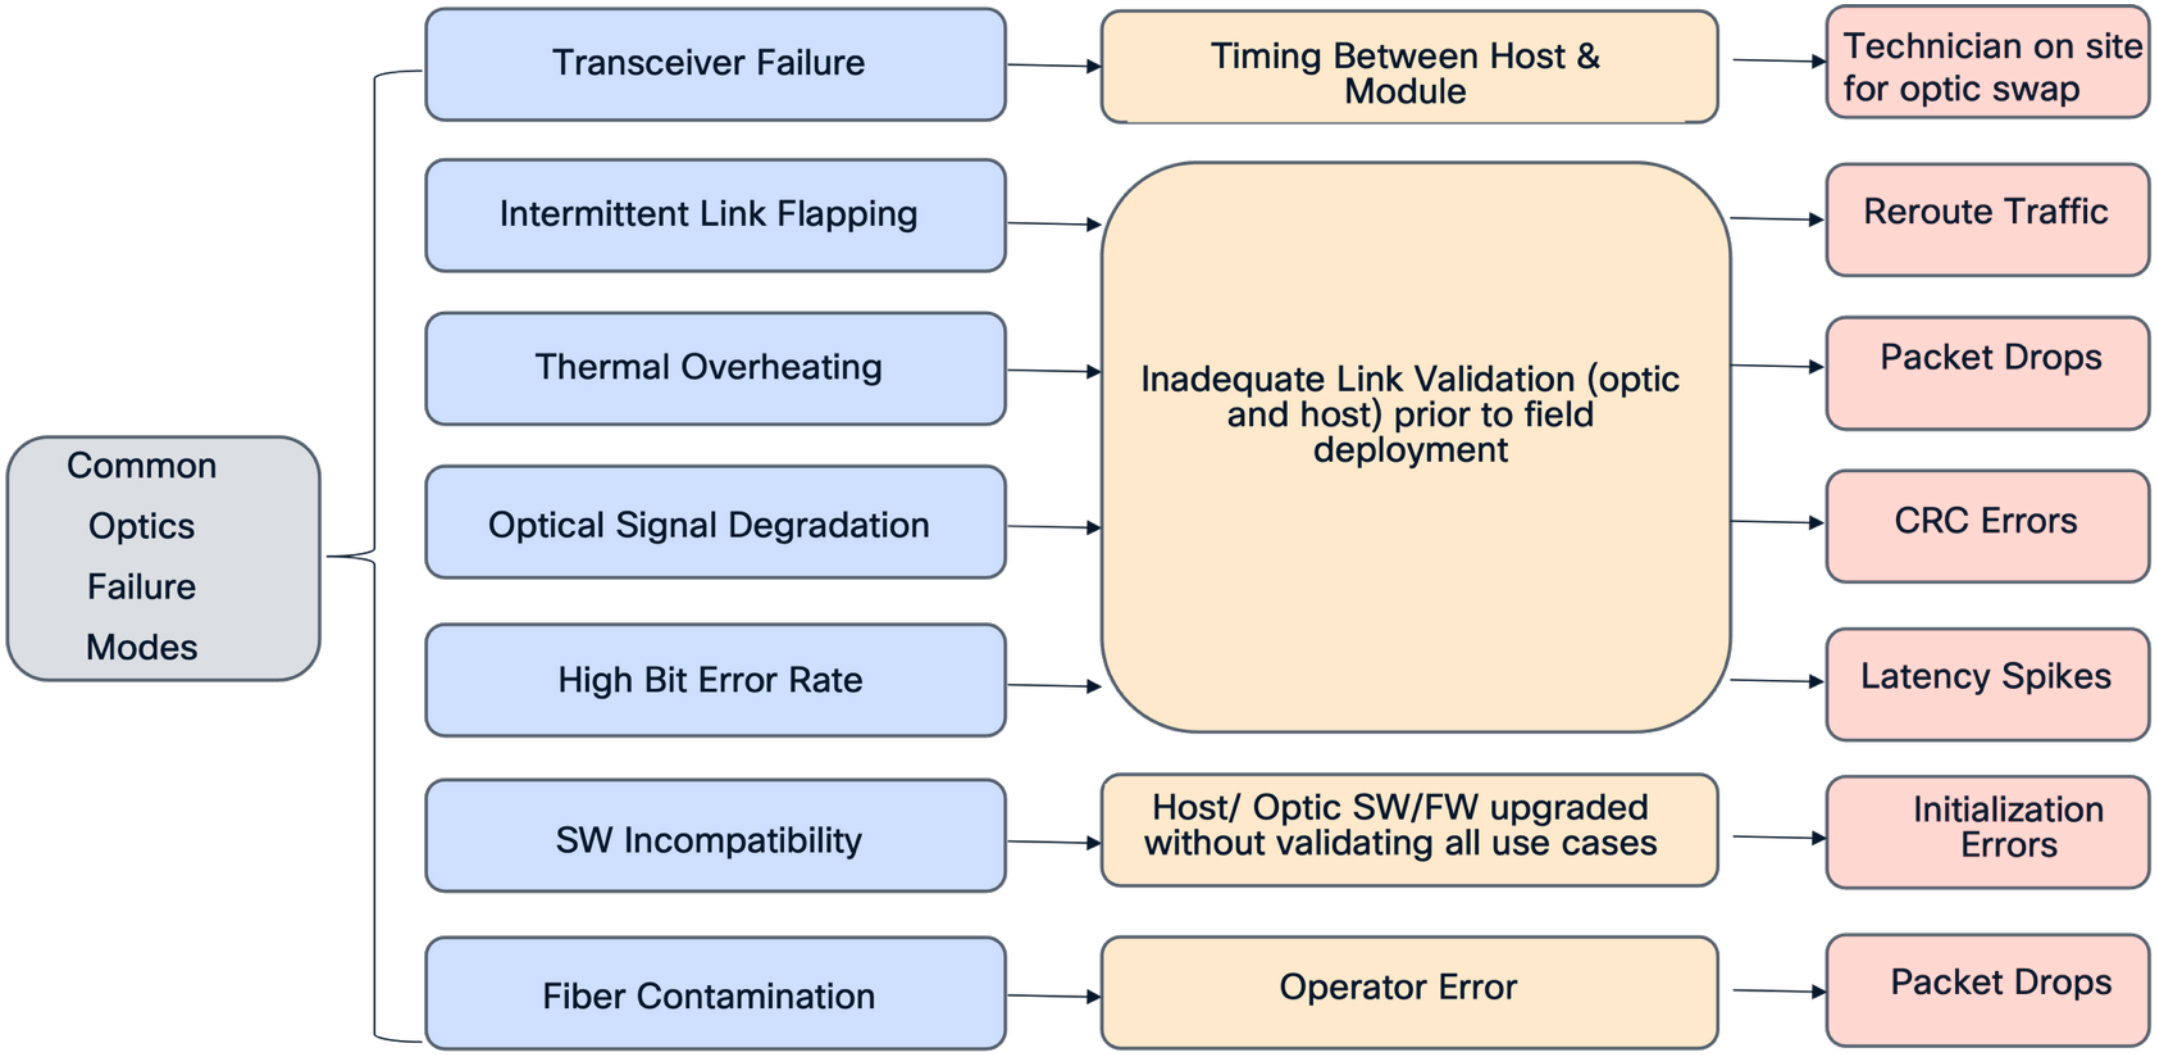

Cisco Optics for AI Datacenter

To meet the needs of AI Datacenters, Cisco leverages innovative Silicon Photonics, a Diverse supply chain, ensures quality with extensive Performance & Reliability Testing

Technology Innovation

Silicon Photonics Platform Innovation

Single Mode and Multimode Optics

Advanced Manufacturing Automation

Optics Design

High Quality Components

Performance margins beyond IEEE

Supply Chain Diversity

GF and TSMC SiPh wafer fab

Multiple DSP sources

Multiple laser sources

Mfg in Thailand and Malaysia

Robust Module Testing Protocols

Robust System Testing Protocols

Advanced Stress Testing

Testing at scale

Unique to Cisco

Reliability

Extensive L1/L2 qual

Unique to Cisco

Field Firmware upgrade

Module FW can be downloaded from software.cisco.com

Example of issue impacting module performance

Issue:
BER variations –up to 5 orders of magnitude- between power cycles.
On the real-world, this could cause link flaps if environmental and use conditions are similar.

BER jumps detected during Cisco Test

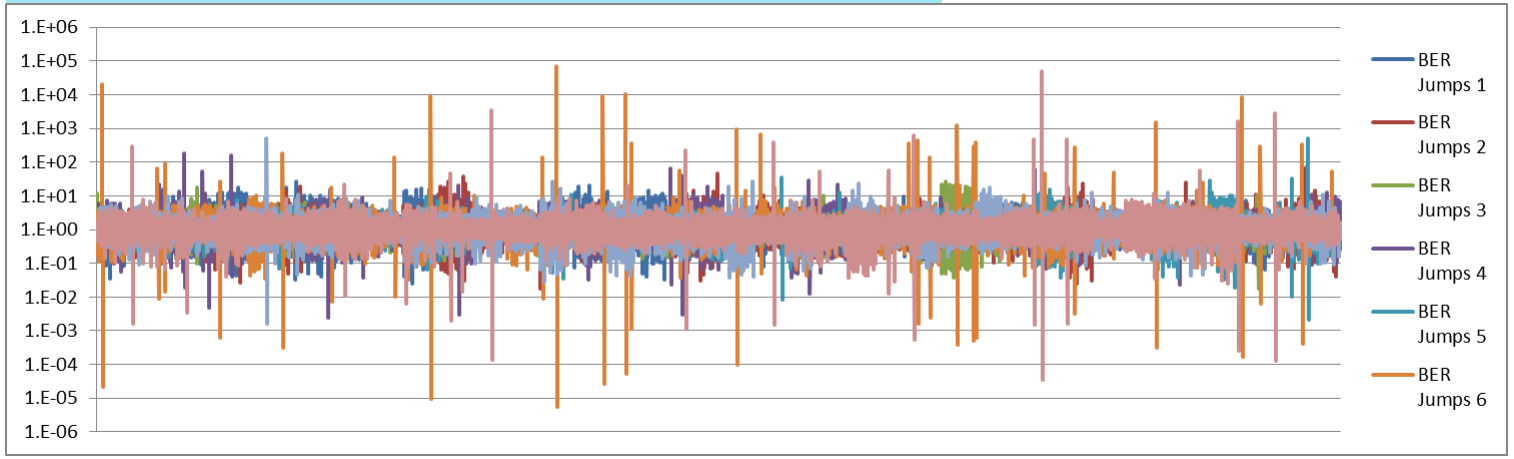

Solution:
Modification to module FW.
BER variations reduced, unlikely to see link flaps on the field after correction.

BER jumps resolved through new FW

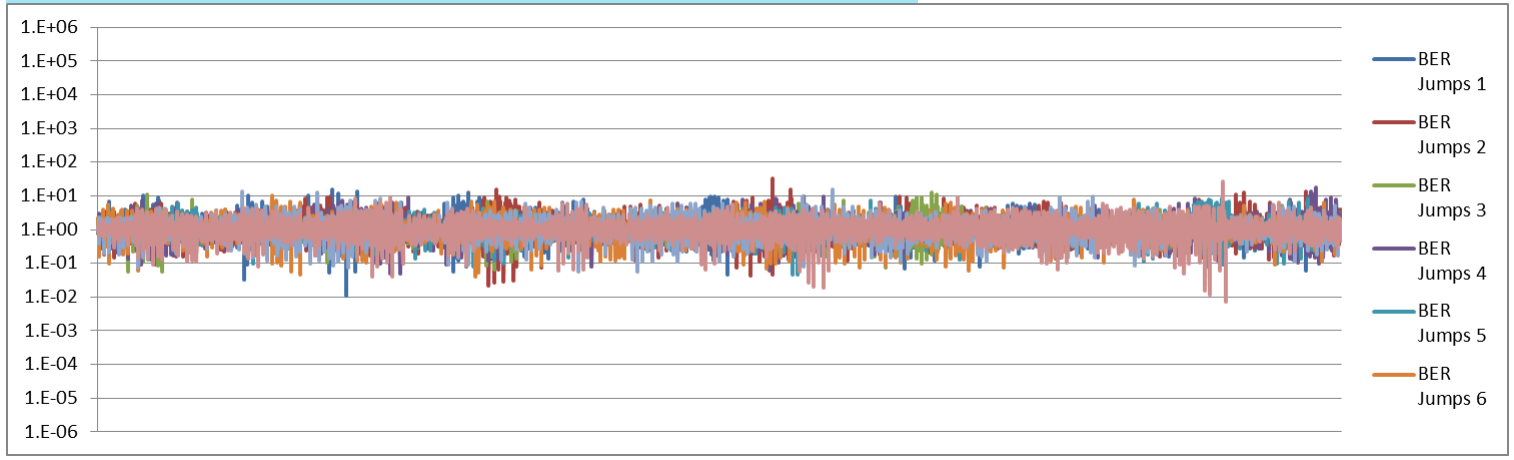

Cisco Optics Reliability

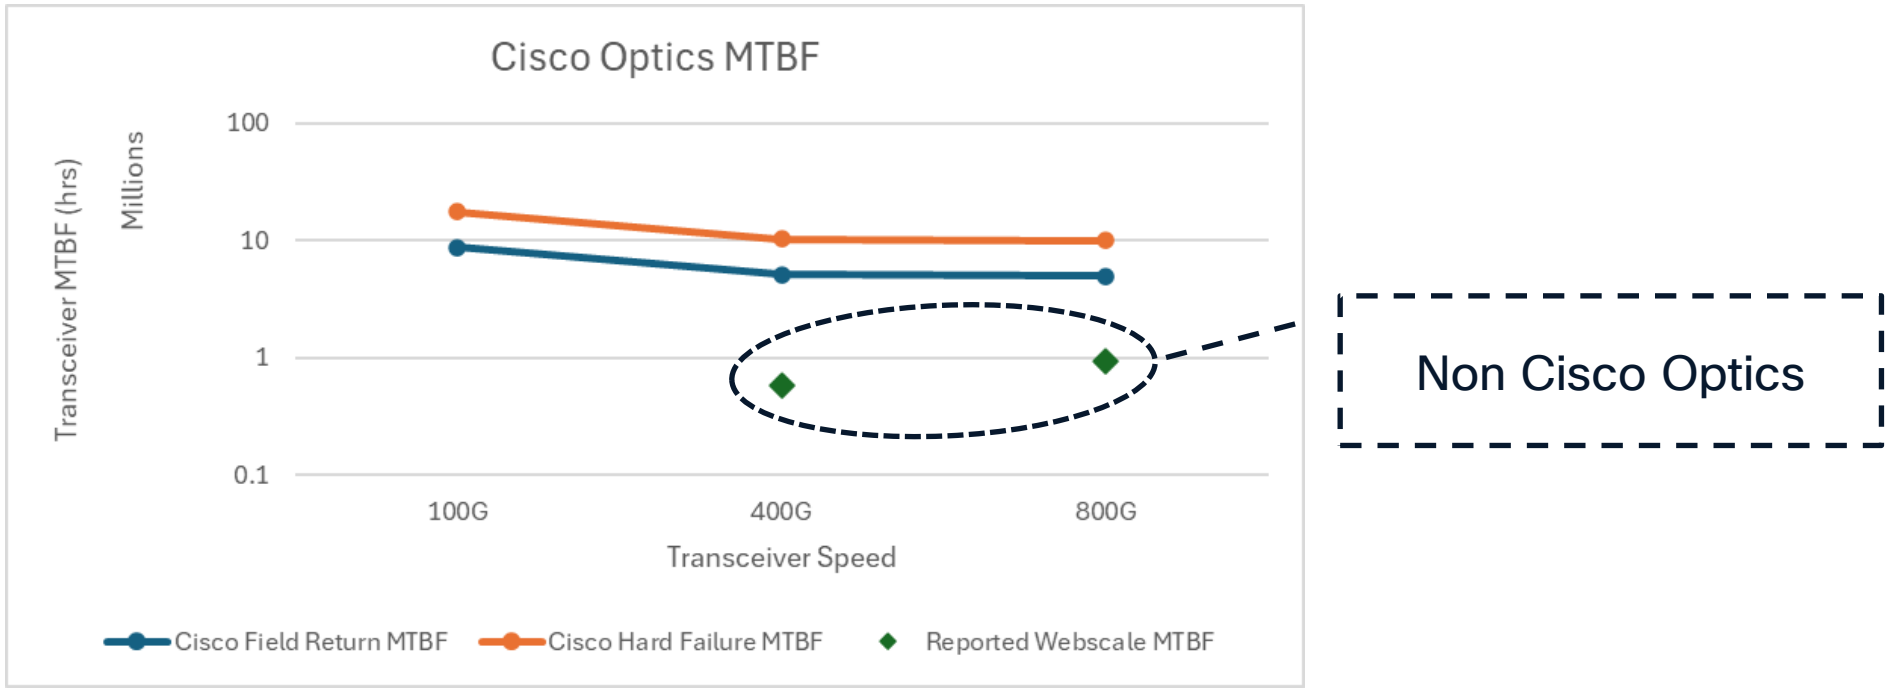

- Cisco Optics are designed for robust performance across many platforms and applications
- Cisco leverages decades of transceiver performance and production experience to design out problems of the past
- Cisco does get some field returns, but roughly half are no trouble found
- Some high-volume industry players only focus on transceiver cost
- AI applications require higher performance than typical DC for optimum ROI on GPUs

Power & Integration

Systems and power scaling | 640x BW, 119x Power

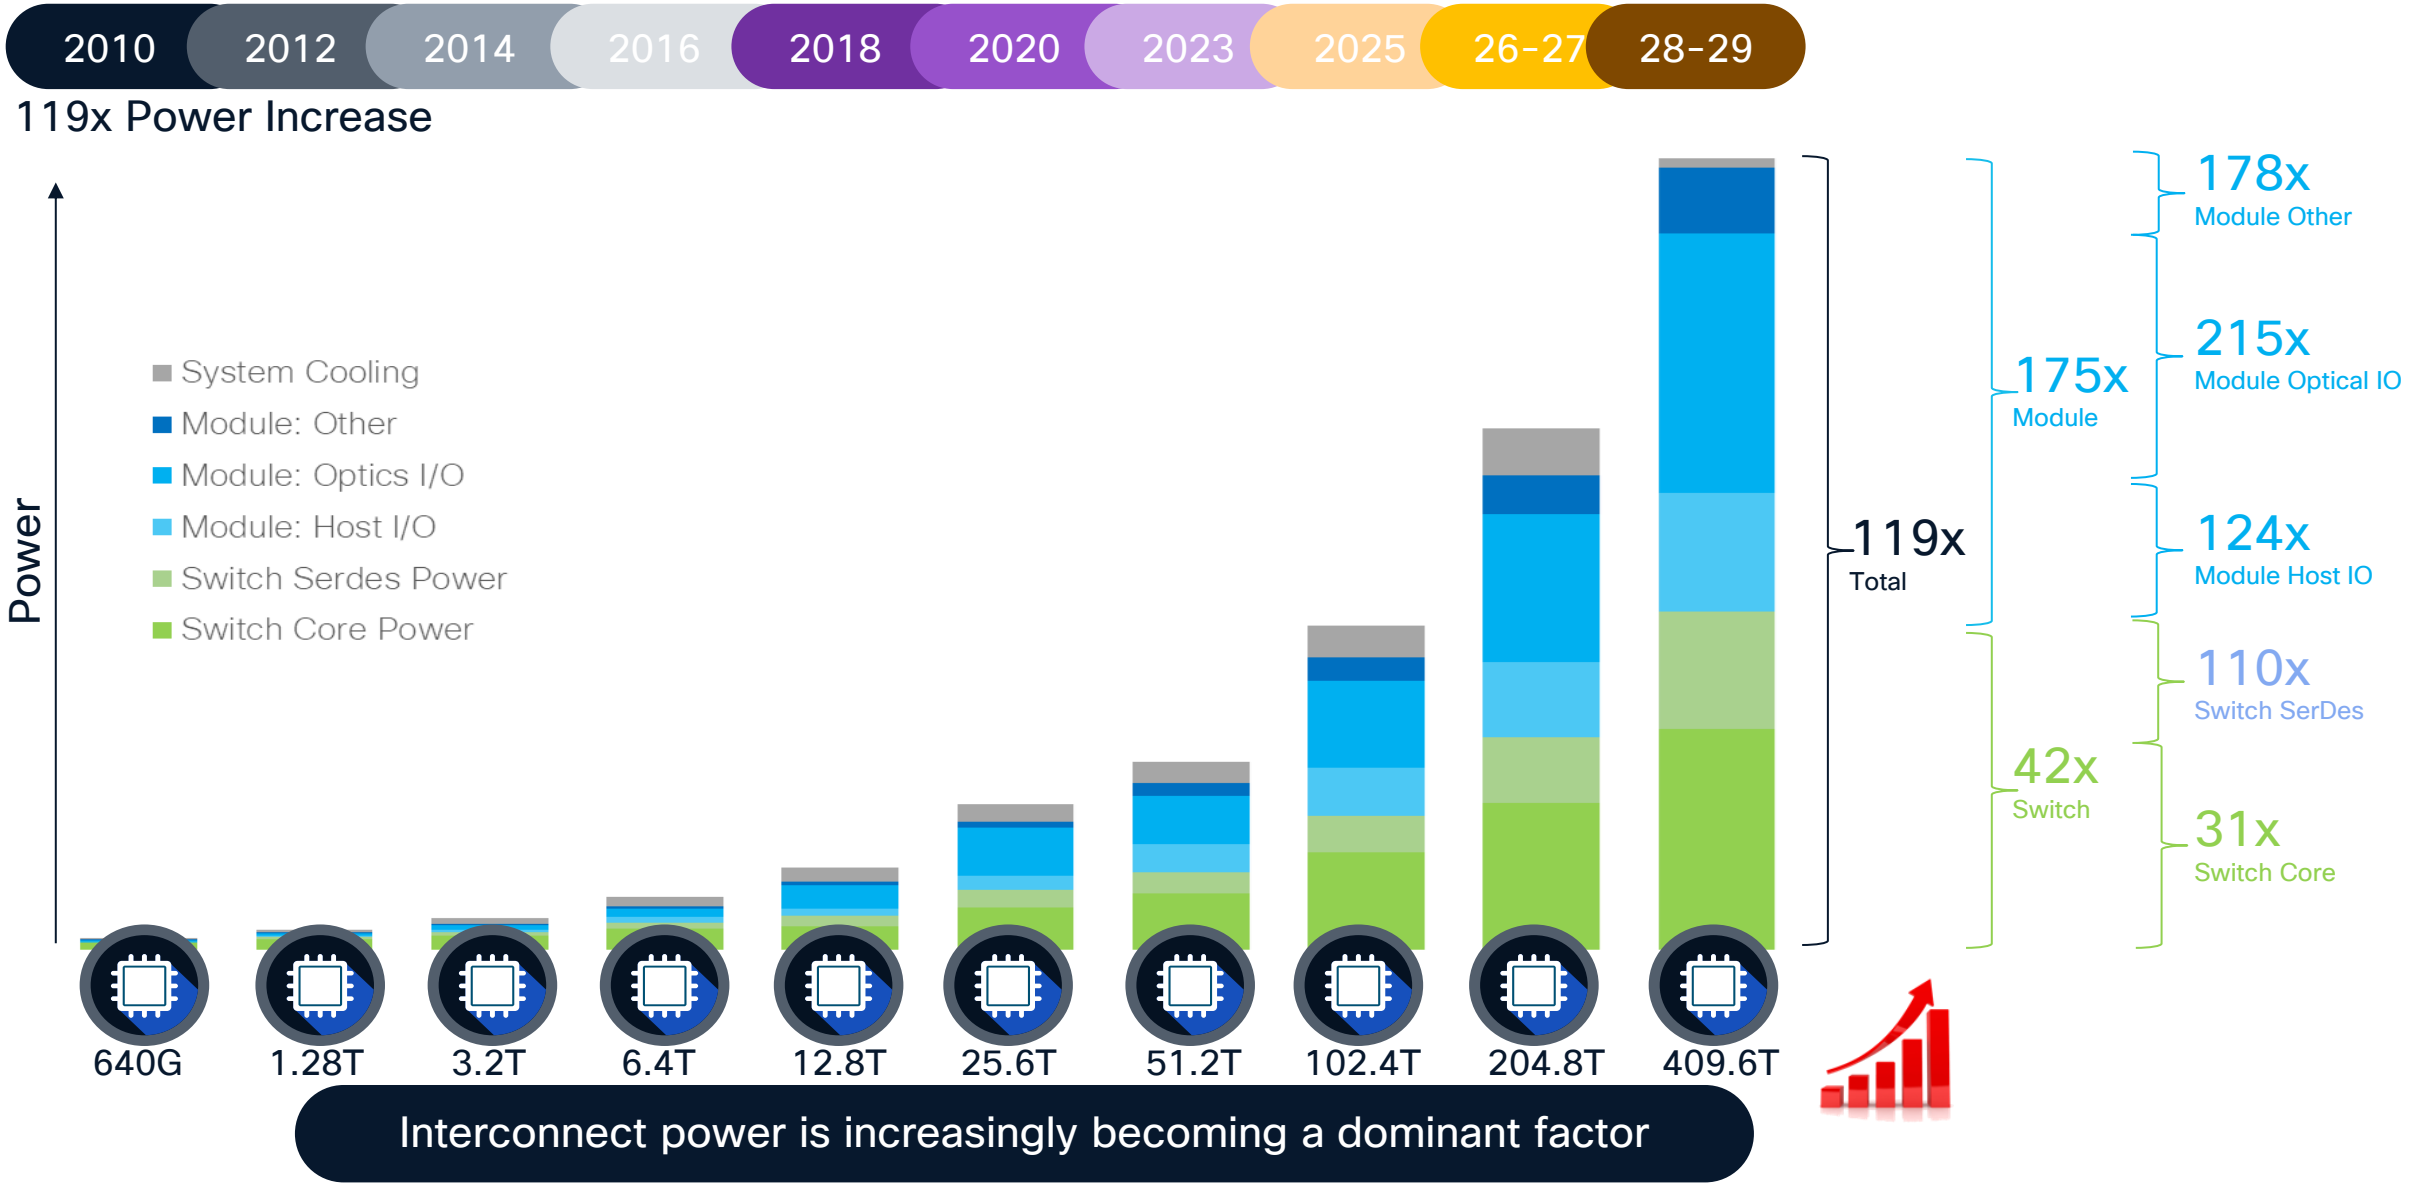

Lowering the power for pluggable interconnect

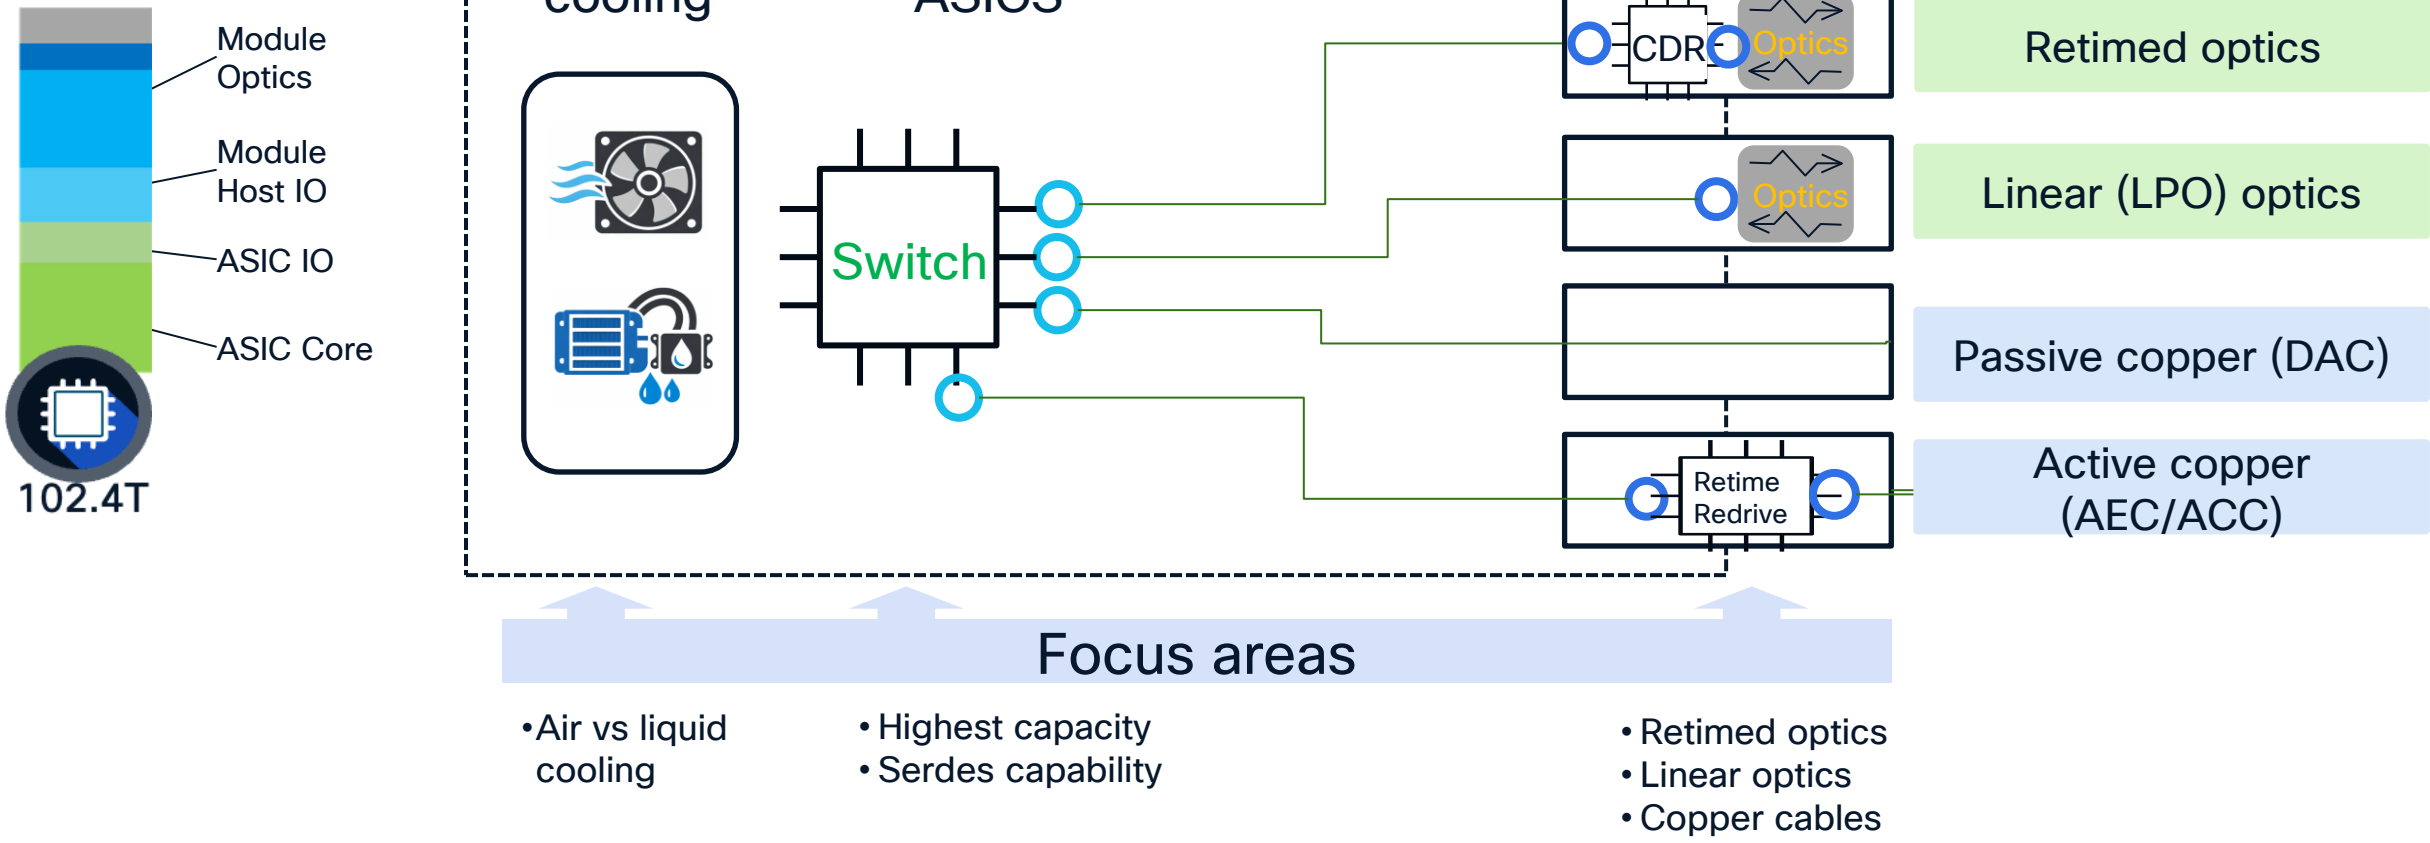

High-Density Interconnect options: Power vs. Reach Tradeoff

Power and Reach @ 1.6T

Maximize use within a rack

Use when inter-rack

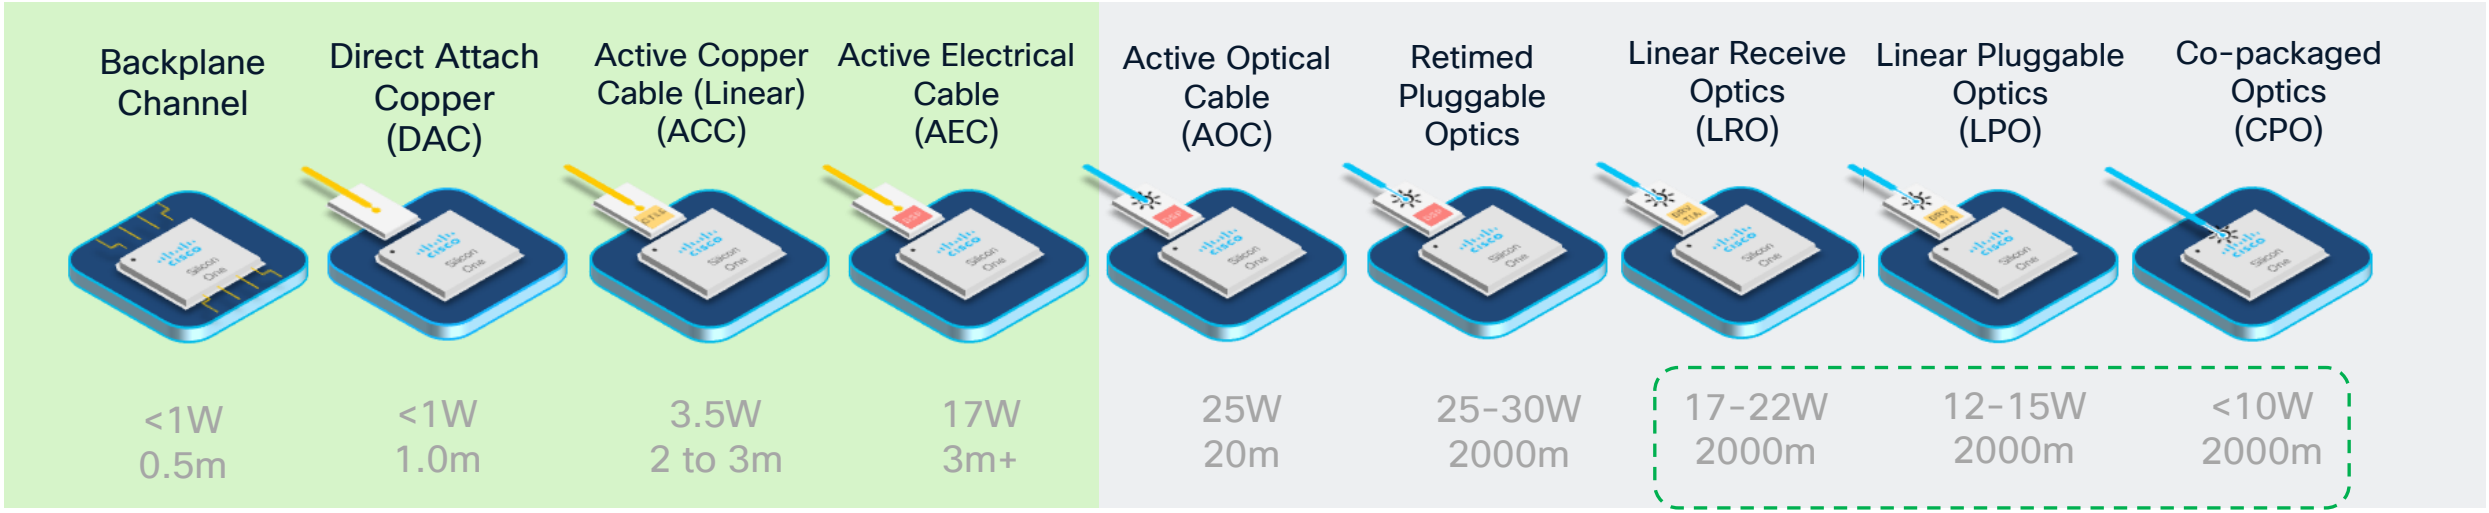

Copper cables and interfaces still look attractive from power perspective

Longer reach = increased power

Power reduction opportunities coming from CPO, LPO & LRO

Linear pluggable (LPO) and Linear Receive Optics (LRO)

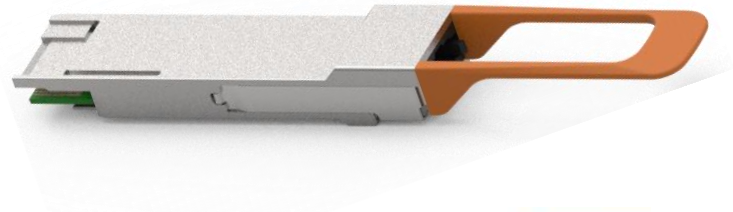

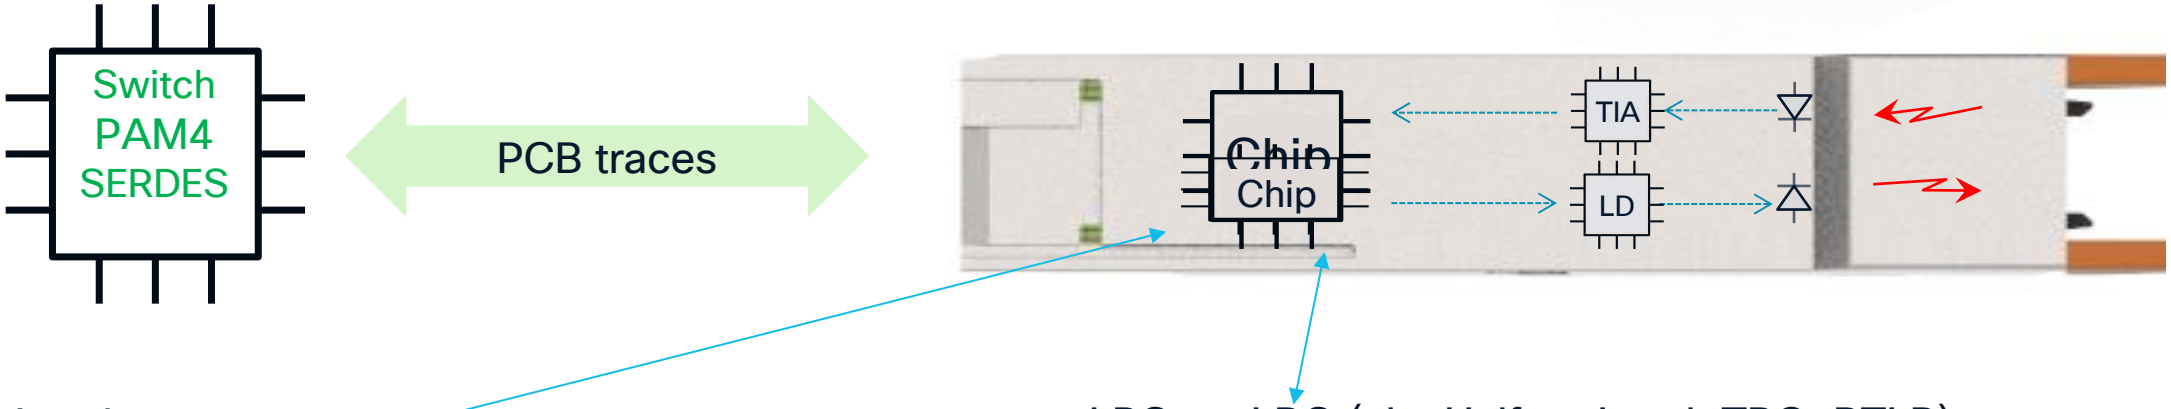

Retimed

- DSP has full digital equalization for both electrical & optical signals
- Enables broad interoperability
  - Host/Port/Module
- Full telemetry & loopbacks possible
- FEC Monitoring or Partitioning possible
- But this **adds** power

LPO or LRO (aka Half-retimed, TRO, RTLR)

- **Power reduction** due to absence of digital equalization
- Performance is based on quality of every component in link since they are concatenated
  - Optics, Serdes, PCB, connectors
  - Varies port to port
  - Some linear gain and equalization in module
- Interoperability becomes more challenging
- Loss of some telemetry, monitoring or loopback

Why does linear optics work?

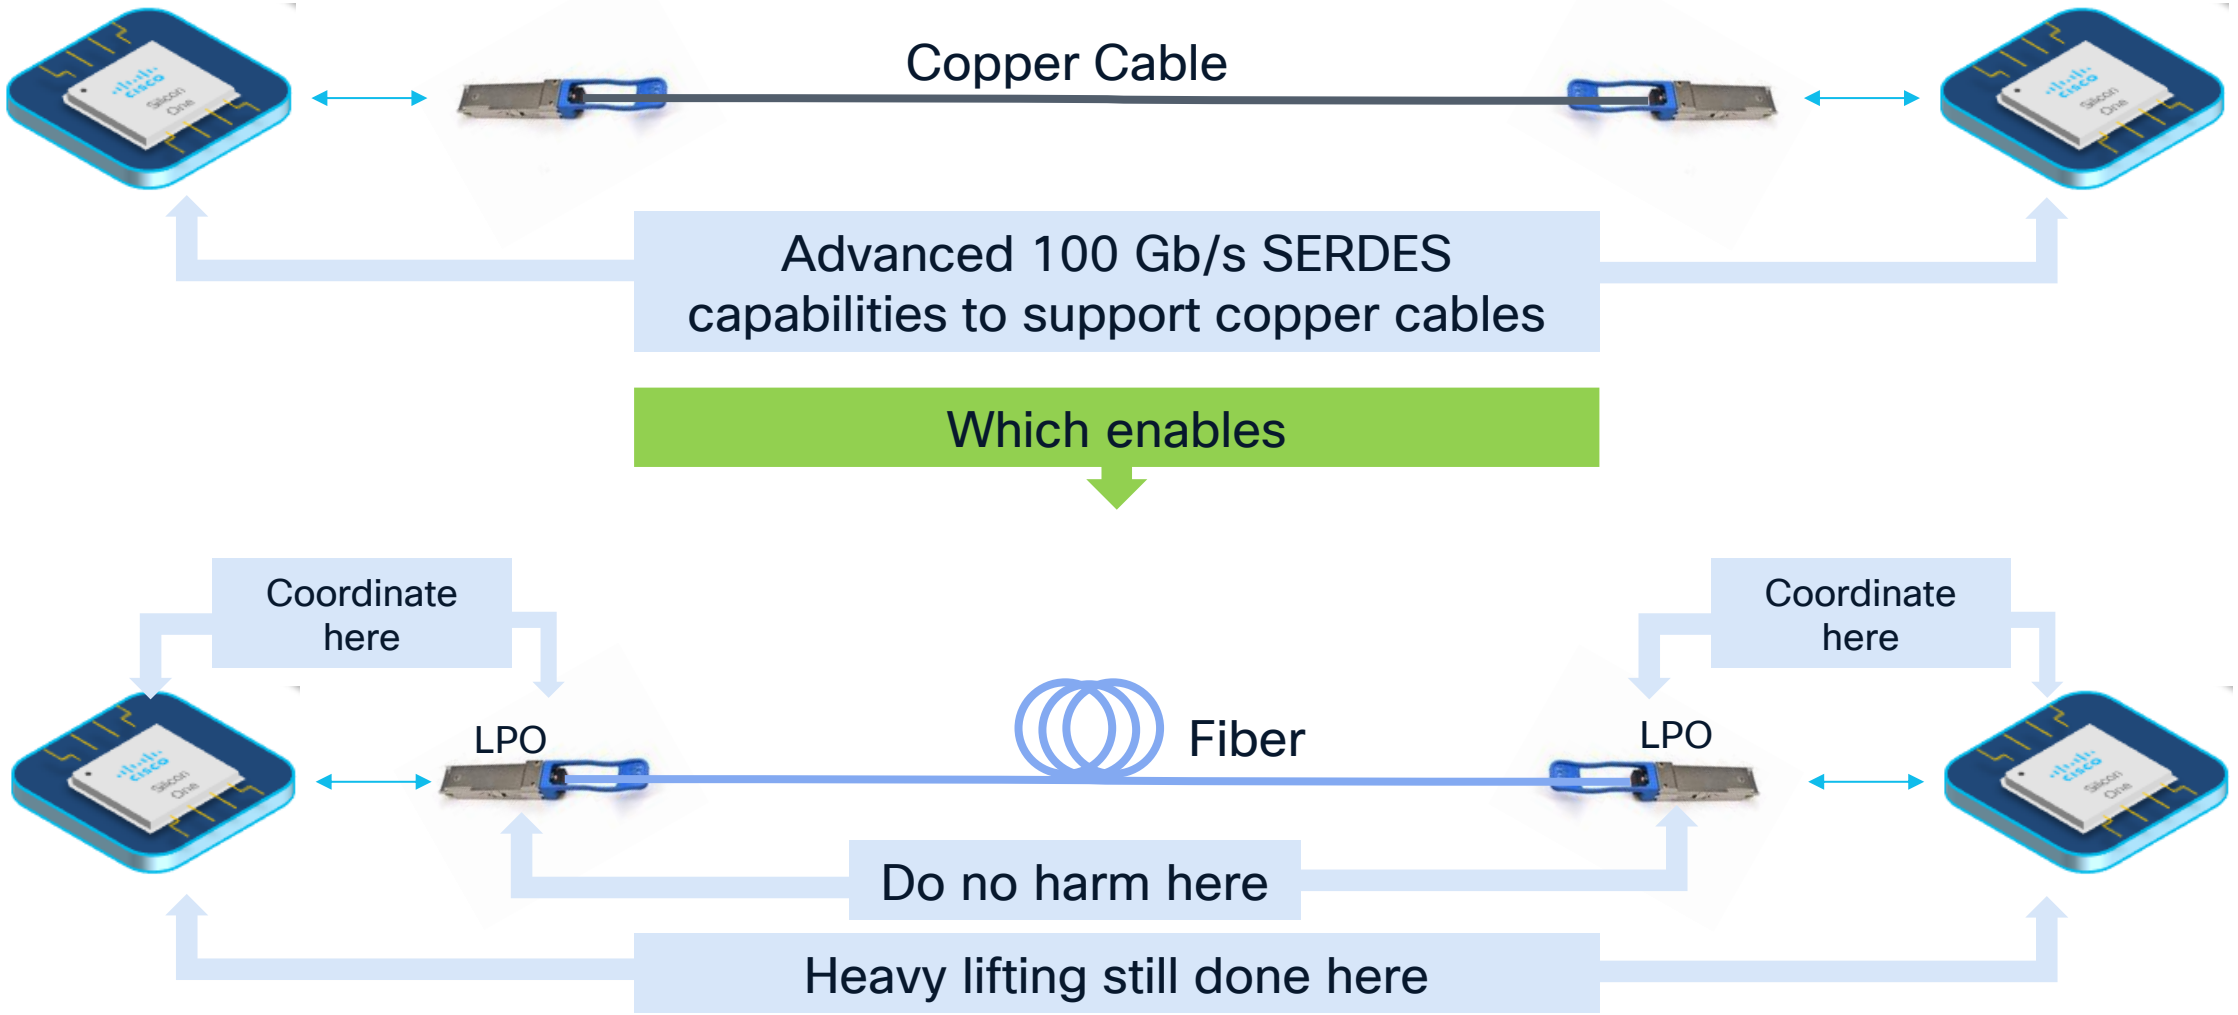

NEW Q3'26

Cisco 800G Linear Pluggable Optics (LPO) Solutions

Powered by **Cisco Silicon Photonics**

AI Scale-out power efficiency and reliability for 800G Networks

Cisco L800G OSFP LPO

50% less power per module

30% less power for integrated systems

Shift power consuming optics DSP functionality to Cisco Silicon One platform

High signal integrity driven by Cisco Silicon Photonics and Cisco Silicon One

High reliability with lower thermal stress

Increase port density without PDU upgrades

Evolution of pluggables

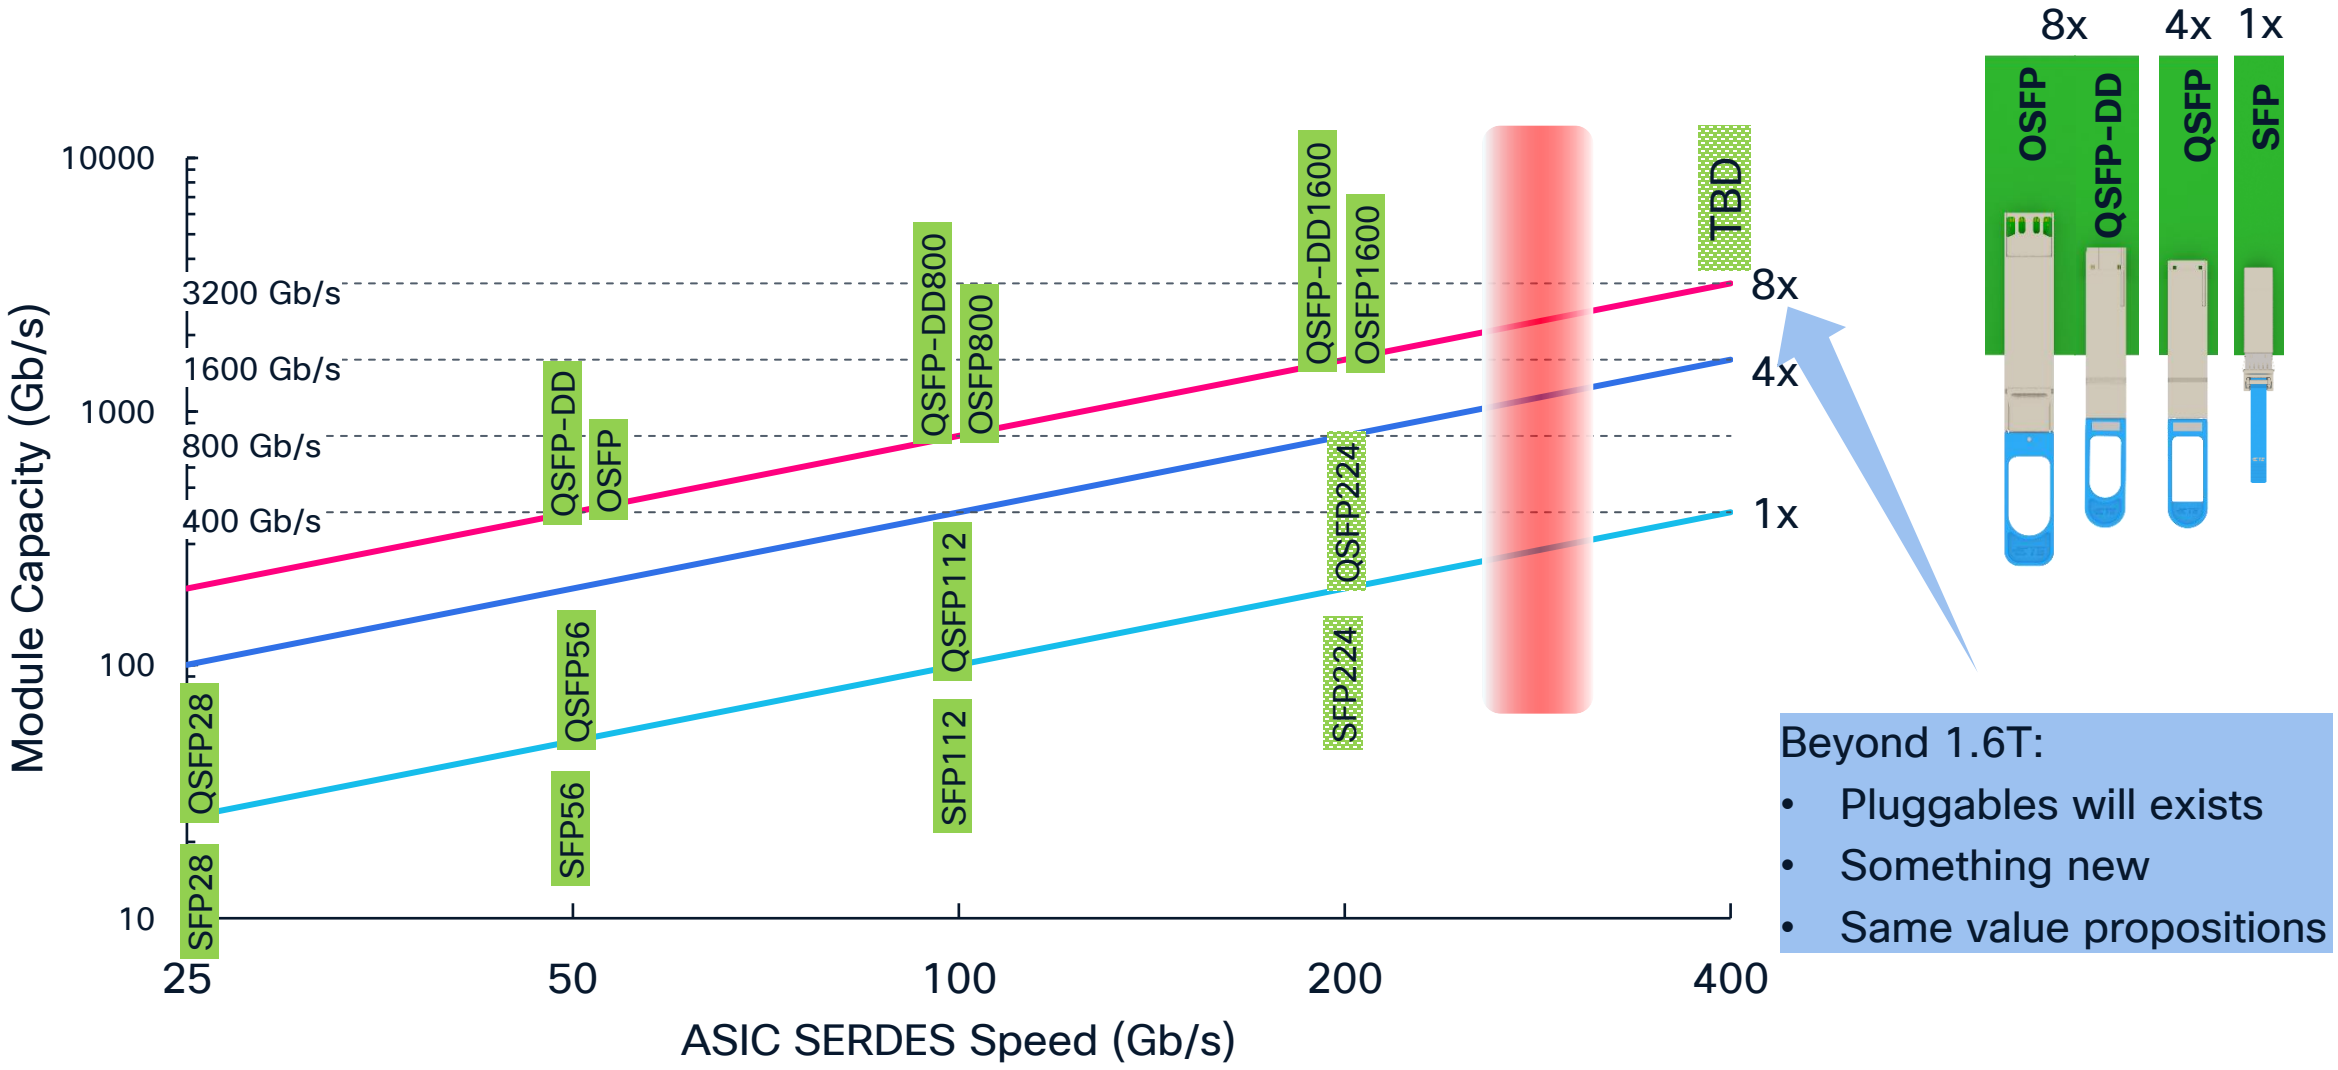

Beyond 1.6T:

- Pluggables will exists
- Something new
- Same value propositions

Alternative to pluggable: Co-packaged Optics

Co-packaged optics (CPO) and Linear Pluggable Optics (LPO) are two implementation variants of the same idea – reduce ASIC to optics power/DSP

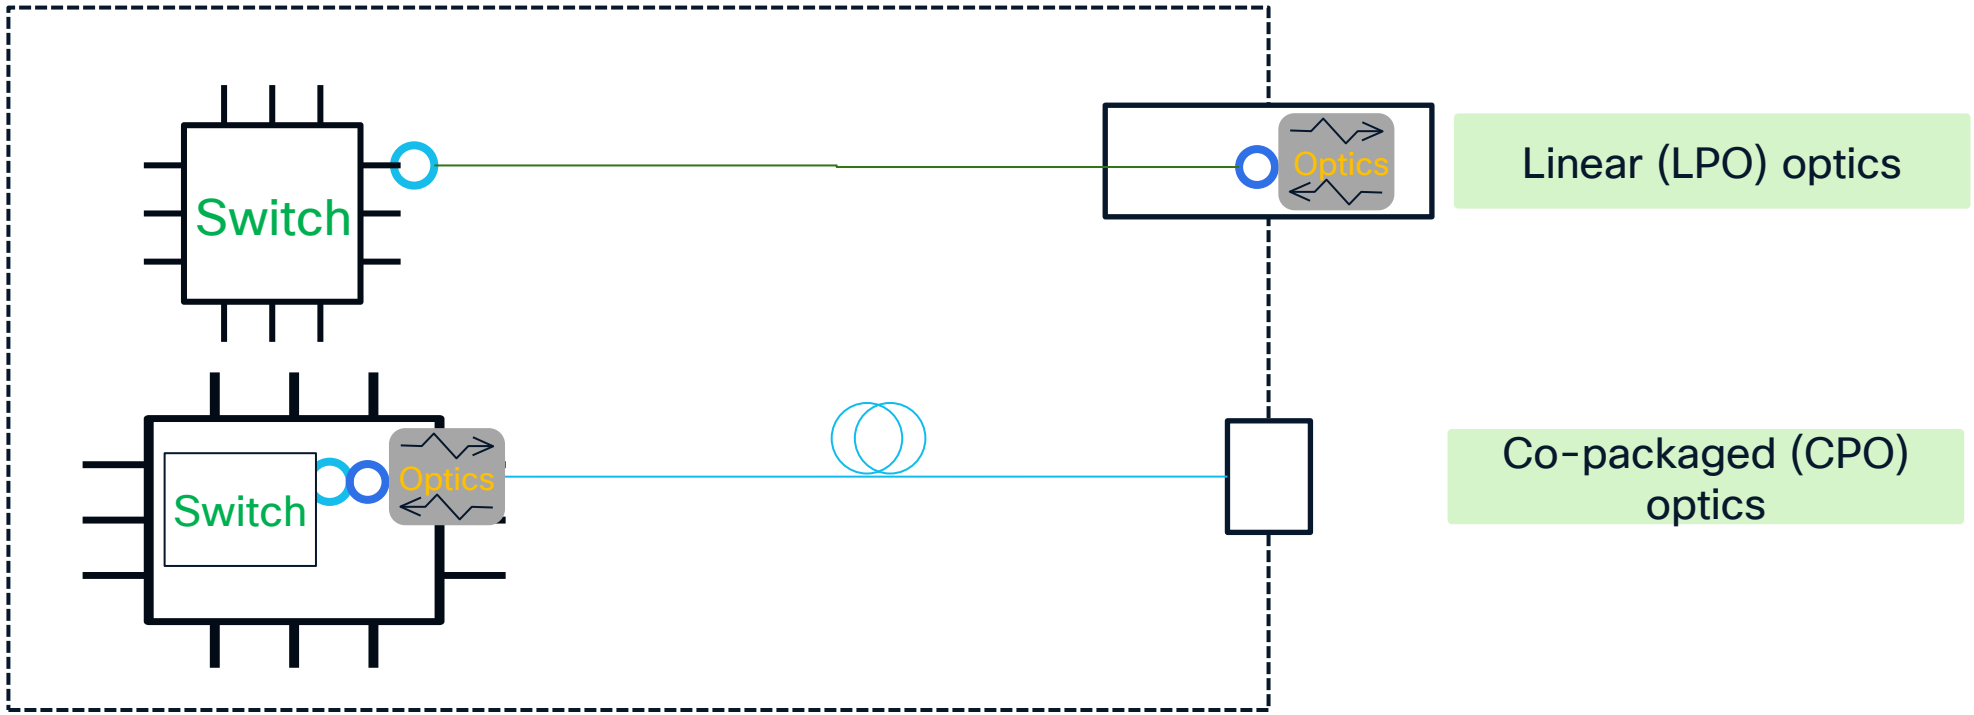

Pluggables and co-packaged optics with co-exist

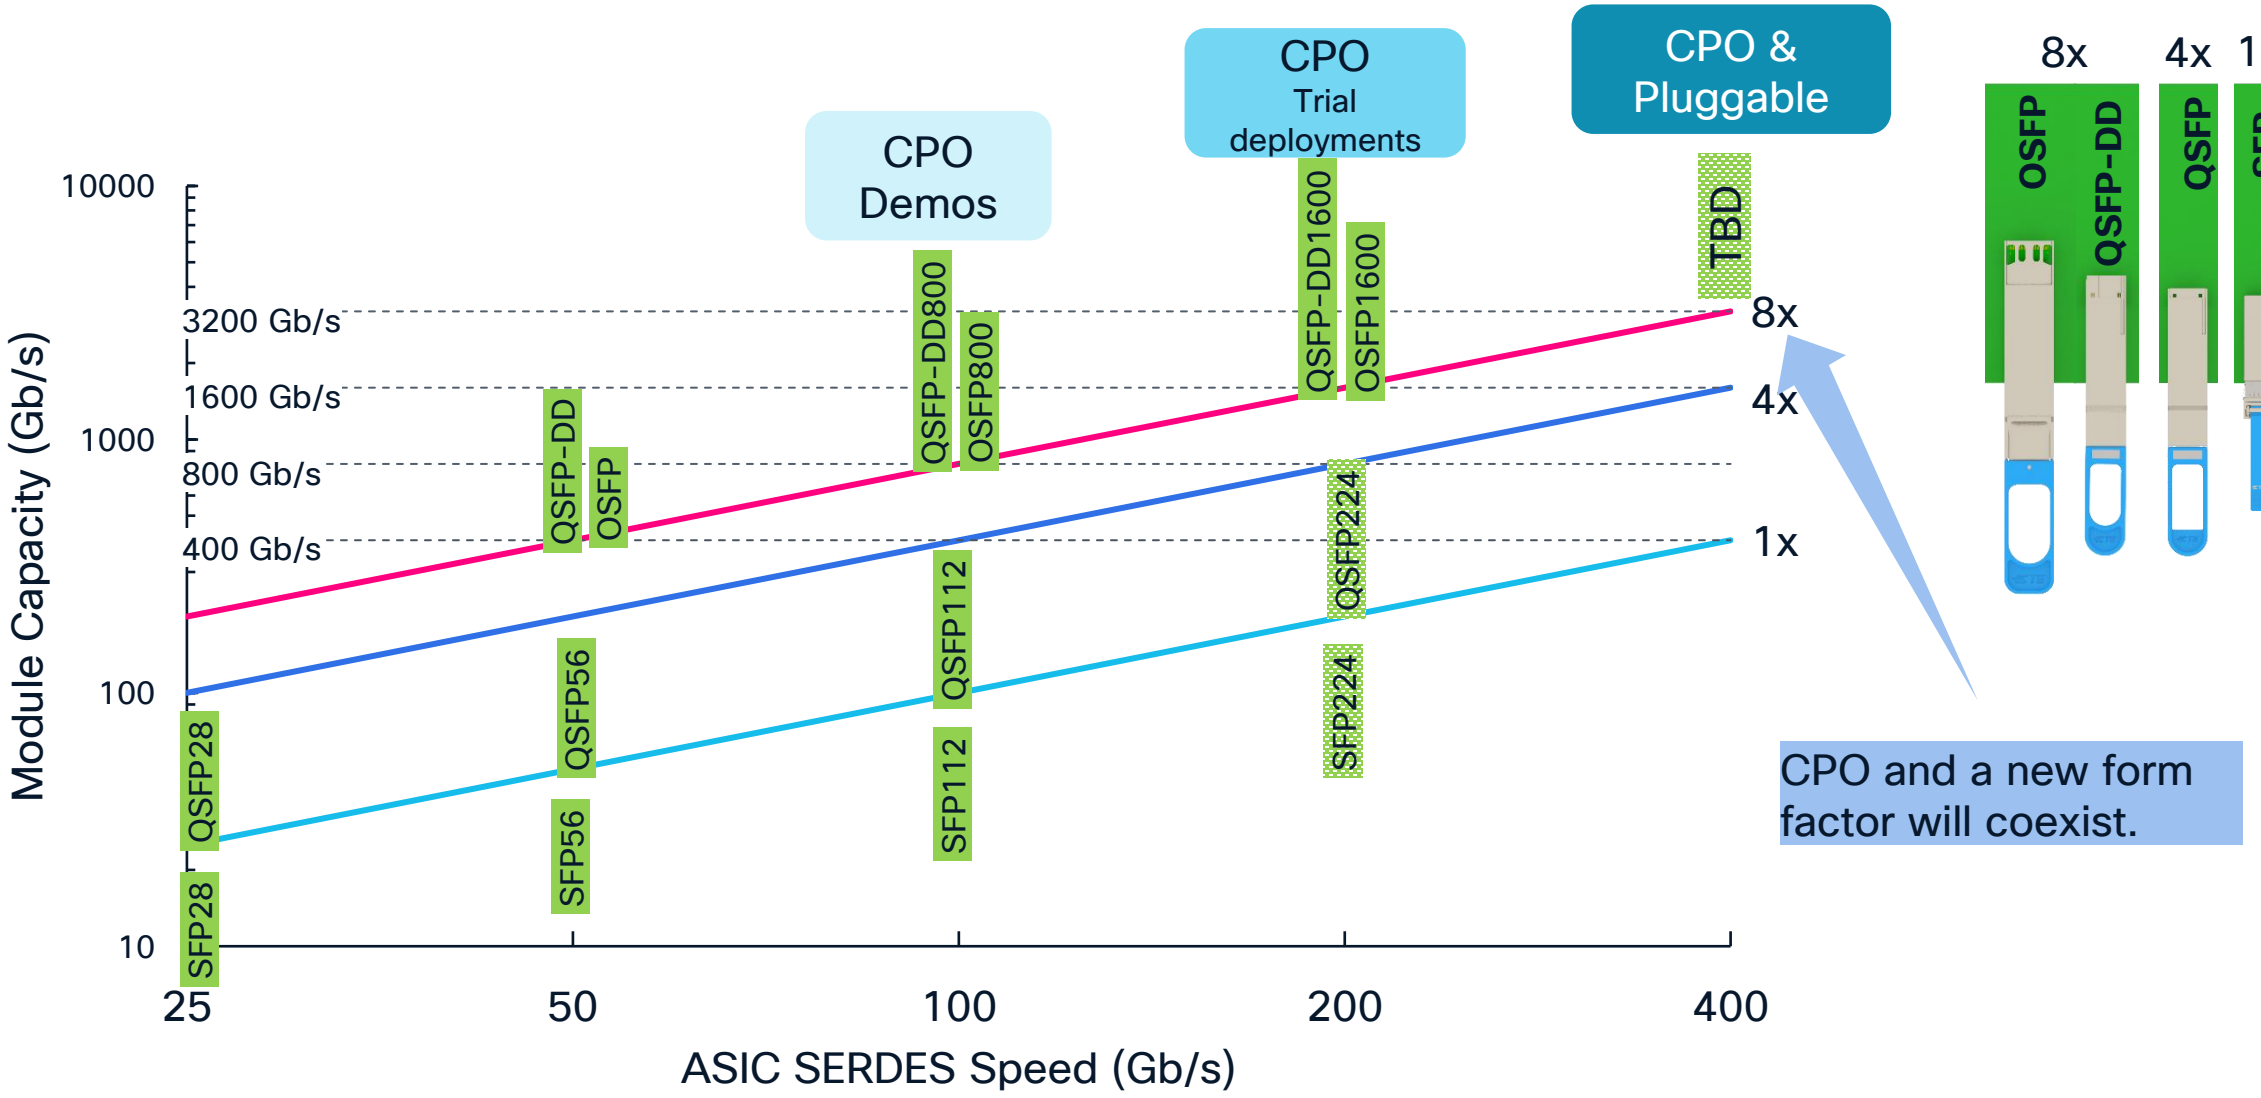

cisco

Summary: Moving the functionality

Retimed Optic
(majority today)

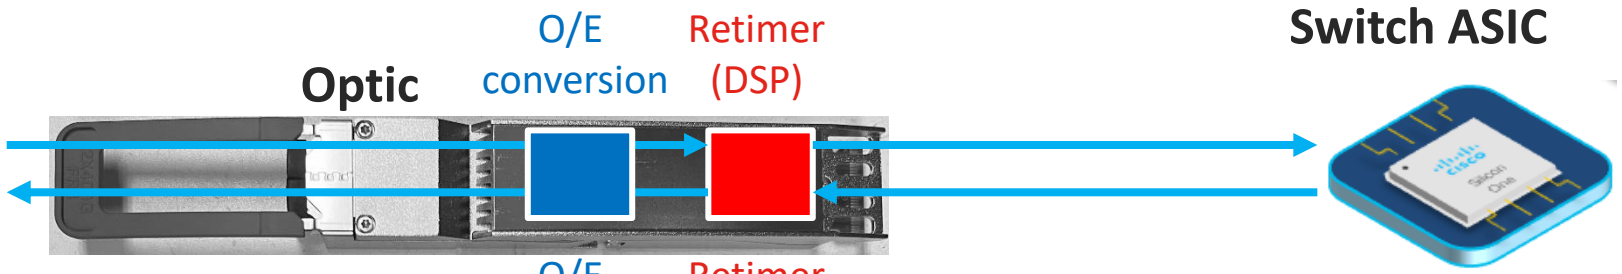

LRO (Linear receiver)
Half-retimed, TRO, RTLR

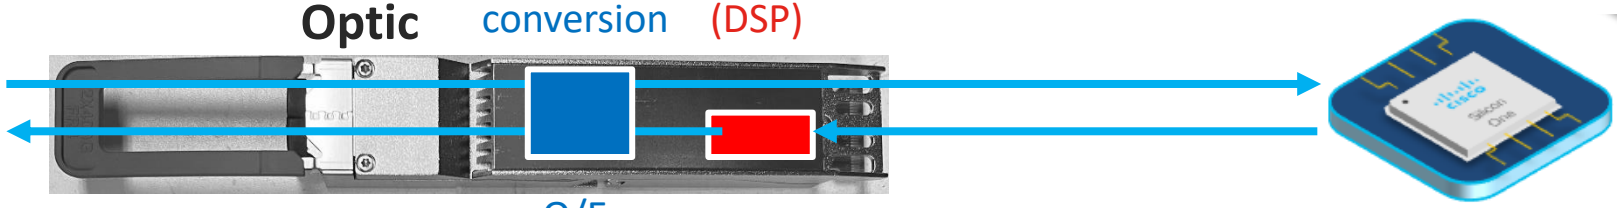

LPO (Fully Linear)

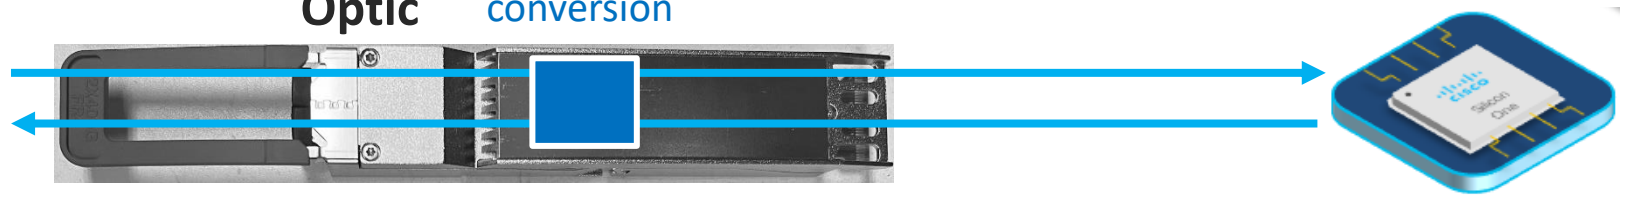

Co-packaged Optic (CPO)
(including laser)

What about interop? CPO, LPO, LRO, retimed

Link Interoperability is independent of implementation

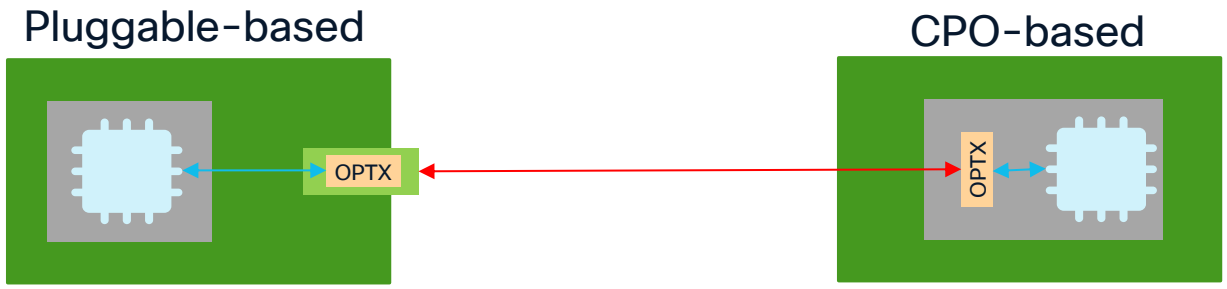

As long as the interface type matches (e.g. 800GBASE-DR8 or 800G-LPO-DR8) the link will interoperate.

- Relevant for pluggable types (OSFP, OSFP-RHS, QSFP-DD) and CPO

The one caveat: LRO optics (half-retimed) is intended to be able to interop with the retimed variant.

- @ 100G/lane, LPO works great and is lower power so little LRO
- @ 200G/lane, LRO is more likely to be deployed initially and should interoperate with retimed modules (but those specs aren't done yet)

Port & Lane Speed

Relentless ASIC Advancement

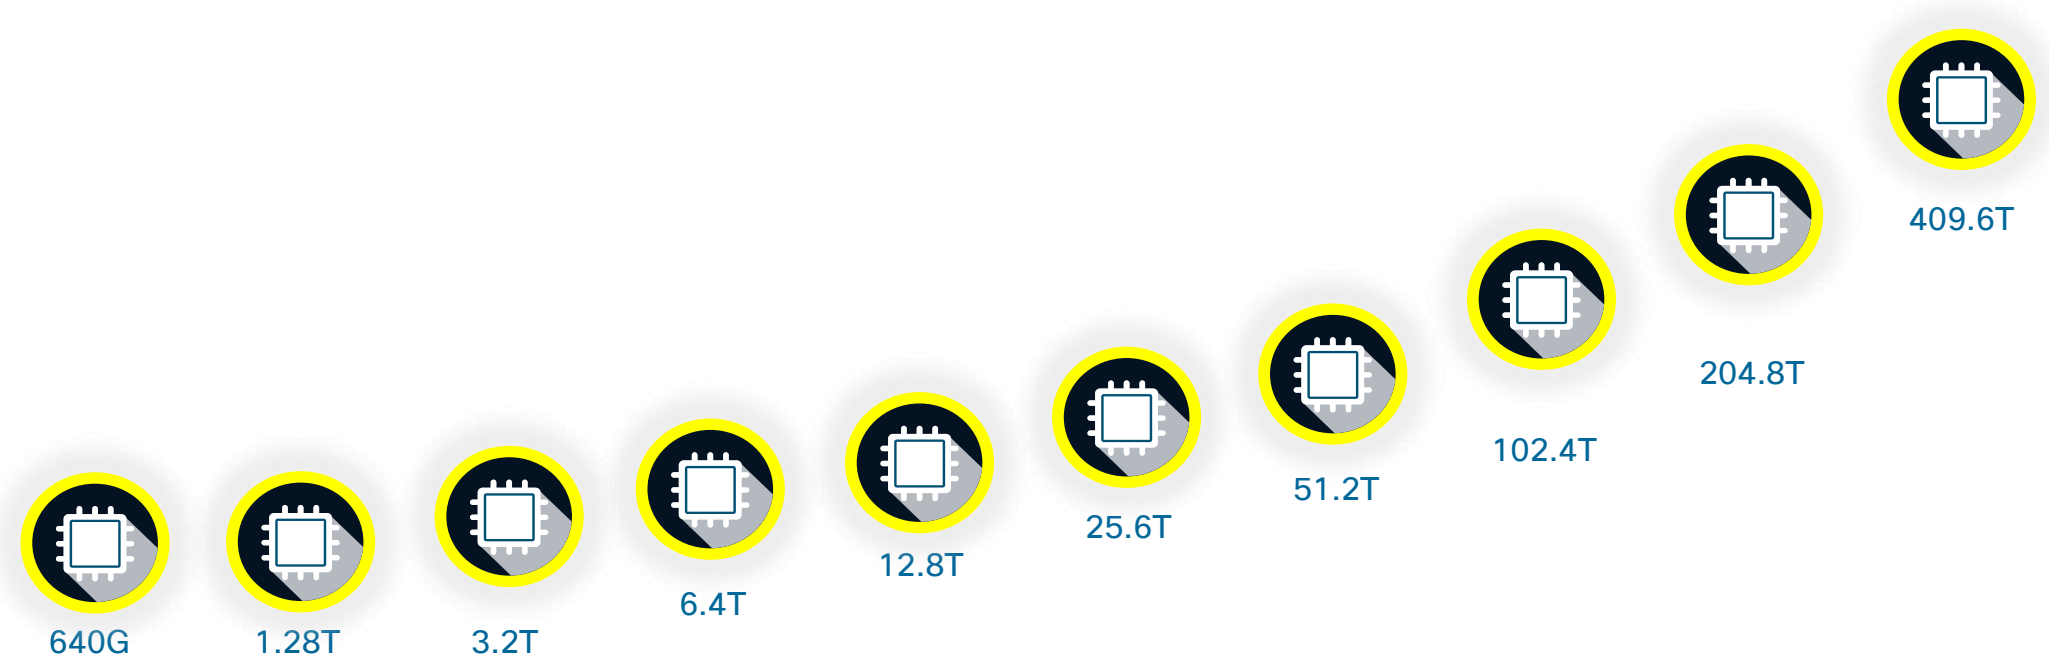

ASIC # SerDes	64			128			256			512					1024						
Serdes Gbps	10G NRZ				25G NRZ				50G NRZ				100G PAM4		200G PAM4				400G		
Port Gbps	40G				100G				400G				800G		1.6T				3.2T		
Ports/Silicon	16			32			64	32		64			64		64			128			

ASIC density continues to redefine how products are built.
SerDes & Interconnect. Optics & wavelengths.

Credit: adapted from slide by Rakesh Chopra & John Chapman
https://blogs.cisco.com/sp/co-packaged-optics-and-an-open-ecosystem

Next generation Ethernet's High-Speed Solutions – reuse, reuse, reuse

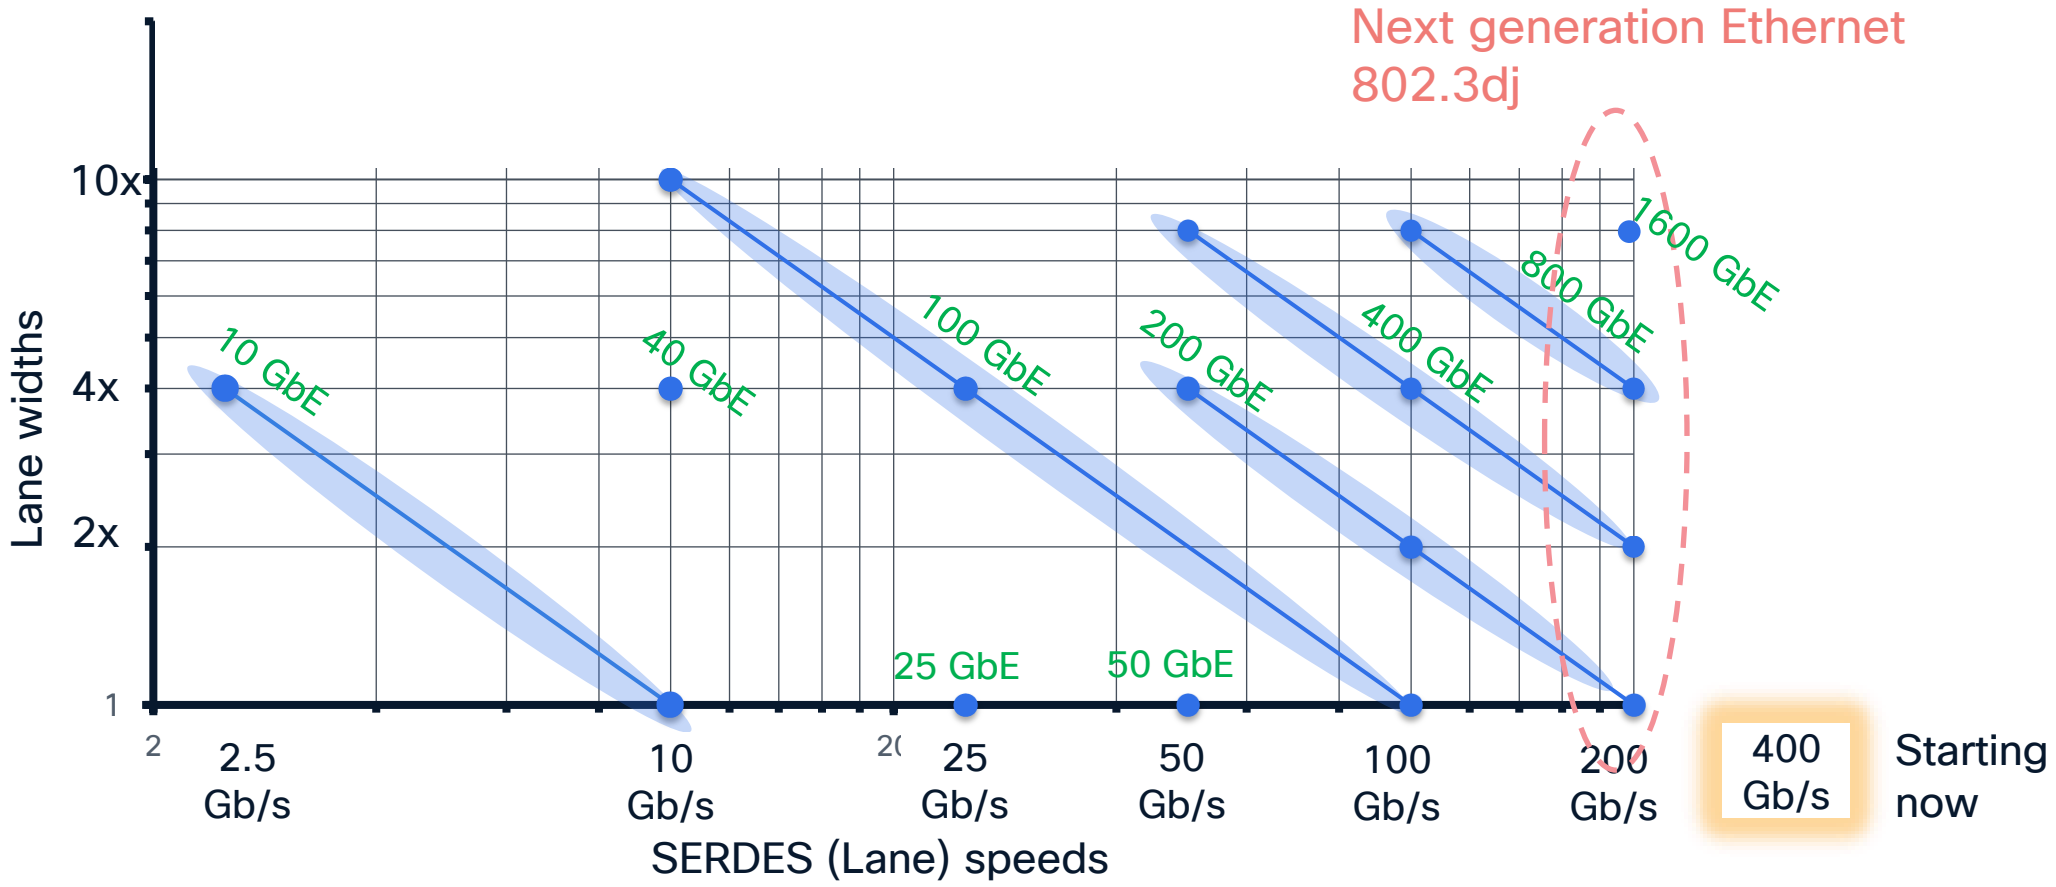

Ethernet's next generation: IEEE 802.3 P802.3dj → 200 Gb/s per lane for 1.6 TbE and below

Key themes :

- **200 GbE, 400GbE, 800GbE and 1.6TbE Ethernet interfaces** (N x 200Gb/s)
  - Consideration of Breakout use-case
- **Electrical interfaces based on 200 Gb/s**
  - ➢Copper Cables
  - ➢Backplanes
  - ➢Chip to module (AUI) interfaces and Chip to chip interfaces
- **Optical Interfaces**
  - ➢Only single mode fiber (Duplex and Parallel)
  - ➢500m to 10km using 200 Gb/s per lane IMDD technology
  - ➢10km to 40 km using 800 Gb/s per lane (16QAM) coherent technology

What will next generation Ethernet provide @ 200G/lane?

Continued flexibility of implementations

- Copper cables
- Multimode Fiber – 30/50 m (just started)
- Single Mode Fiber inside DC – 500m & 2km
- Single Mode Fiber Campus – 10 km
- Outside plant, DCI – up to 40 km

Compatibility with known implementations was a strong factor in the Ethernet project's decisions

- Compatible with current CMOS
- Compatible with popular system configs or pluggable modules

Next generation Ethernet will maintain Ethernet's backwards compatibility

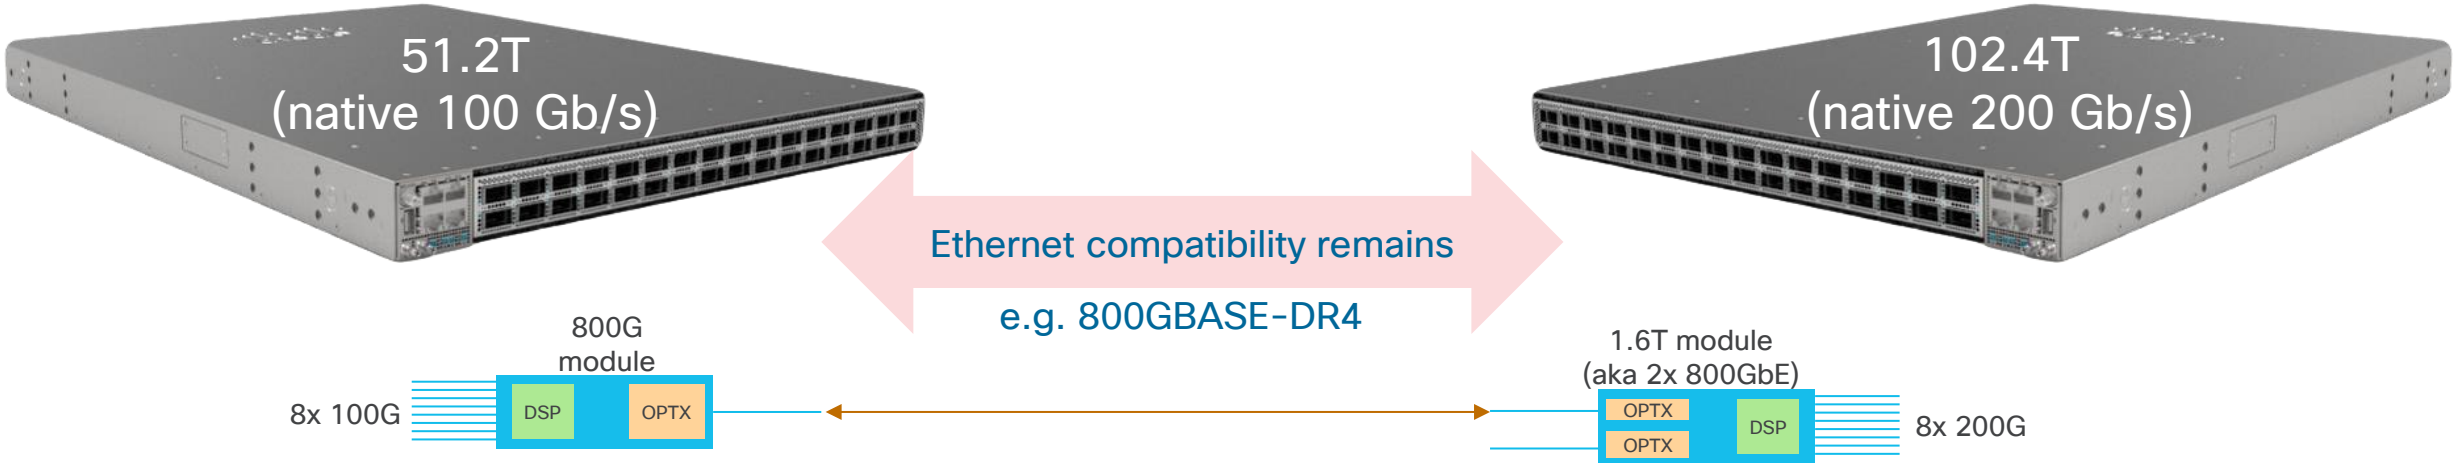

IEEE 802.3 P802.3dj's logic and architecture ensures compatibility with "legacy" systems and implementations.

Any 800 GbE capable port will operate with any other 800 GbE capable port:

- As long as the same PMD interface is used
- For example, 800GBASE-FR4 in both ports

Key elements of 200 Gb/s technology

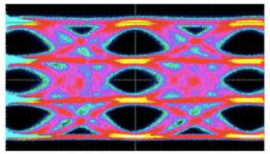

Compatible with current CMOS, SiPhotonics, EML, Cable, Fiber technology

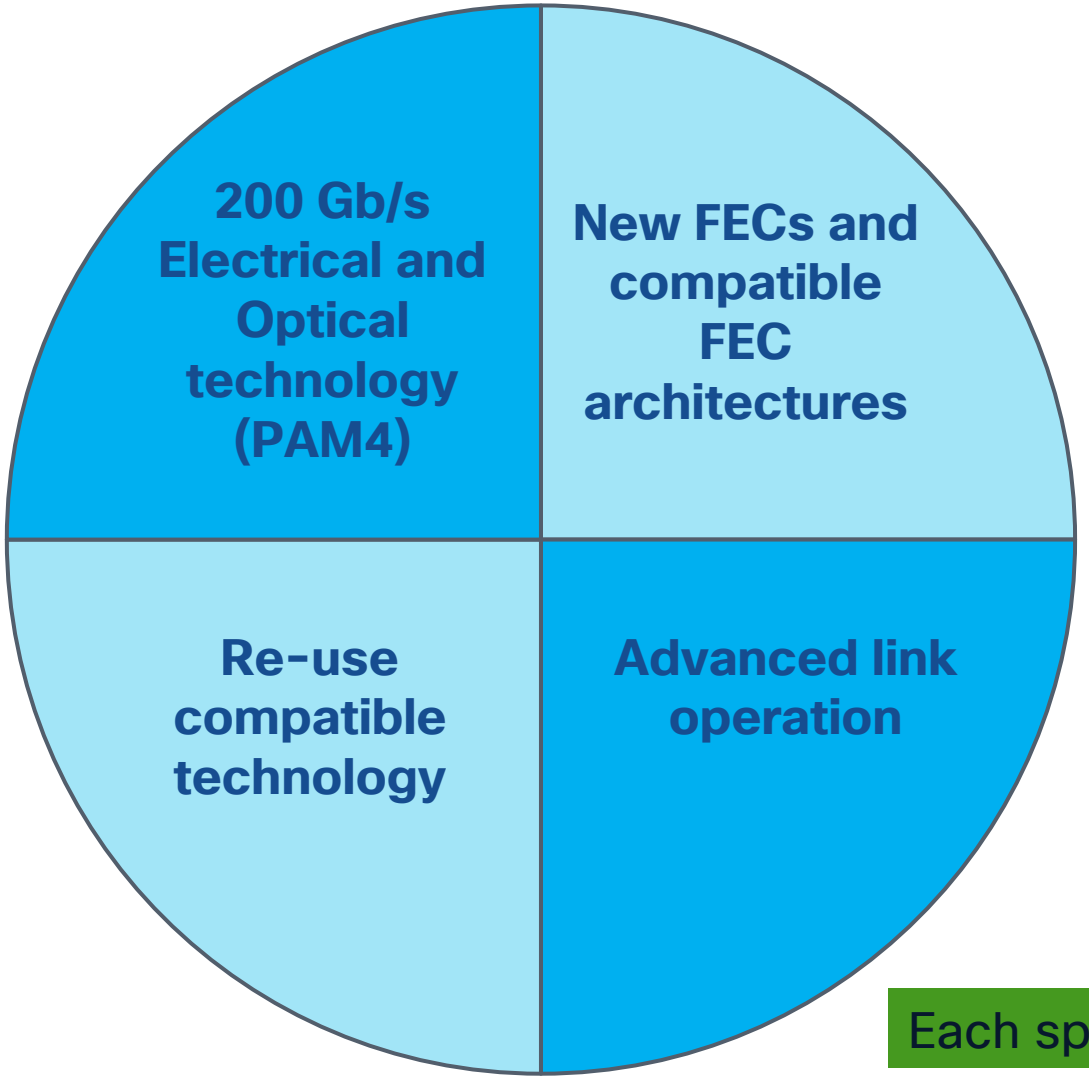

Multiple FEC schemes for different reaches. Common architecture to enable compatible implementations

Implementations can use same technology for 8x, 4x, 2x or 1x variants

New (optional) Optical Link Training and proposed Auto-negotiation

Each speed step is more challenging!

Different reaches need different FECs

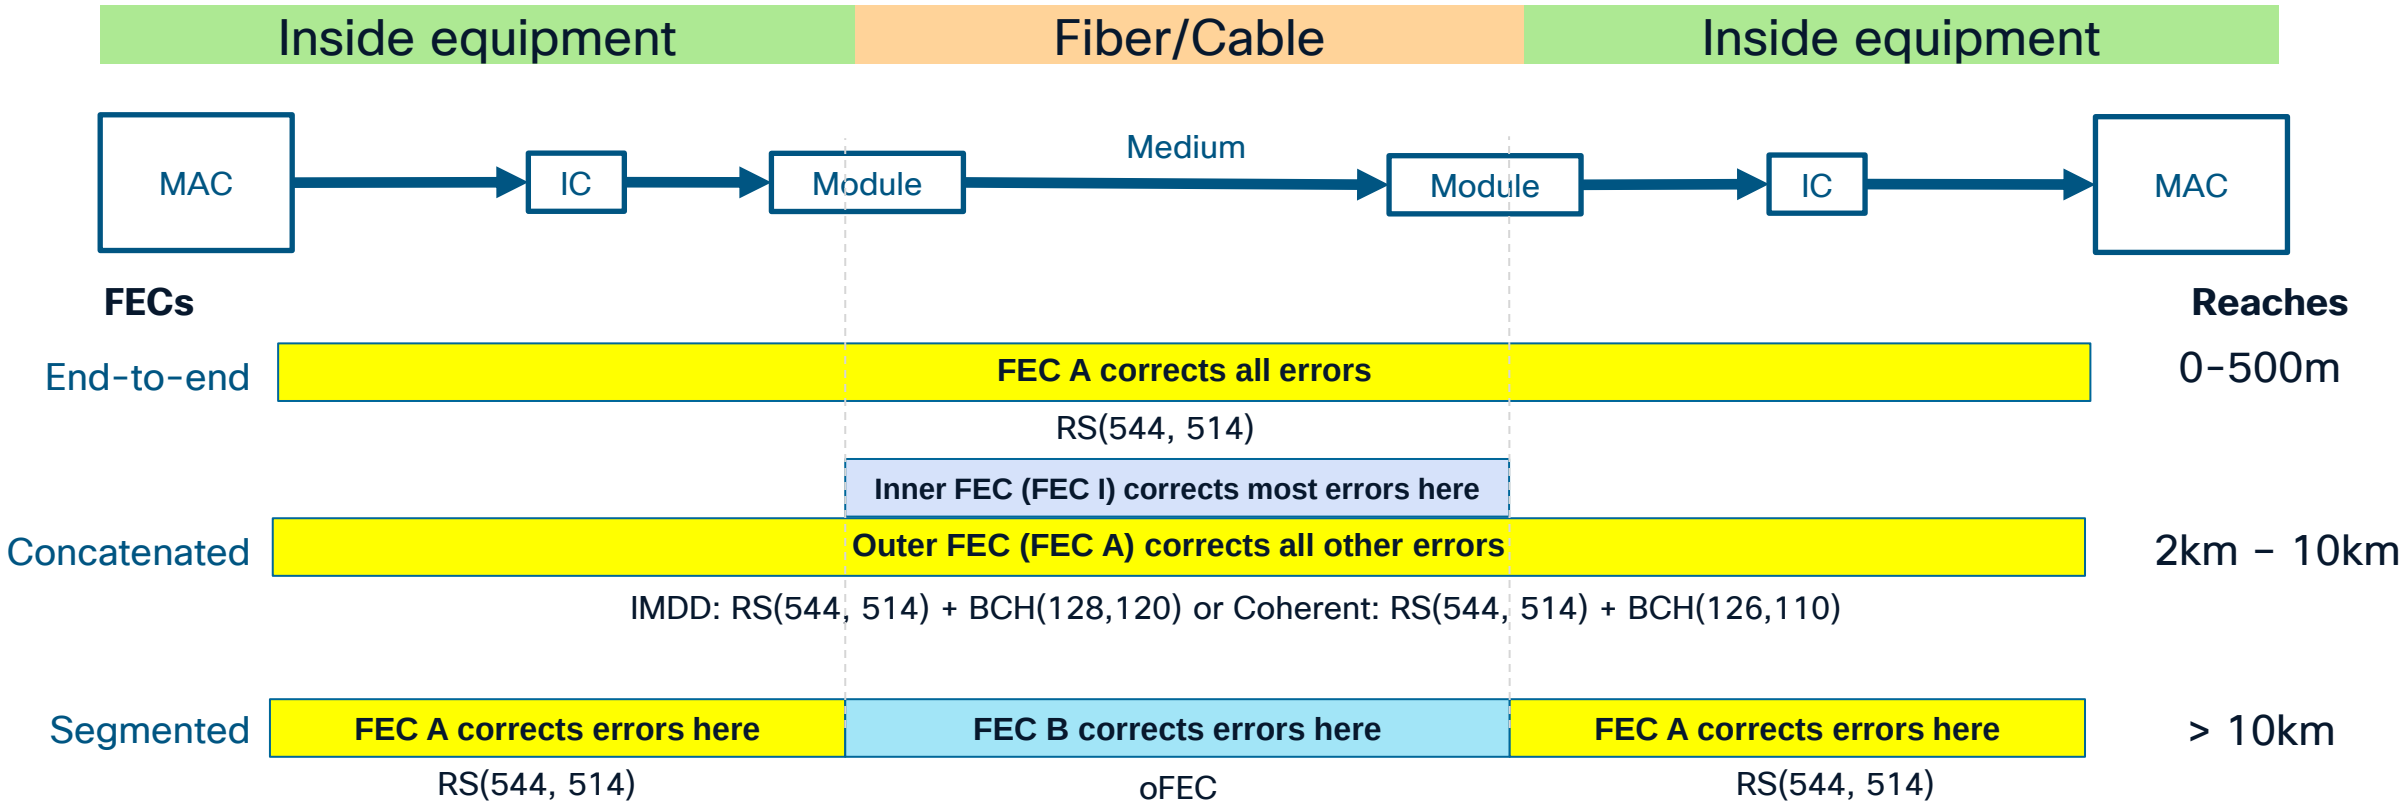

They key point:

- The switch platform will always have integrated FEC
- Modules may have additional FEC to reach a longer distance

Optics for AI Infrastructure

Requirements around AI are driving interconnect technology

- New network **architectures** are being adopted to optimize for AI Infrastructure
  - Scale Up, Scale Out, Scale Across
- Faster **speeds**:
  - 100 Gb/s → 200 Gb/s → ...
  - 800 GbE → 1.6 TbE → ...
- Lowering **power** is key
  - Copper cables still live @ 200 Gb/s
  - Novel implementations like LPO / CPO
  - Faster speeds yield better energy efficiency (W/Gb)
- Heightened focus on **reliability**
  - Core Cisco strength

AI networking will dominate the future optics design priorities for data centers

NEW Q3'26

Reliable connectivity for AI data centers

Cisco OSFP 1.6T Pluggable Optics

Powered by **Cisco Silicon Photonics**

High-performance and reliability
Advanced DSP-based diagnostics and monitoring

High bandwidth AI scale-out

1.6T switch-to-NIC connectivity and high-density fabrics for massive AI cluster builds (2x800G, 4x400G, and 8x200G links)

Diversified supply chain

Seamless AI cluster scaling with diversified supply chain and interoperability across Cisco and 3rd party DSPs

Efficiency and flexibility

Optics for power-efficiency and deployment flexibility in liquid cooled and air-cooled platforms

Wrap up

Key takeaways – thank you for your attention

Pluggable optics are crucial for today's needs and for the AI era. They continue to be:

- Flexible, Scalable, Performance, Reliable
- IMDD & Coherent optics broadly deployable for all use cases

The industry is actively exploring innovative solutions for further optimization for AI's unique demands:

- Co-packaged optics, Linear pluggable optics, Silicon Photonics
- The future will be a mix of these technologies as AI infrastructure continues to evolve

Much **tighter co-design** between ASIC, System Architecture, System design, and Optics is necessary to deliver more efficient solutions

Cisco continues to lead the optics evolution building on our ASICs, Systems and SiPhotonics strengths

Simplify AI with Cisco Optics tools

Optics Selector

https://optsel.cisco.com

Optics-to-Device Compatibility

https://tmgmatrix.cisco.com/

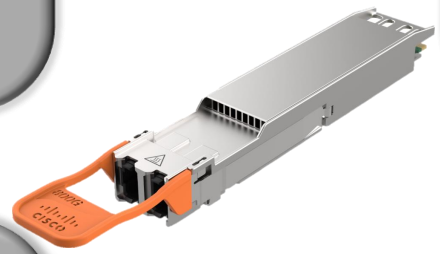

Optics-to-Optics Interoperability

https://tmgmatrix.cisco.com/iop

Product Information (Dynamic Datasheet)

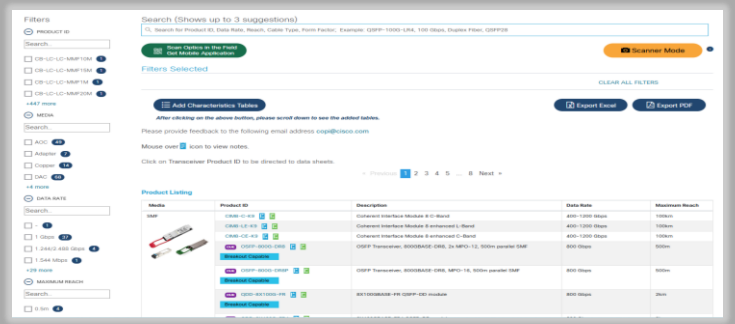

https://copi.cisco.com/

Complete your session surveys

Complete your surveys in the Cisco Events App.

Complete a minimum of 4 session surveys and the overall event survey to receive a unique Cisco Live t-shirt.
(from 11:30 on Thursday, while supplies last)

Continue your education

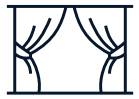

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

Visit the Technical Solutions Clinics to discuss your technical questions

Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs

Visit the On-Demand Library for more sessions at CiscoLive.com/On-Demand

Thank you

cisco Live !


CISCO