400G, 800G, and Terabit Pluggable Optics:

What You Need to Know

Mark Nowell
Cisco Fellow

Errol Roberts
Distinguished Architect

BRKOPT-2699

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile 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 June 13, 2025.

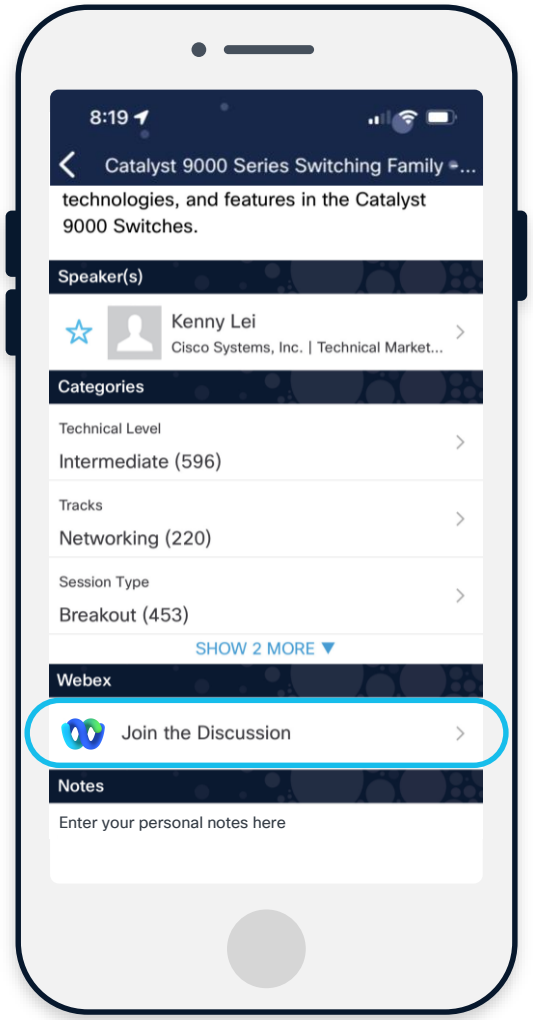

Are pluggable optics dead or alive for the AI era?

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?

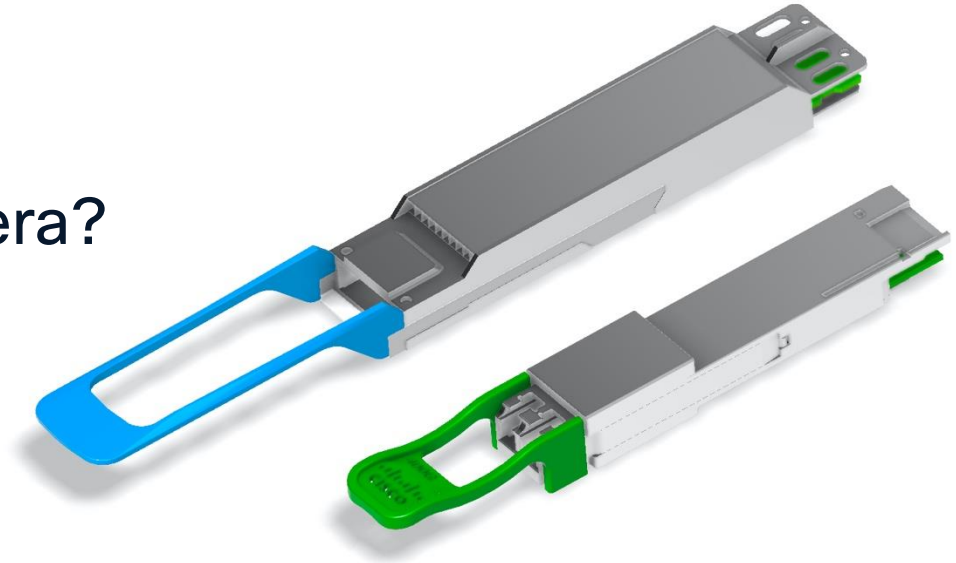

AI Infrastructure & Pluggable Optics

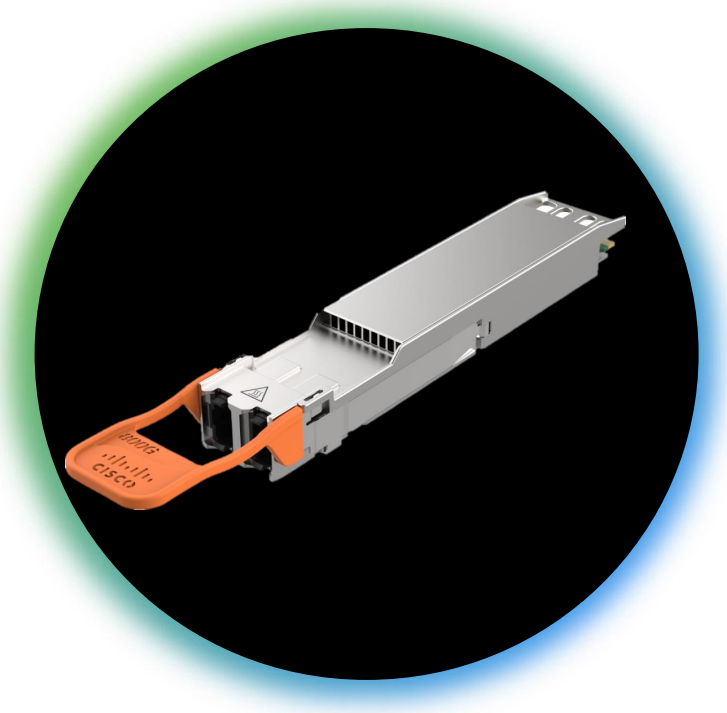

- AI infrastructure demands are a disruptive transformation
- AI is driving innovations in pluggable optics
- Demands higher – reliability, performance, integration
- Cisco leadership in SiPhotonics addressing AI infrastructure demands

Agenda

01 Market Dynamics

02 400G & 800G Pluggable IMDD and Coherent Status

03 Beyond 800G – how is AI/ML impacting interconnect

04 Wrap Up

Market dynamics

Ethernet Data Center Switch Port Speed Transitions

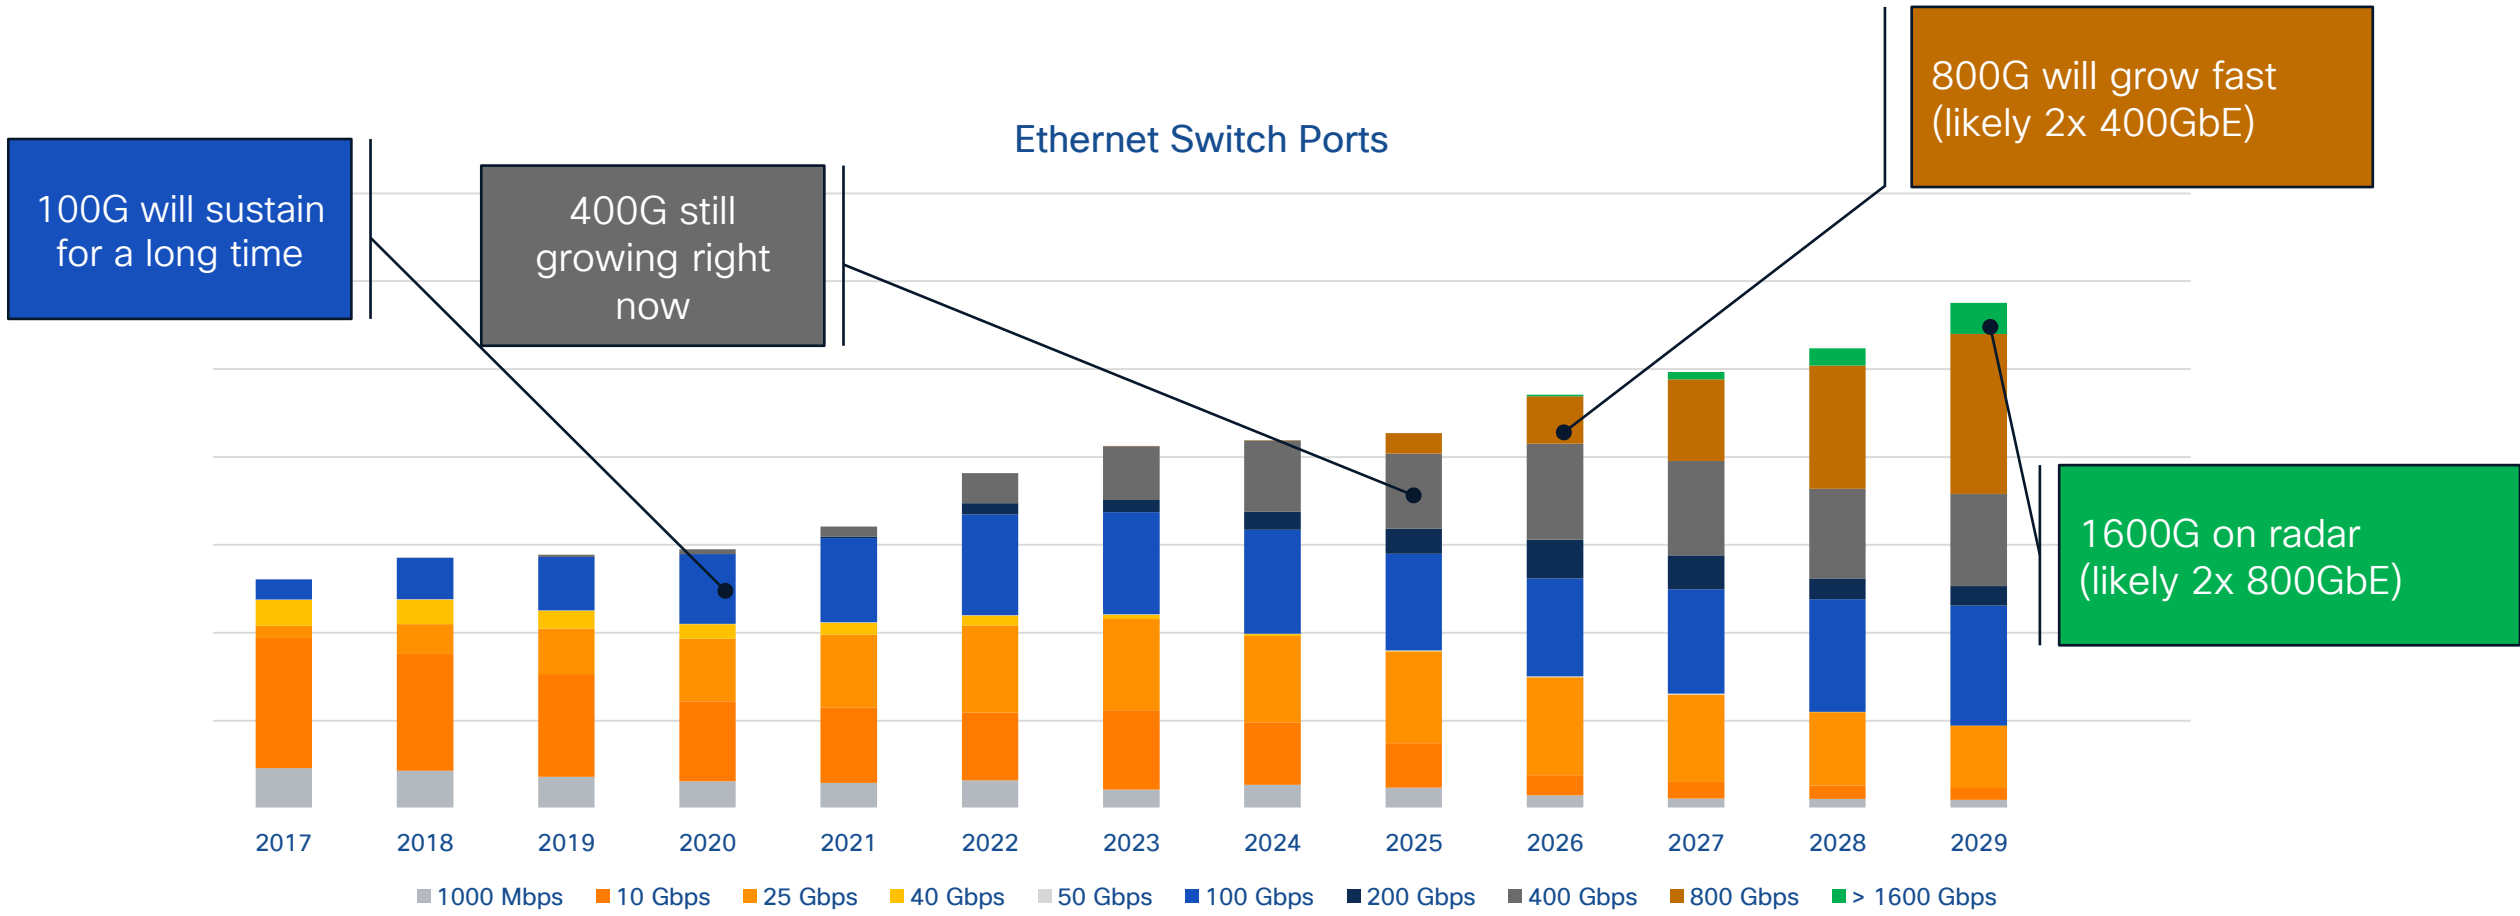

- Majority of the highest speed transitions are webscale (top 8) customers
- Webscale will drive the speed transitions quickly to scale
- Rest of the market will leverage that scale in their own time frame

Source: Dell'Oro's Ethernet Switch – Data Center Five Year Forecast Report 2025 – 2029

Speed evolution in a data center

100% pluggable optics

Inter-Data Center

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

Switch fabric

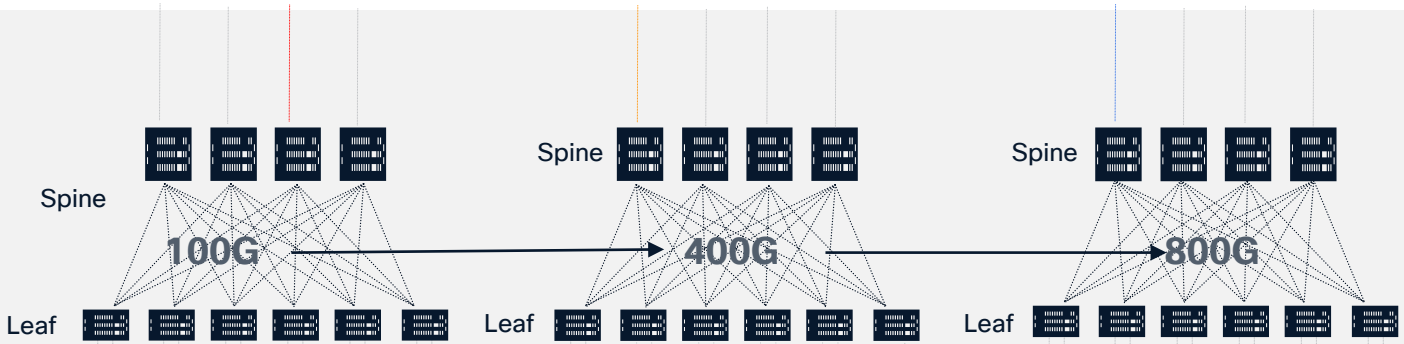

- Switch silicon bandwidth growing due to higher Radix and faster Serdes speeds
- Switch ASIC throughput growing: 6.4 Tbps to 12.8 Tbps to 25.6 Tbps to 51.2 Tbps
- Optics increasing from 40Gbps to 100G Gbps to 400Gbps to 800Gbps

Access

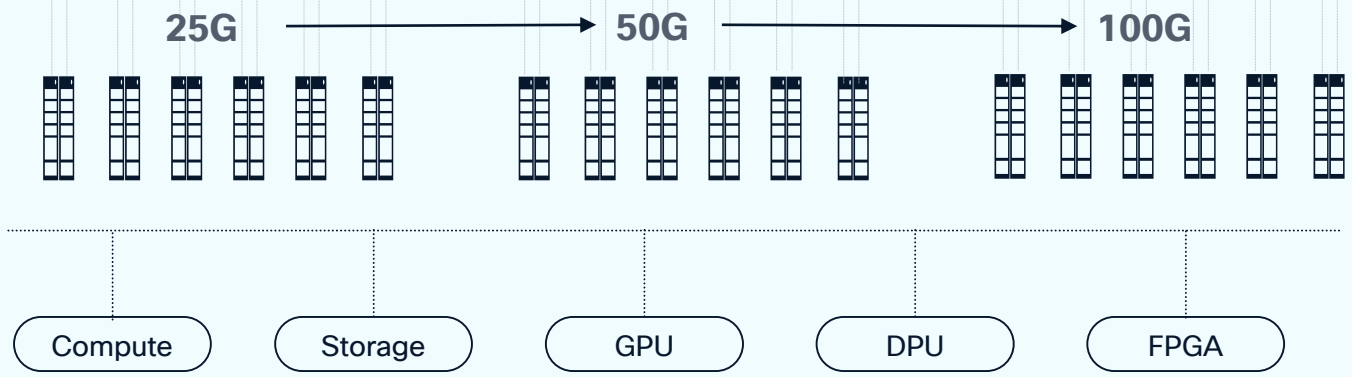

- Server network connectivity evolves with server processor upgrade cycles as data center traffic grows
- Server port speed is transitioning from 1/10 Gbps to 25 Gbps to 100 Gbps
- Storage, GPU, DPU, FPGA driving connectivity bandwidth, PCIe speed increase

Ethernet Speed Transitions in AI Networks

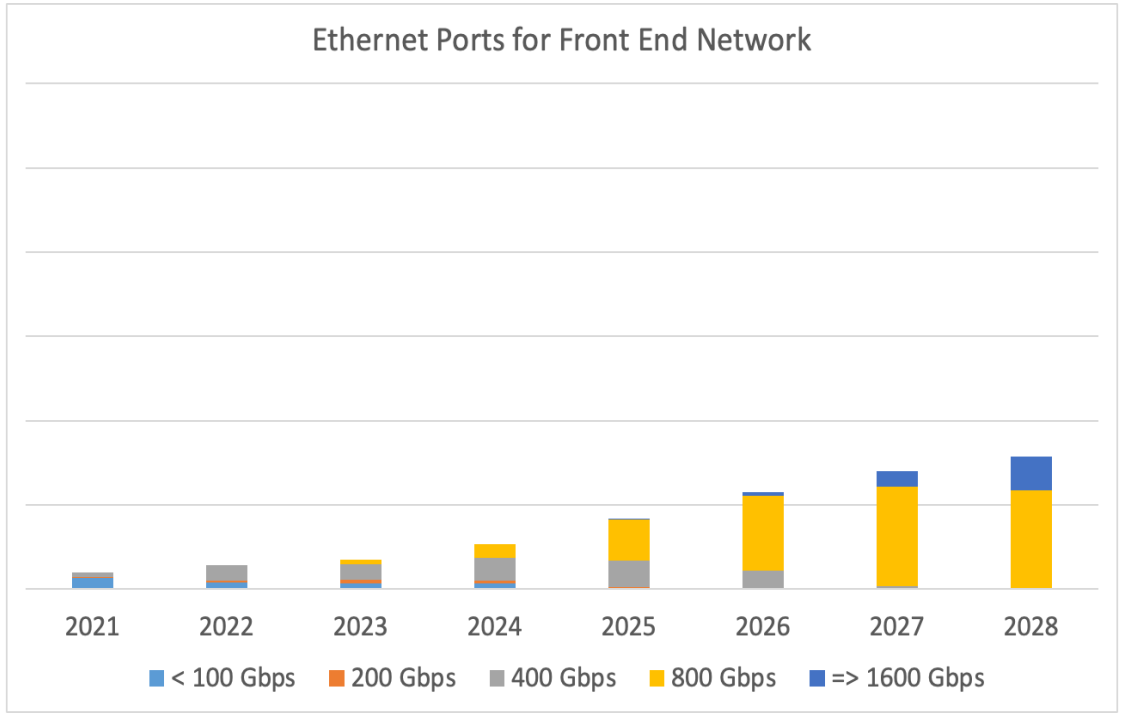

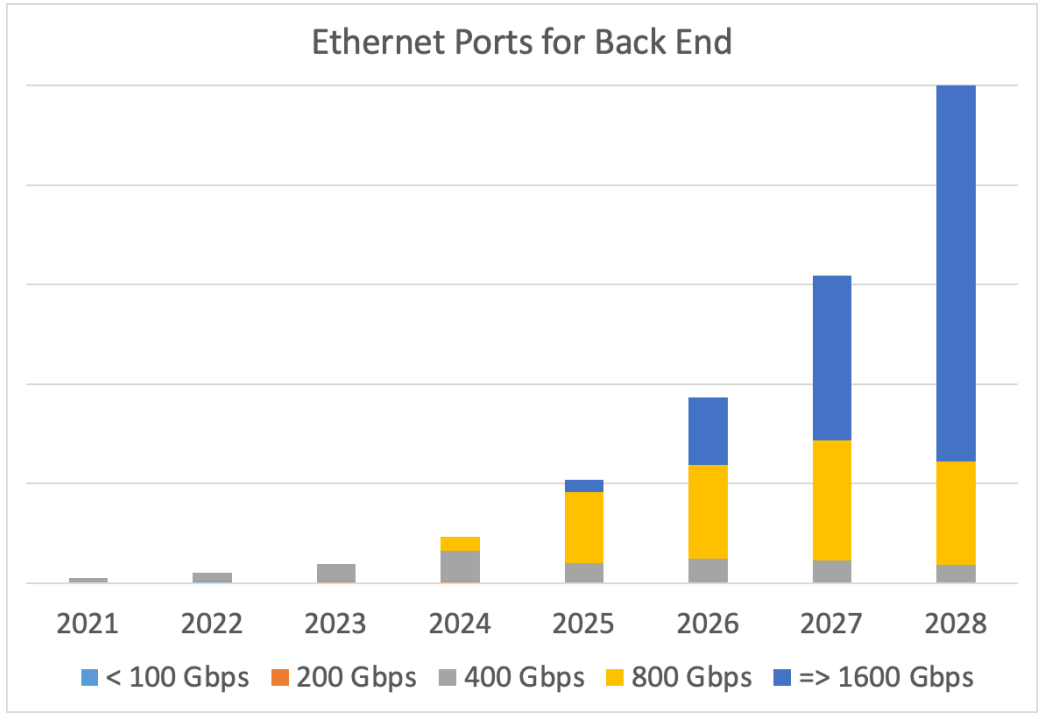

Majority of the switch ports in AI back-end Networks to be 800 Gbps in 2025 and 1600 Gbps in 2027, showing a very fast migration to the highest speeds available in the market.

Dell'Oro: 95N01_Advanced_Research_AI_Networks_For_AI_Workloads_Report_3Q24

AI Network Type Fundamentals

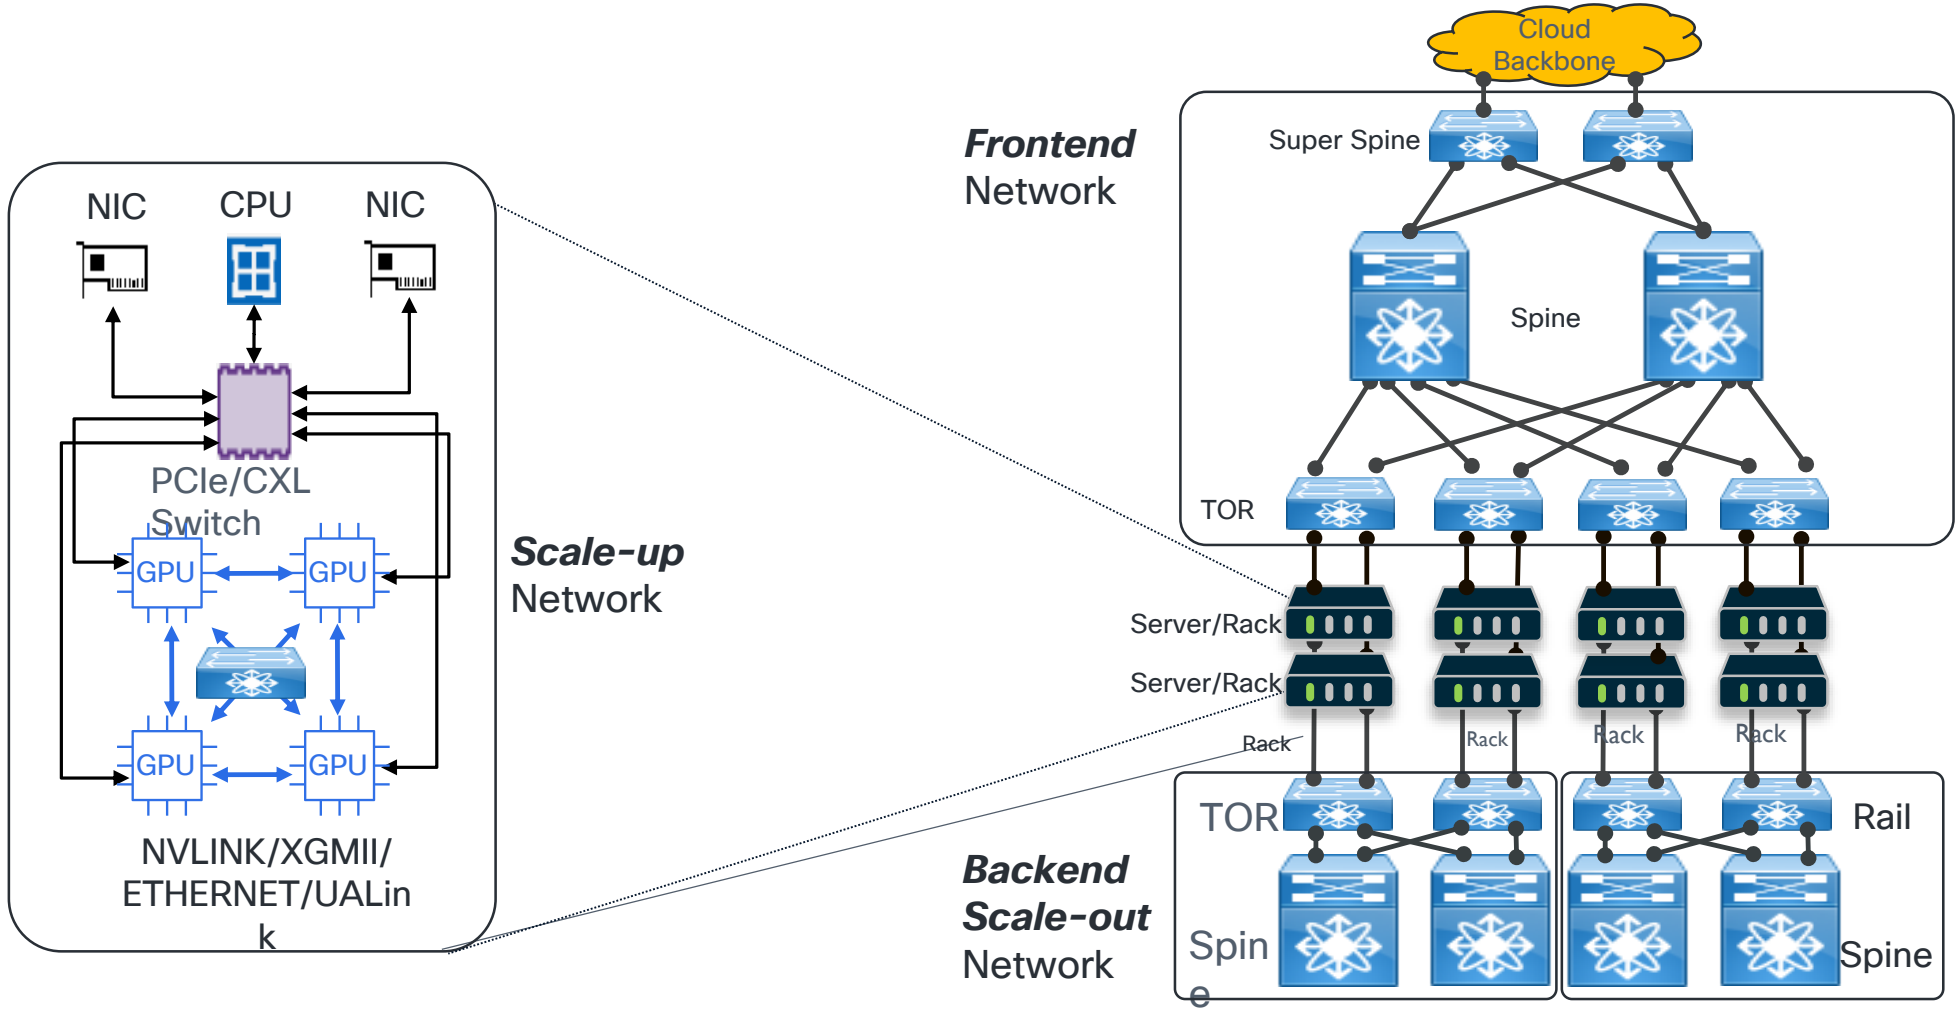

AI/ML is a disruptive event for equipment design

		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

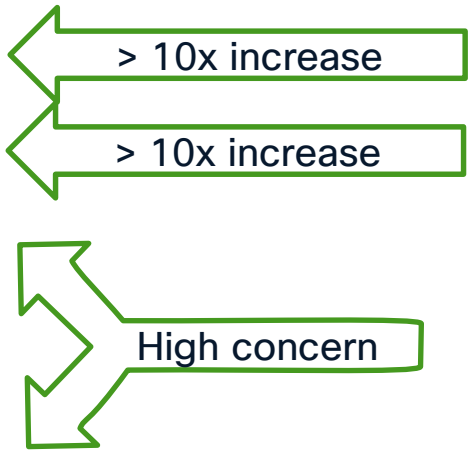

These challenges are forcing innovation to happen at all levels, including pluggable modules. But pluggable modules still offer:

- 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

Cloud & GPU service providers

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

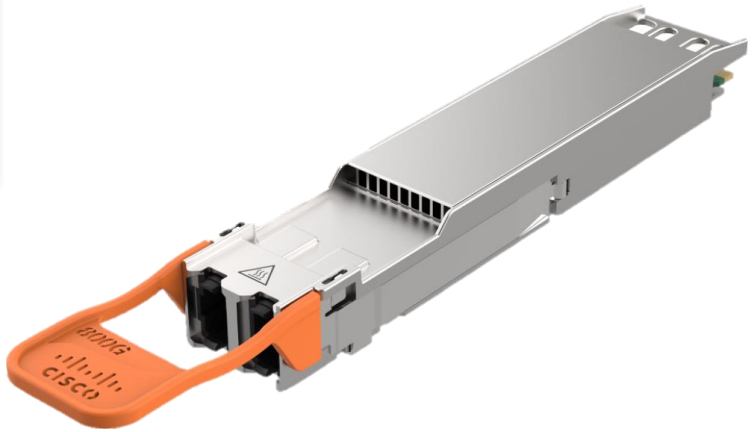

High speed and ultra-long haul
coherent pluggable optics

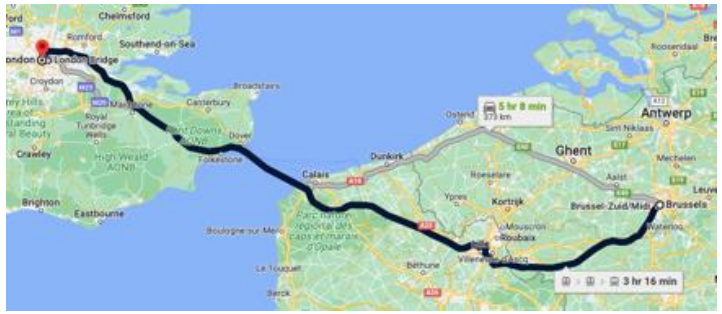

Optics firmware updates onsite
Leadership in standards development

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, 800GbE	Mature standards	400G (OIF, Open ZR+)
100 Gb/s	Mature technology	400ZR: 16QAM @ 60 Gbd 800ZR: 16QAM @ 118 Gbd
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 ✓		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)
	OIF400ZR ✓	cisco	400 GbE Coherent 120km
	OIF800ZR		800 GbE Coherent 120km
	OIF1600ZR		800 GbE Coherent 120km
	Open ZR+ ✓	cisco	Long Distance 400G
	CMIS	cisco	
	802.3df ✓	cisco	200G/400G/, 800G Ethernet Task Force @ 100Gb/s per lane
	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
	OSFP MSA ✓	cisco	400G/800G/1.6T Form factor
	400G-BiDi MSA ✓		400 GbE MMF BiDi

Key elements of 400G & 800G optics

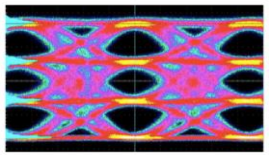

New Pluggables with 8-wide connector to support 400G and 800G ports
Multiple form factors

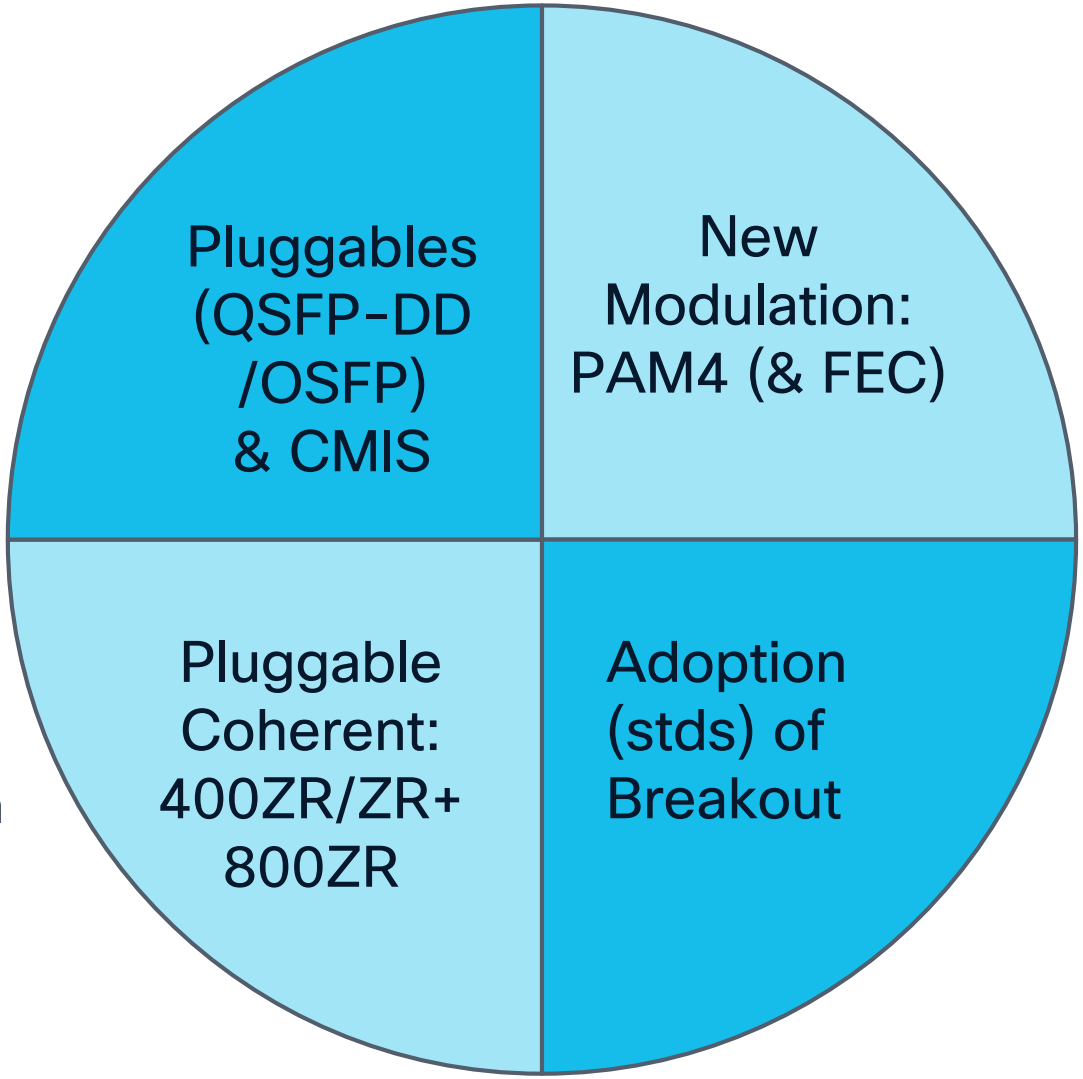

Higher speed interfaces adopted PAM4 modulation.
Ubiquitous use of FEC. & continued FEC innovation

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

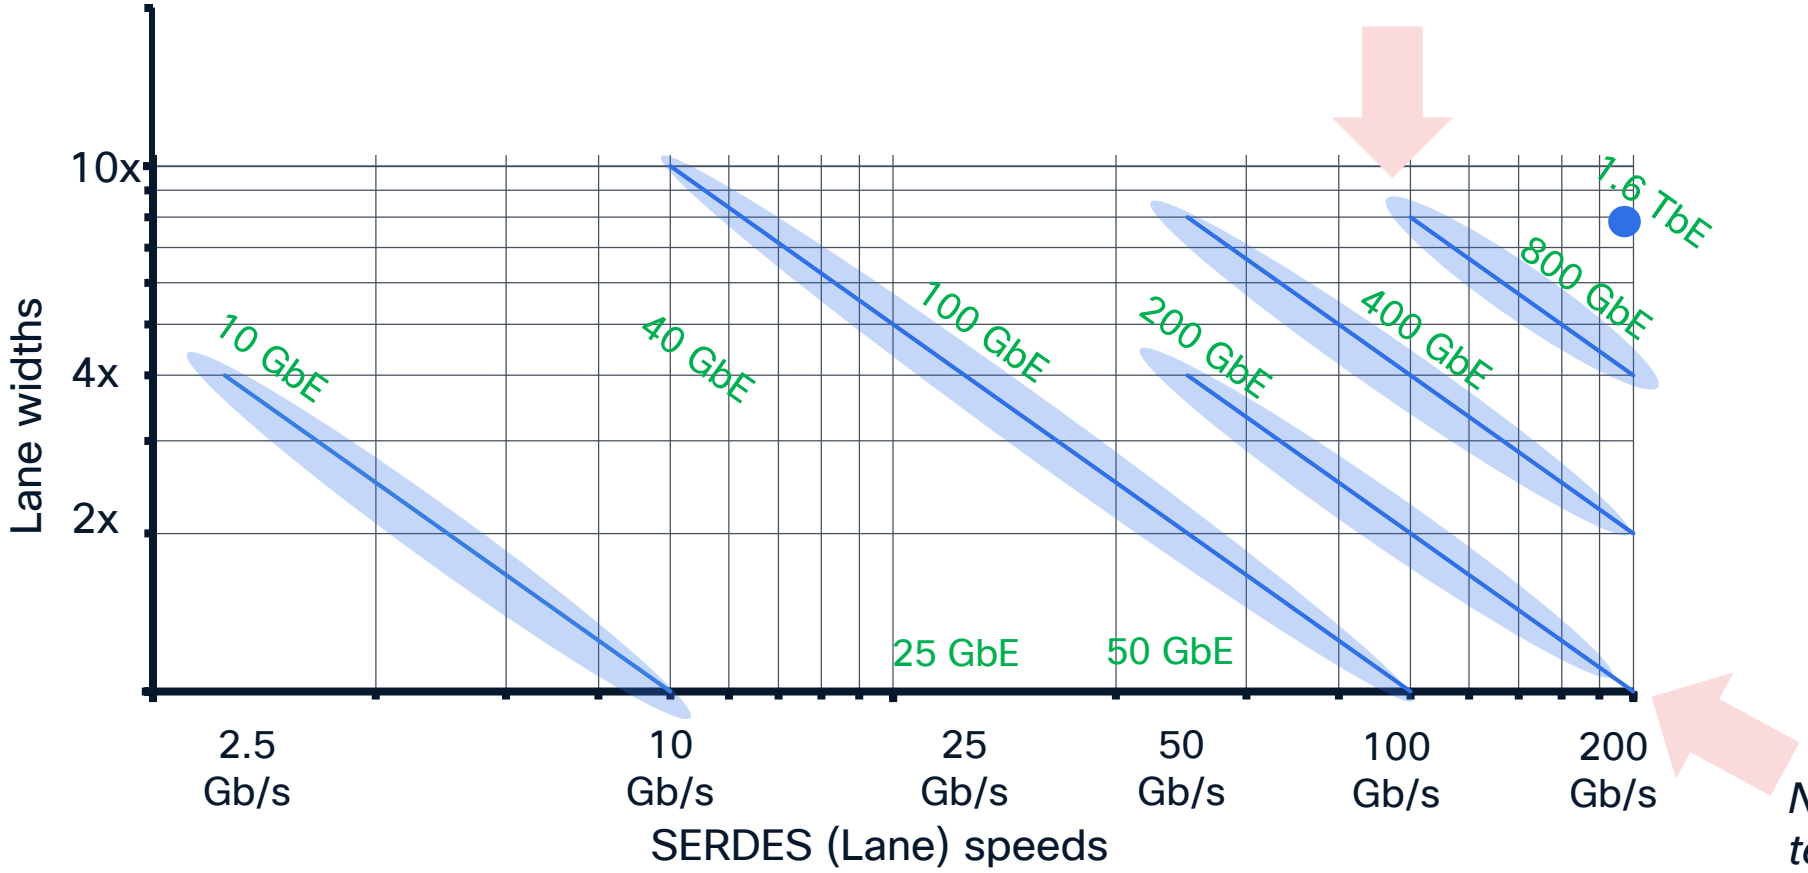

Newer 200Gb/s technology in development now. See later.

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)

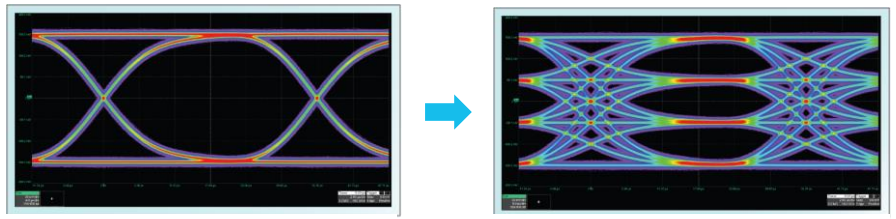

NRZ
1 bit/sec/symbol

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...

Optics firmware updates enabled by CMIS

Firmware updates onsite without unplugging the optics from the Switch:

- Optics complexity…
- Enhance optics performance by adding new firmware features
- Maintain firmware uniformity across install base
- Update optics firmware with minimum disruption with very fast firmware upgrade

Enter the CLI command and load new optics FW onsite

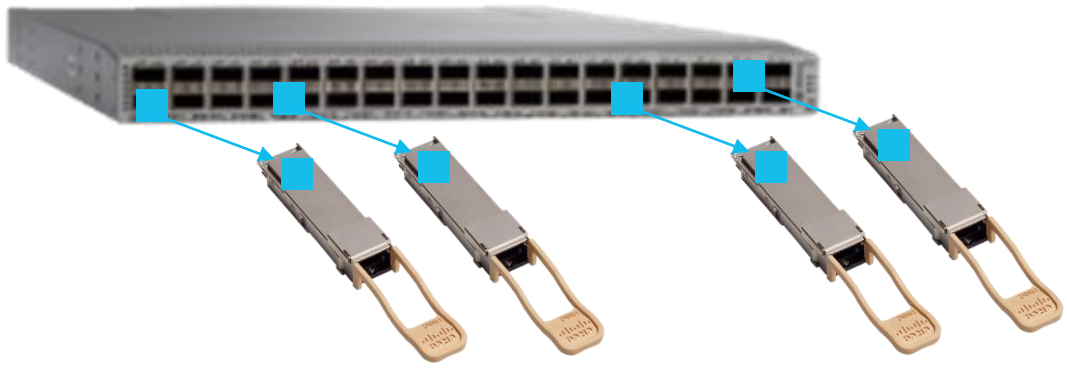

See demo at WoS

100G Single Lambda Optics migration to 400/800G

100G Lambda MULTI-SOURCE AGREEMENT

100G QSFP28

400G or 800G QSFP-DD

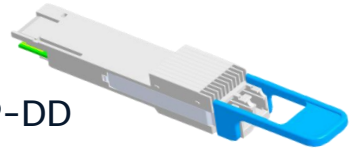

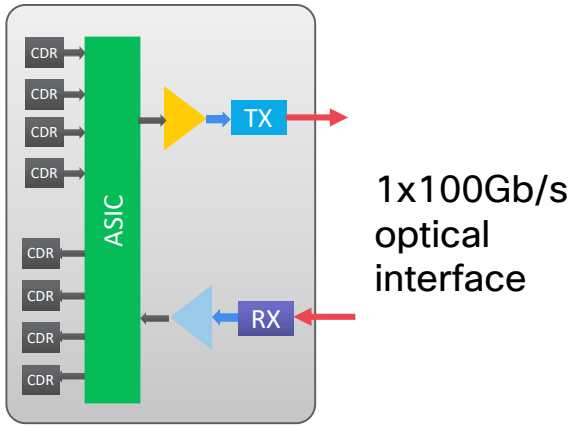

4x25Gb/s electrical interface

1x100Gb/s optical interface

nx100Gb/s optical interface

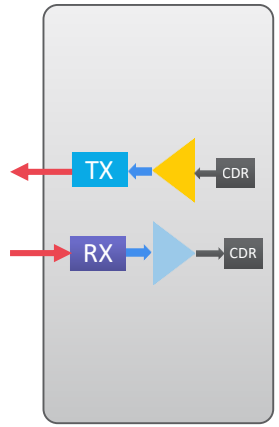

nx100 Gb/s nx50 Gb/s electrical interface

As your network evolves, the optical interface is THE point of interoperation.

Equipment and electrical serdes can evolve through 3 generations (25 Gb/s, 50 Gb/s or 100 Gb/s) without changing the optical interface that interconnects your equipment.

Pluggable Optical Modules Options

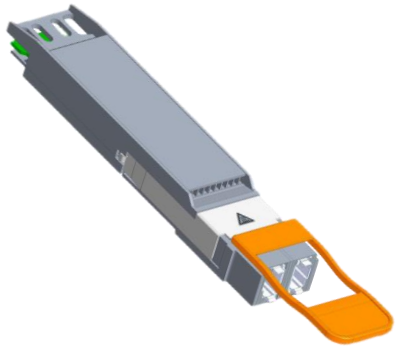
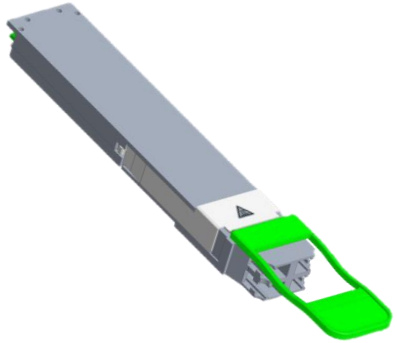
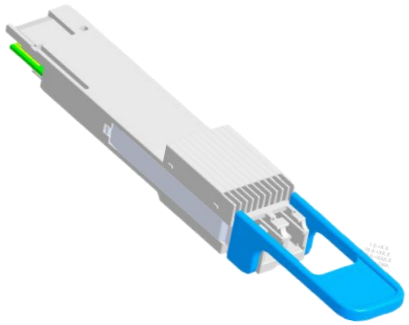
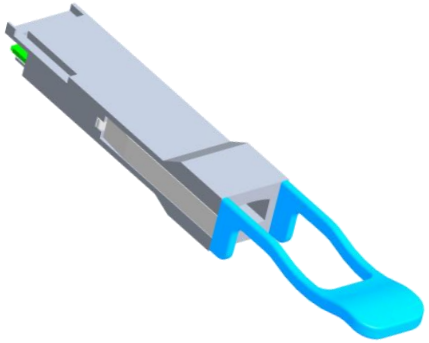

	OSFP	OSFP-RHS	QSFP-DD	QSFP
Host high-speed pins	8	8	8	4
Module capacity (Gbit)	400, 800, 1600	400, 800, 1600	400, 800, 1600	40, 100, 200, 400
Cooling	Integrated heatsink	Riding heatsink	Riding heatsink	Riding heatsink
High-speed IO Compatibility (Gbps)	50-100-200	50-100-200	50-100-200	1—25-50-100
Port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

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

Current deployment trends

400G:

- QSFP-DD saw broad adoption across all markets (Hyperscaler → Enterprise)
- QSFP112 happening now (use primarily in NICs)

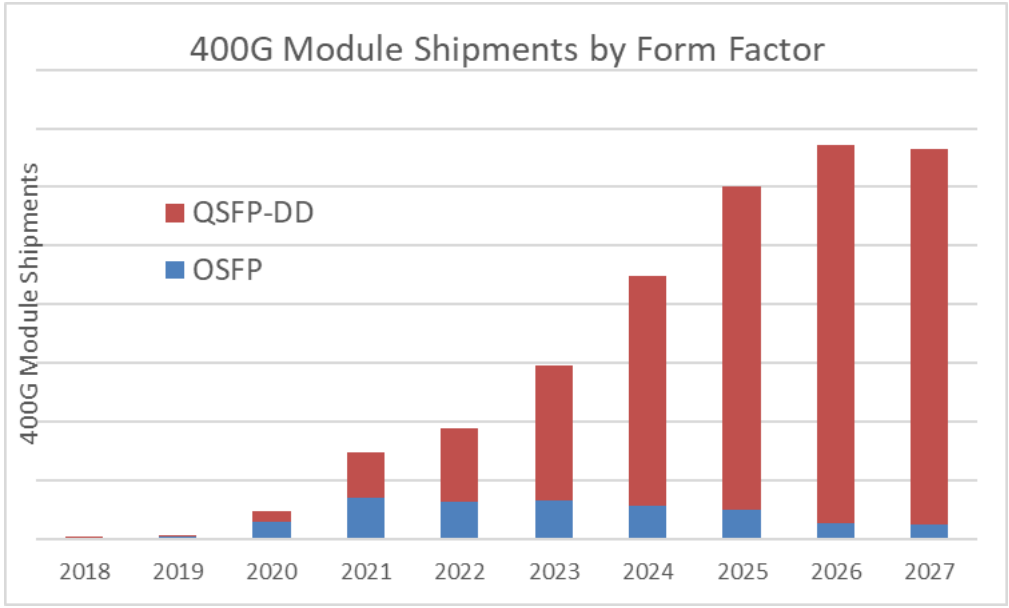

800G:

- OSFP adoption growing. The perceived advantage of integrated heatsink swayed decisions at early adopter hyperscalers
- Nvidia NICs based on OSPF-RHS
- Survey of Cisco's Enterprise DC customers shows strong interest in QSFP-DD due to backwards compatibility

Pluggable form factors – a wealth of options

OSFP

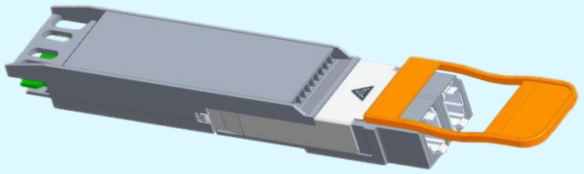

400G/800G/1.6T

OSFP-RHS

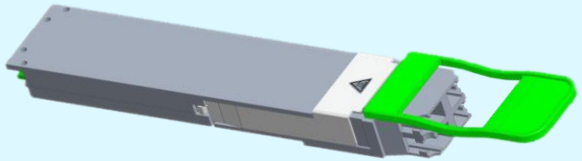

400G/800G/1.6T

OSFP-XD

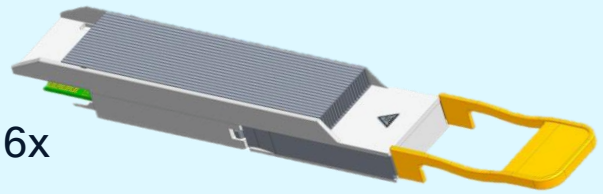

16x

1.6T

Switches

Compute NICs

QSFP-DD

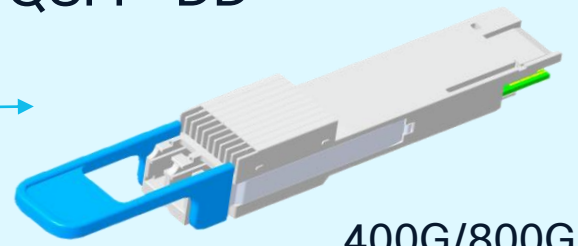

400G/800G/1.6T

QSFP

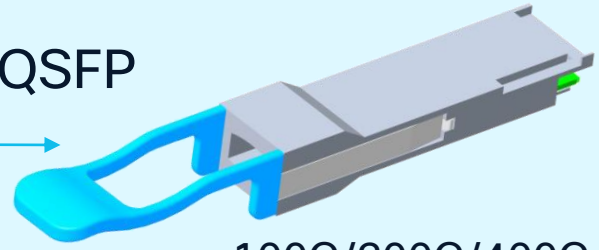

100G/200G/400G
QSFP28, QSFP56, QSFP112

Port compatibility

Everything within the square is compatible with each other

Pluggable form factors – interface interoperability

OSFP

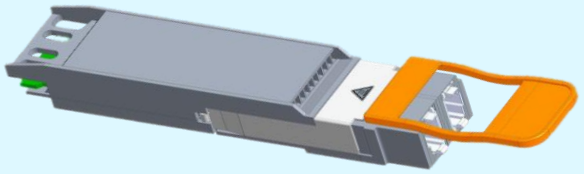

400G/800G/1.6T

OSFP-RHS

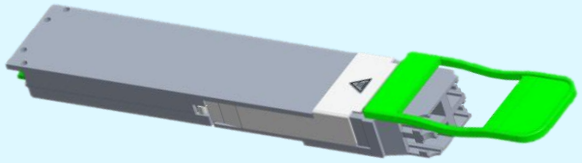

400G/800G/1.6T

OSFP-XD

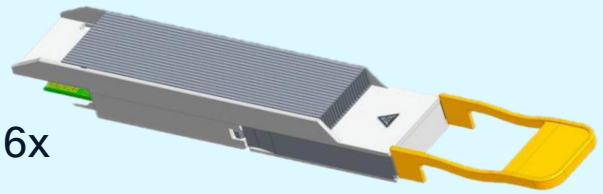

16x

1.6T

Ethernet plug and play as normal.

Breakout as well!

QSFP-DD

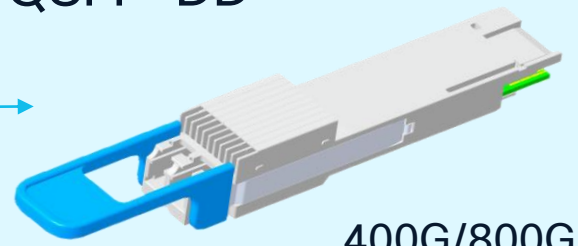

400G/800G/1.6T

QSFP

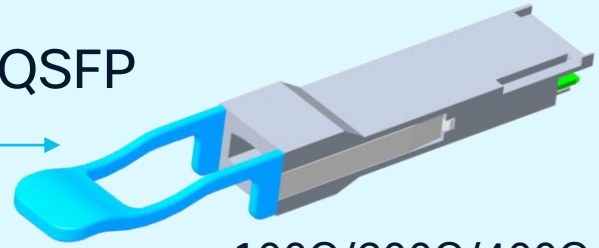

100G/200G/400G
QSFP28, QSFP56, QSFP112

Ethernet Interface compatibility

Port compatibility
Everything within the square is compatible with each other

QSFP-DD Extends to 1.6T

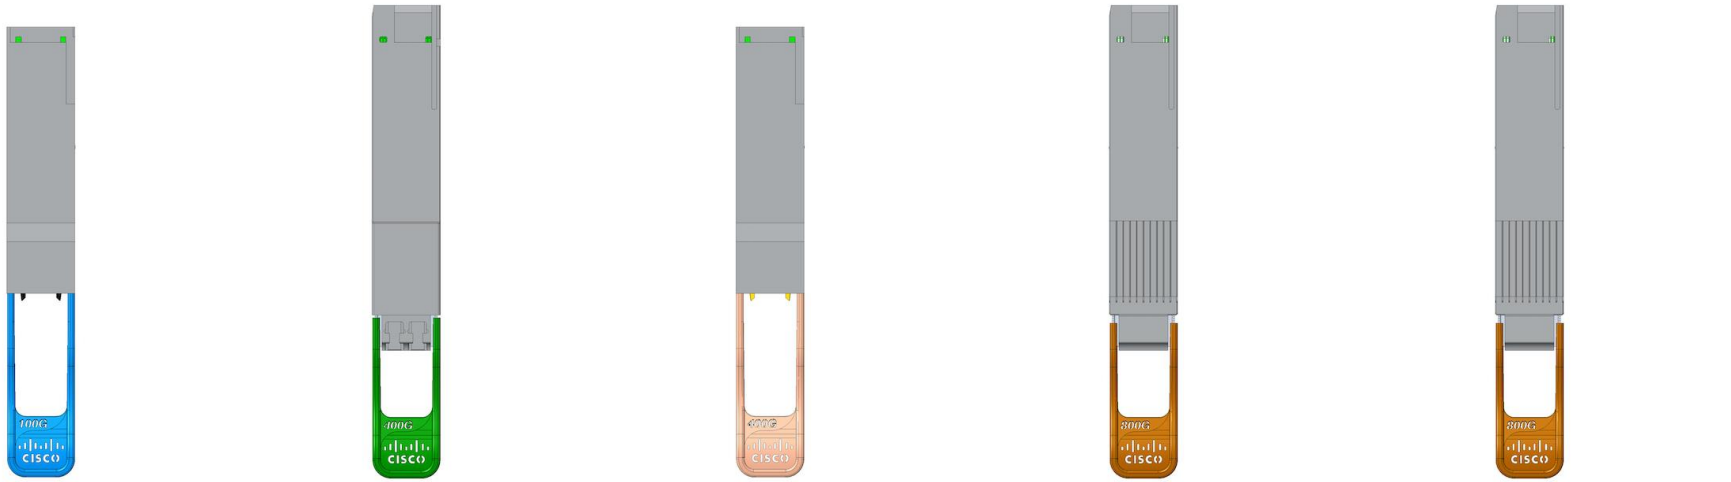

100 Gb/s 25 Gb/s SerDes **QSFP28**	400 Gb/s 50 Gb/s SerDes **QSFP56-DD**	400 Gb/s 100 Gb/s SerDes **QSFP112**	800 Gb/s 100 Gb/s SerDes **QSFP-DD800**	1600 Gb/s 200 Gb/s SerDes **QSFP-DD1600**

Complete backwards compatibility

Powerful value proposing to enable seamless network growth and investment protection. Technically superior solution.

Module Roadmap for 200/400/800G

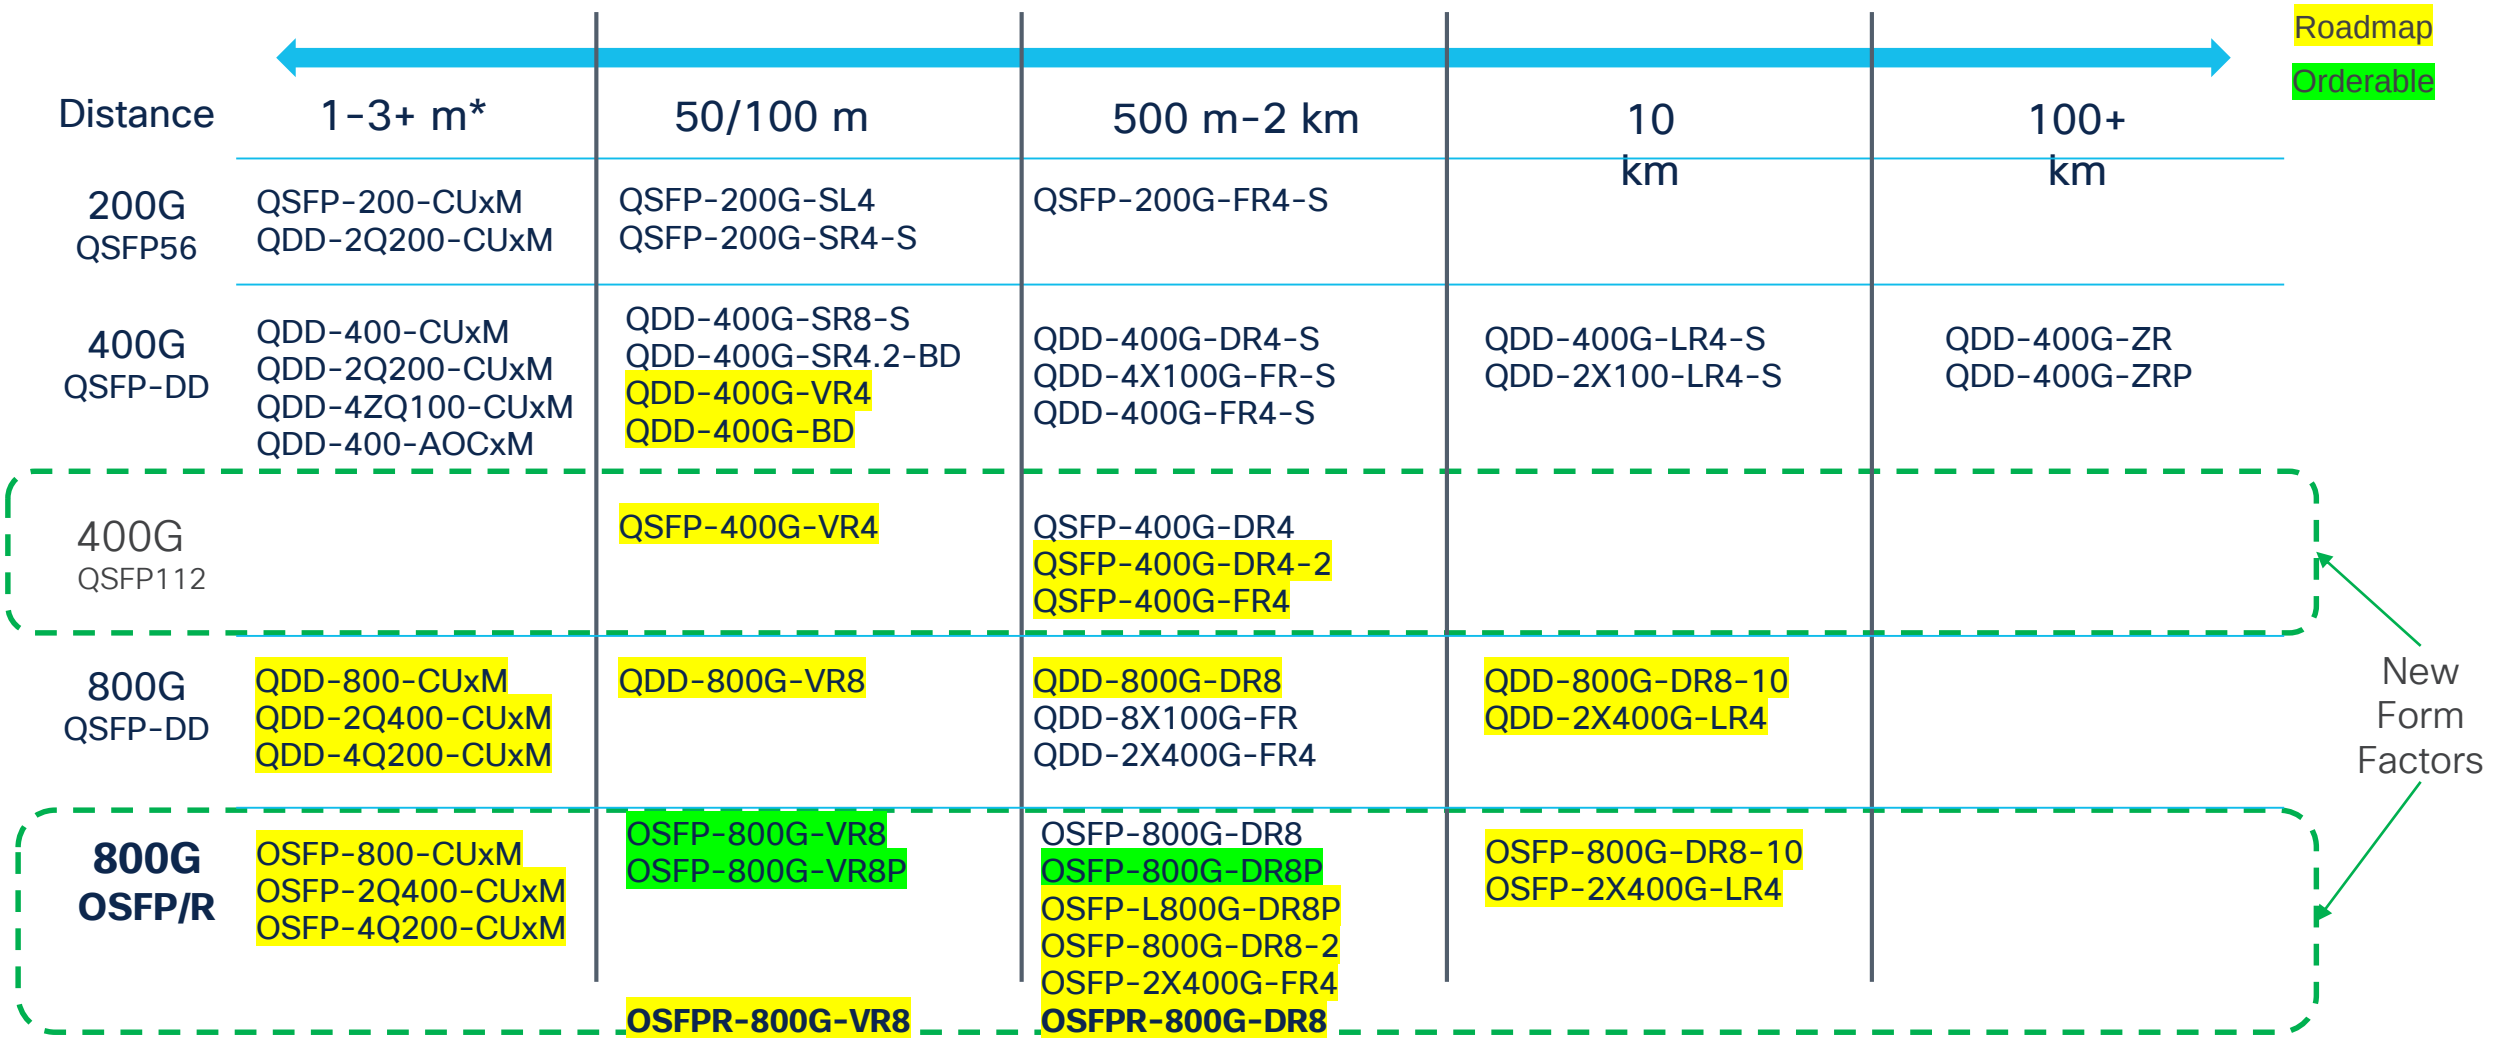

Current trend: 800G Pluggables supporting dense 400 GbE

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

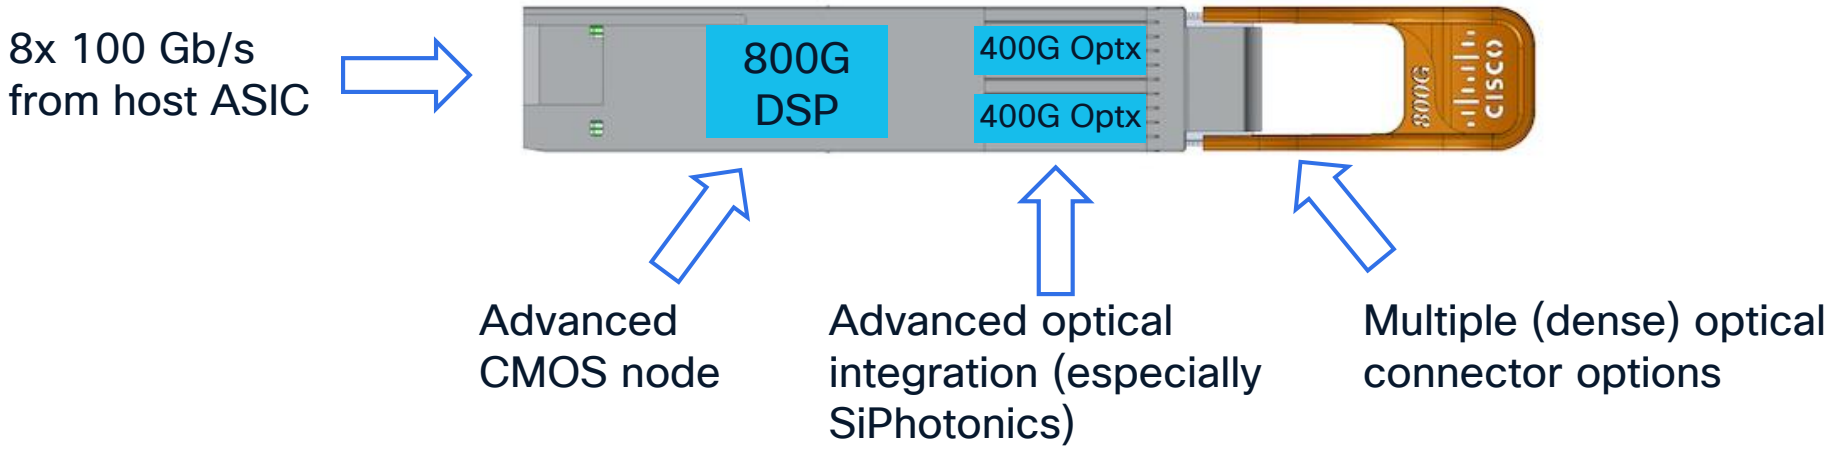

Implementing Dense 400 / 800 GbE

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

Examples using QSFP/QSFP-DD

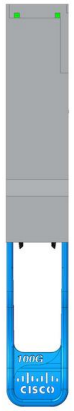

100 Gb/s
25 Gb/s SerDes
QSFP28

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

400 Gb/s
100 Gb/s SerDes
QSFP112

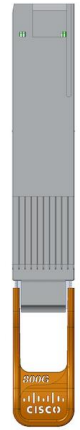

800 Gb/s
100 Gb/s SerDes
QSFP-DD800

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

What is a Breakout?

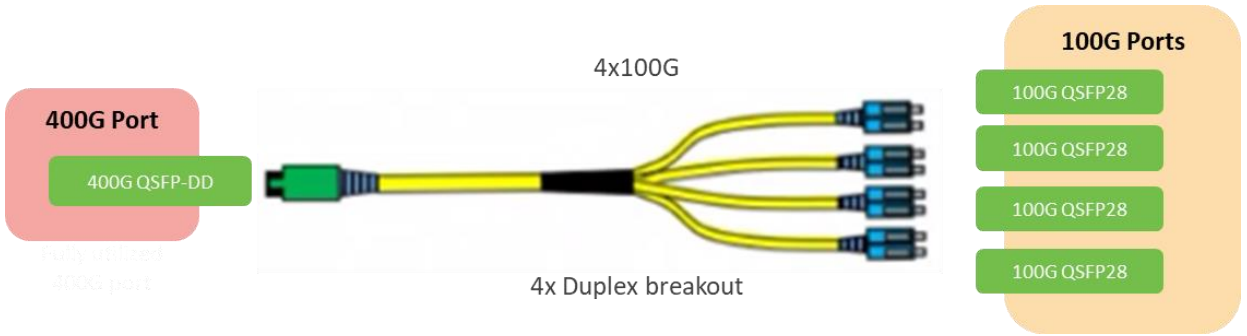

- Breakouts take advantage of ports with multiple optical lanes for both the Tx and Rx
  - Optical lanes in this context means pairs of fibers
  - e.g. 400G DR4 optical connector has 4 pairs of fibers, each pair can be configured as a 100G-DR
- A breakout is when a port (module) is configured as multiple lower speed interfaces
- Power efficient (12W @ single 400G module vs 17W @ 4x 100G modules)
- Breakout transceivers generally use MPO connectors which have multiple fibers for both the Tx and Rx
  - The port controls how the module will be configured either for breakout or non breakout operation
  - Structured cabling solutions preferred for deployment
  - For 2x breakout, module's support regular LC connectors too
- Breakouts can also be done for copper cables and AOCs
  - Cables and AOCs are fixed for either breakout or non-breakout applications

Breakouts Promote Sustainability

32 1RU servers per cabinet w/ 16 cabinet row

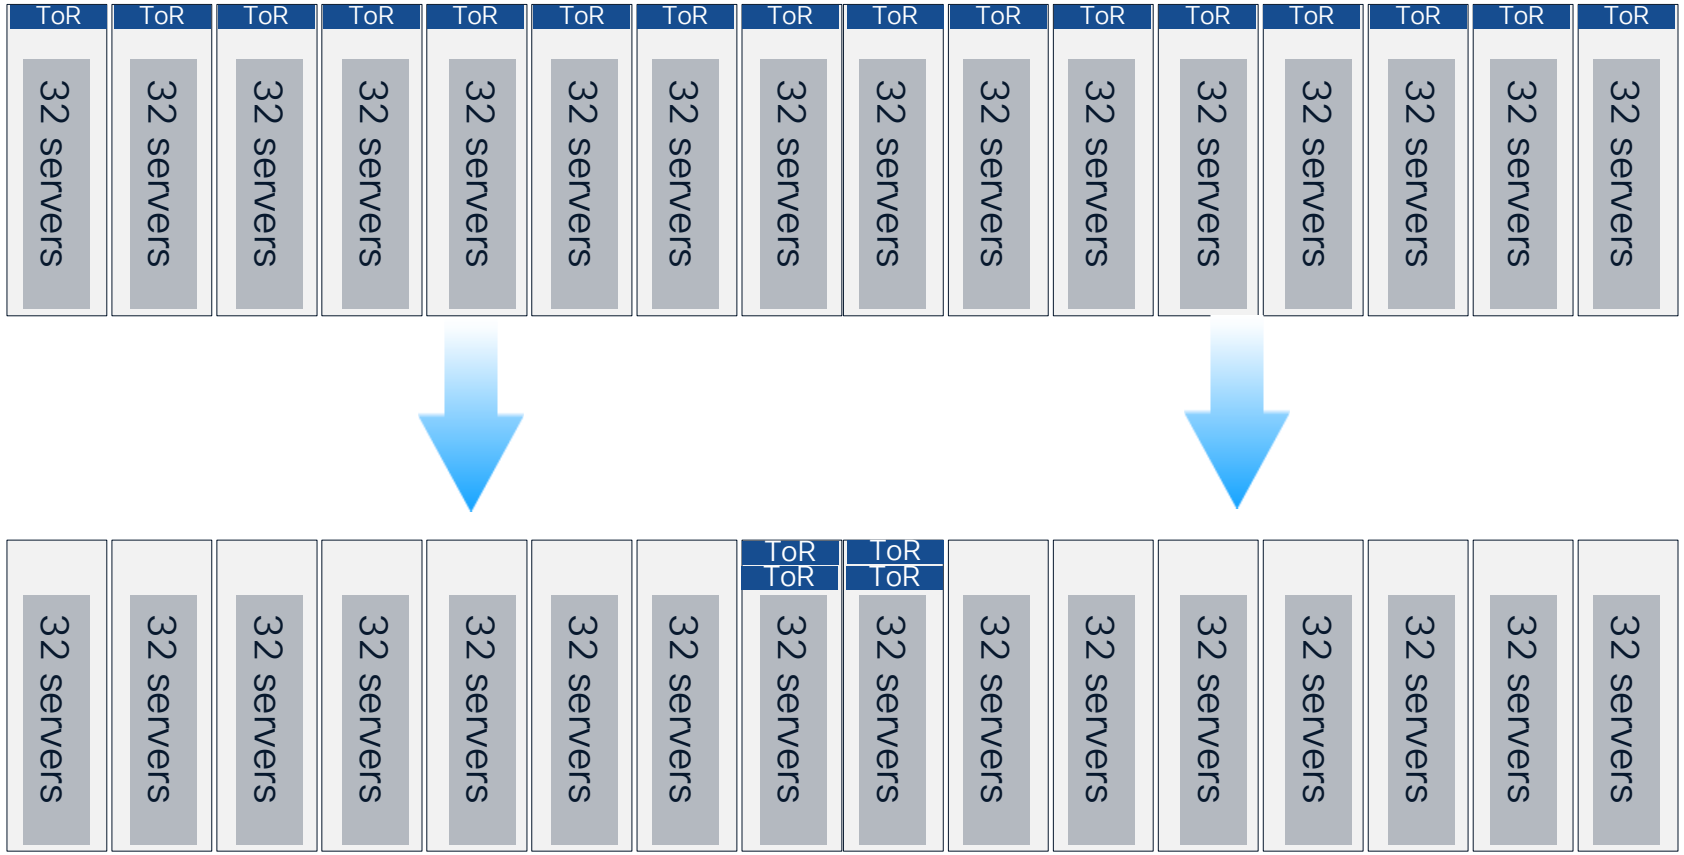

Non Breakout Architecture
- 16 TOR = 811W x 16 = 12976W
- 512 DAC 0.5W per end = 512W
- Total power 13,744W

Breakout Architecture using 64 port switches
- 4x switch: 1324W x 4 = 5296W
- 128 DR4 modules 9W x 128 = 1152W
- 100G-DR modules 3W x 512=1536W
- Total Power 7984W

40% savings in power

Energy optimized architectures using breakout are explored in this whitepaper.

Introducing Cisco 400G BiDi for duplex fiber

Increase the capacity of installed multimode fiber

Cisco first to market with BiDi technology

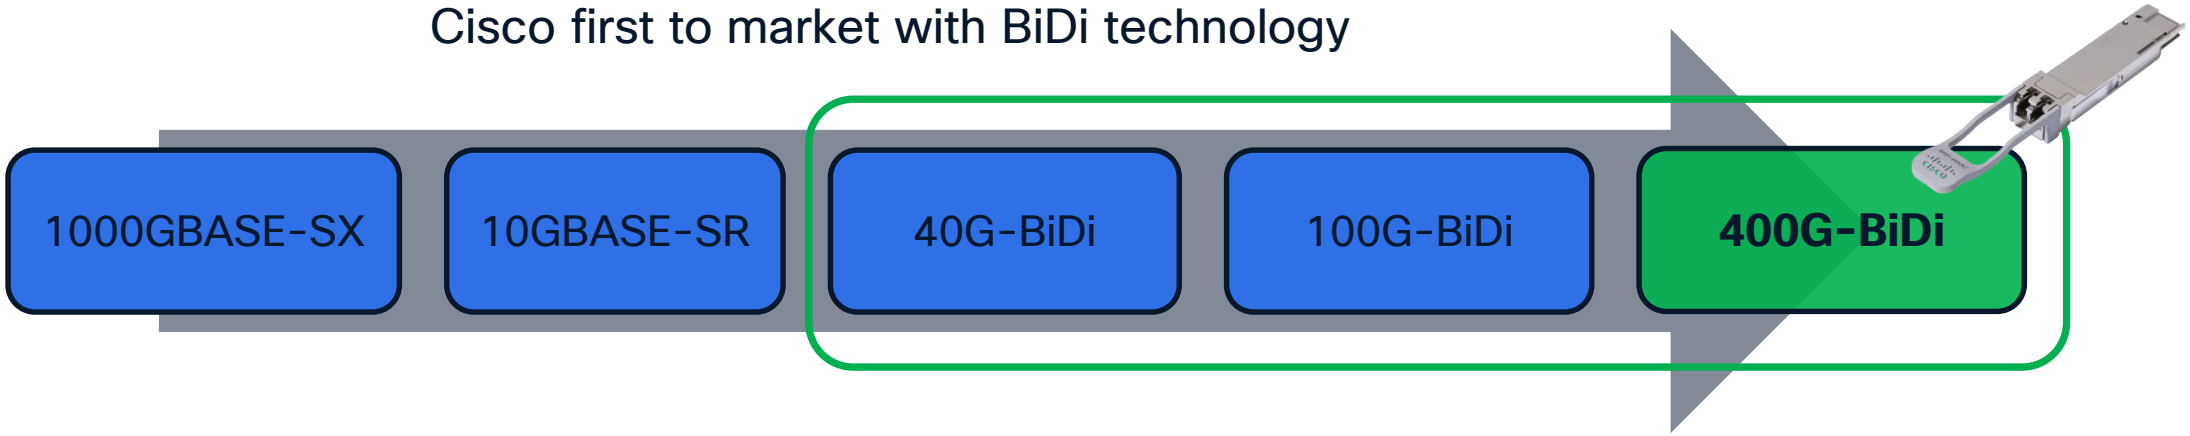

Cisco BiDi innovations extend the life of duplex multimode fiber infrastructure

GLC-SX-MMD	SFP-10G-SR	QSFP-40G-SR-BD	QSFP-40/100-SRBD	QDD-400G-BD
SFP	SFP	QSFP+	QSFP28	QSFP56-DD
1GE	10GE	40GE	100GE	400GE
Duplex LC	Duplex LC	Duplex LC	Duplex LC	Duplex LC
1000m	400m	150m	100m	100m

400G migration options for duplex multimode fiber

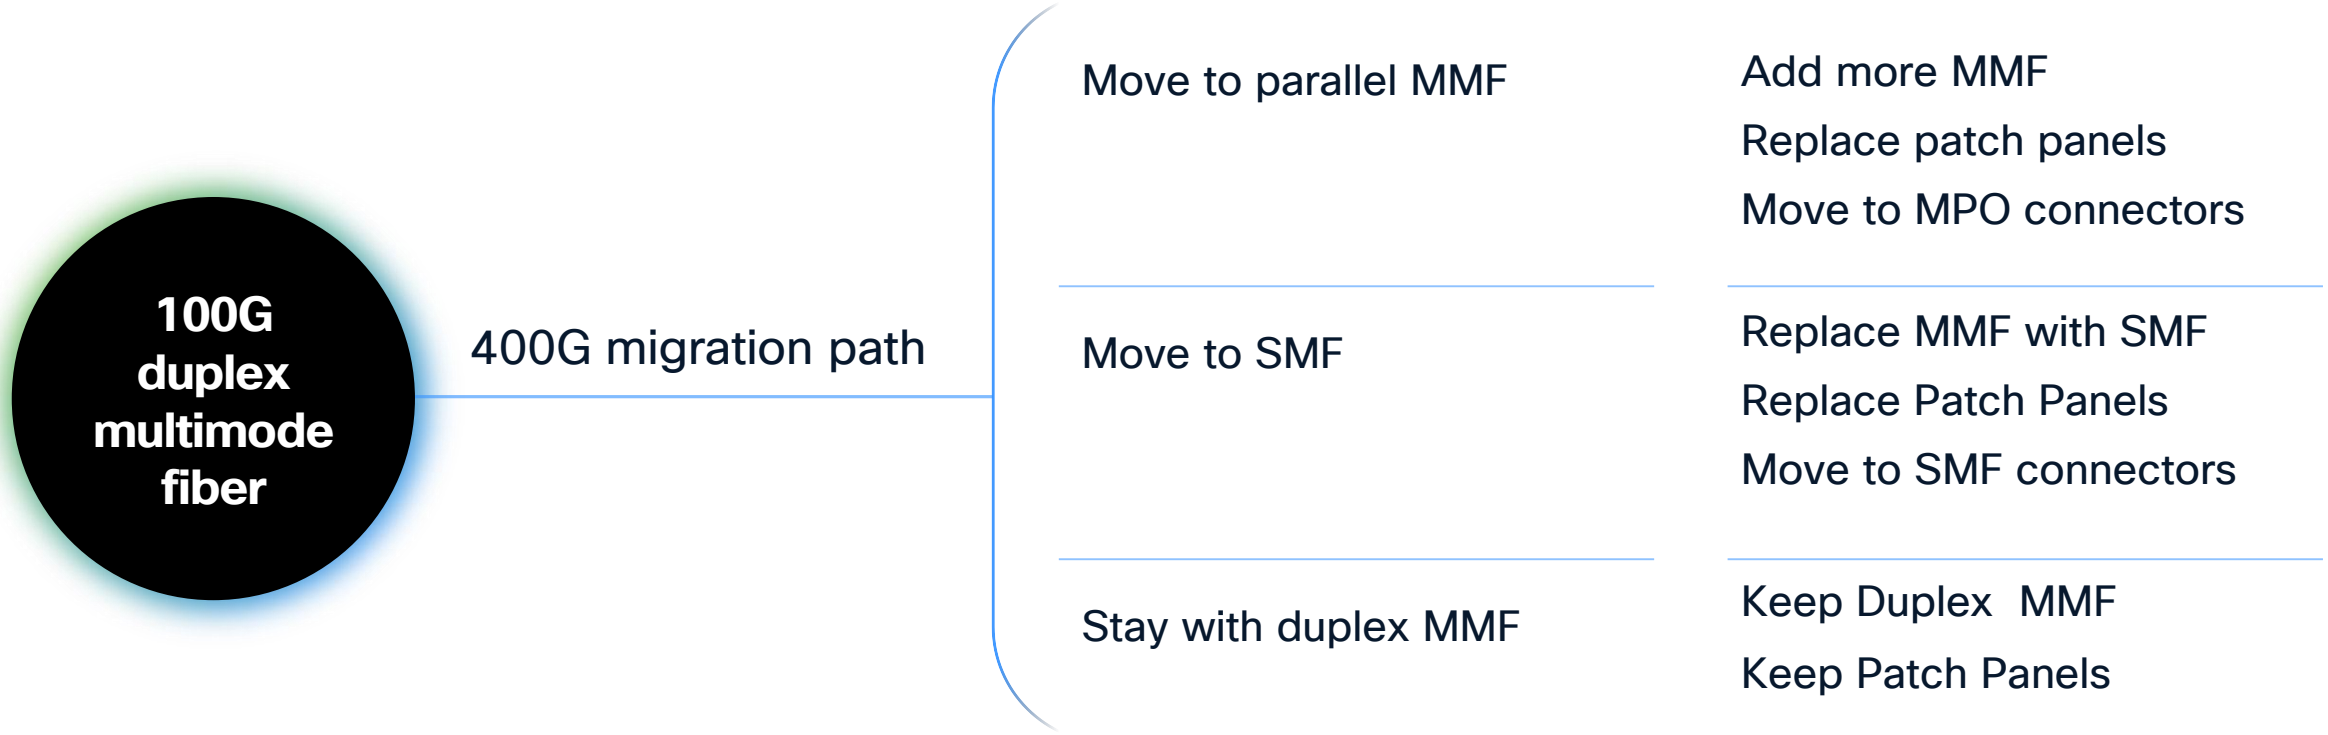

Investment protection with 400G BiDi optics

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 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

Cisco has the migration path for your data center network

Cisco optics portfolio for data center speed migration

		1G	10G	40G	100G	400G
MMF	Duplex	GLC-SX	SFP-10G-SR	QSFP-40G-SR-BD	QSFP-40/100-SRBD QSFP-100G-SR1.2 QSFP-100G-SR	QDD-400G-BD
	Parallel			QSFP-40G-SR4	QSFP-100G-SR4-S	QDD-400G-SR8-S QDD-400G-SR4.2-BD QDD-400G-SR4
SMF	Duplex	GLC-GE-DR-LX	SFP-10G-LR-S	QSFP-40G-LR4	QSFP-100G-DR-S QSFP-100G-FR-S QSFP-100G-LR4-S	QDD-400G-FR4
	Parallel			QSFP-4X10G-LR	QSFP-100G-PSM4-S	QDD-400G-DR4-S QDD-4X100G-FR-S

400G BiDi vs 400G Parallel Optics

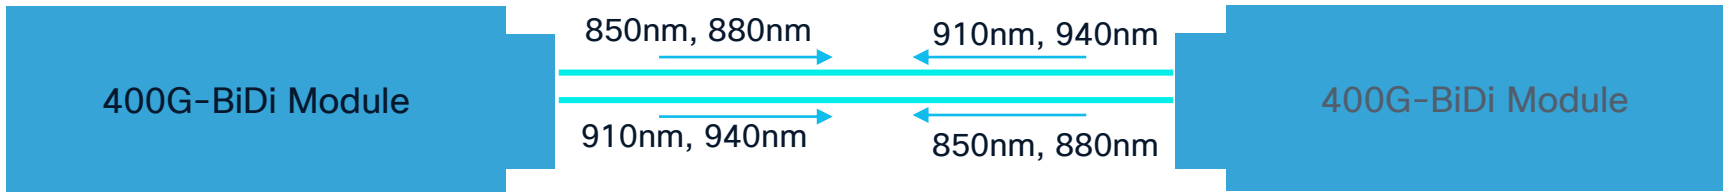

- Only 1 fiber pair
- 4 wavelengths
- Duplex LC connectors
- Compatible with 100G BiDi

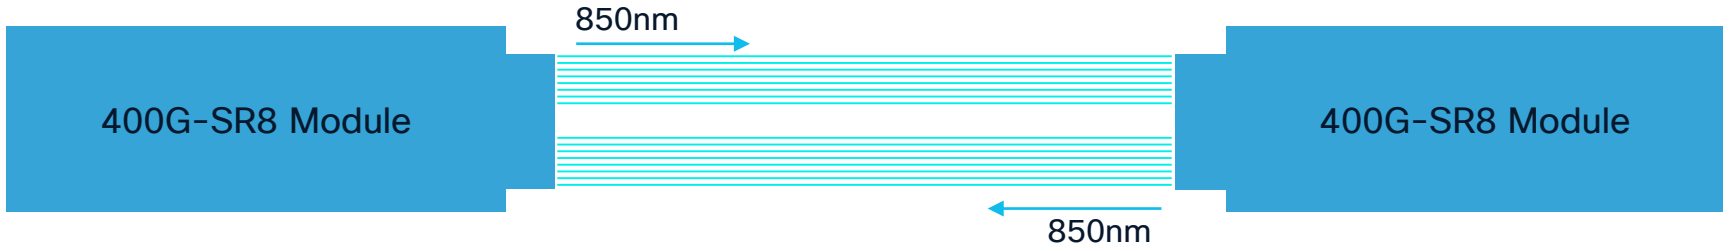

- 8 fiber pairs
- 1 wavelength
- MPO16 connectors
- Angled-polished connectors

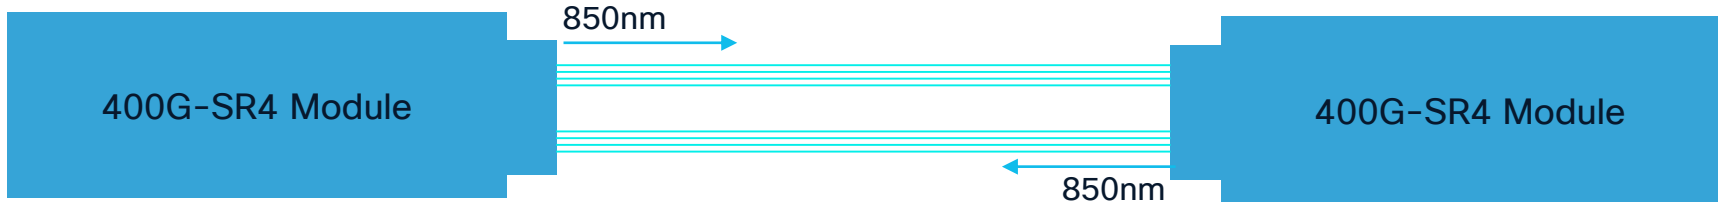

- 4 fiber pairs
- 1 wavelength
- MPO12 connectors
- 4 lanes x 100G

Key takeaways – client (IMDD) reaches

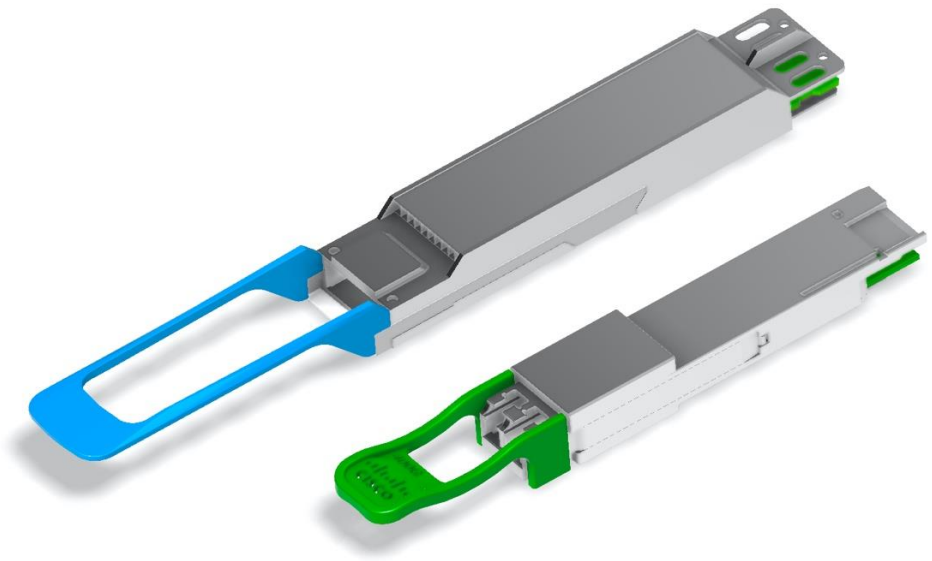

400G & 800G Standards are mature. Ethernet interop is guaranteed regardless of form factor

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

Sustainable and 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
  - 100G, 400ZR/ZR+, 800ZR/ZR+

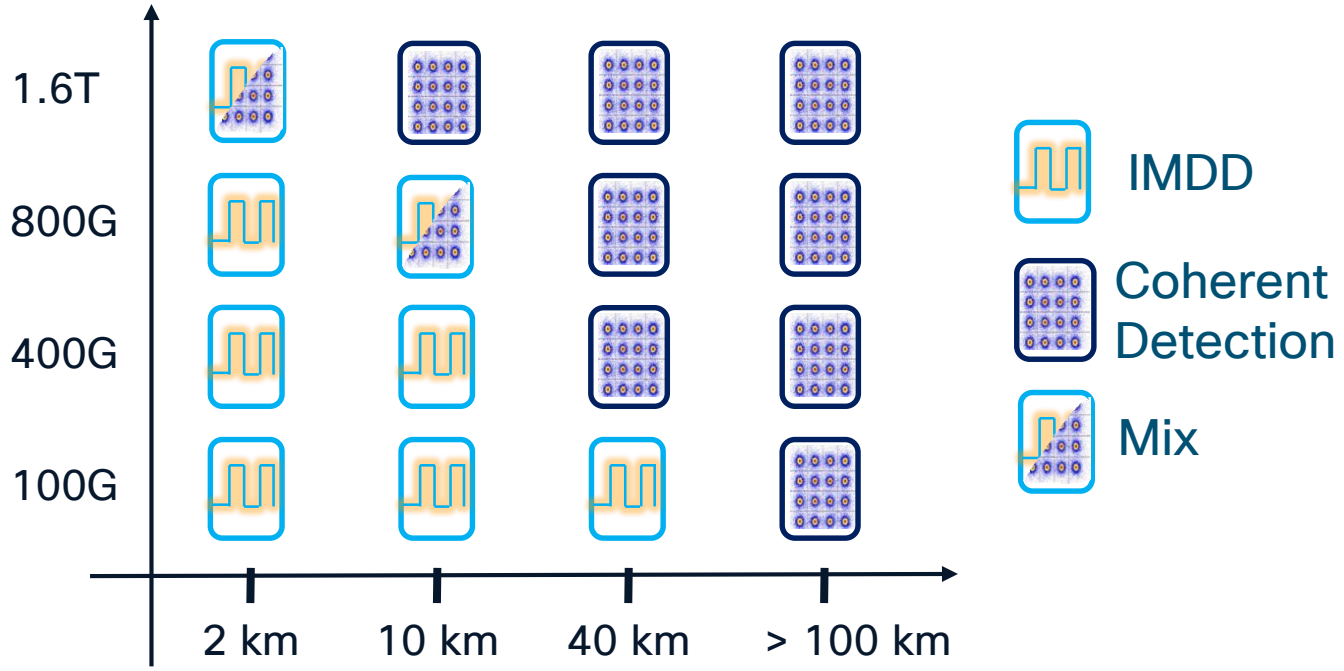

As speed increases, the crossover distance between IMDD & coherent decreases

The rise of “Coherent Pluggables”

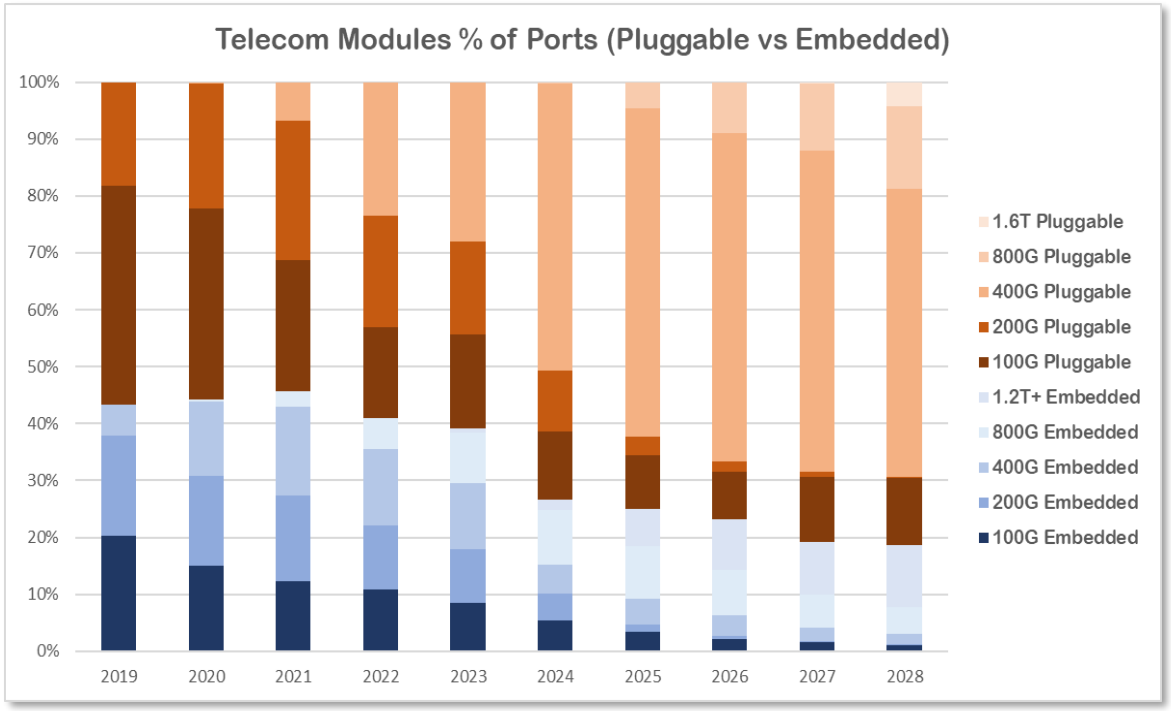

Today >70% of the coherent ports are pluggable coherent optics

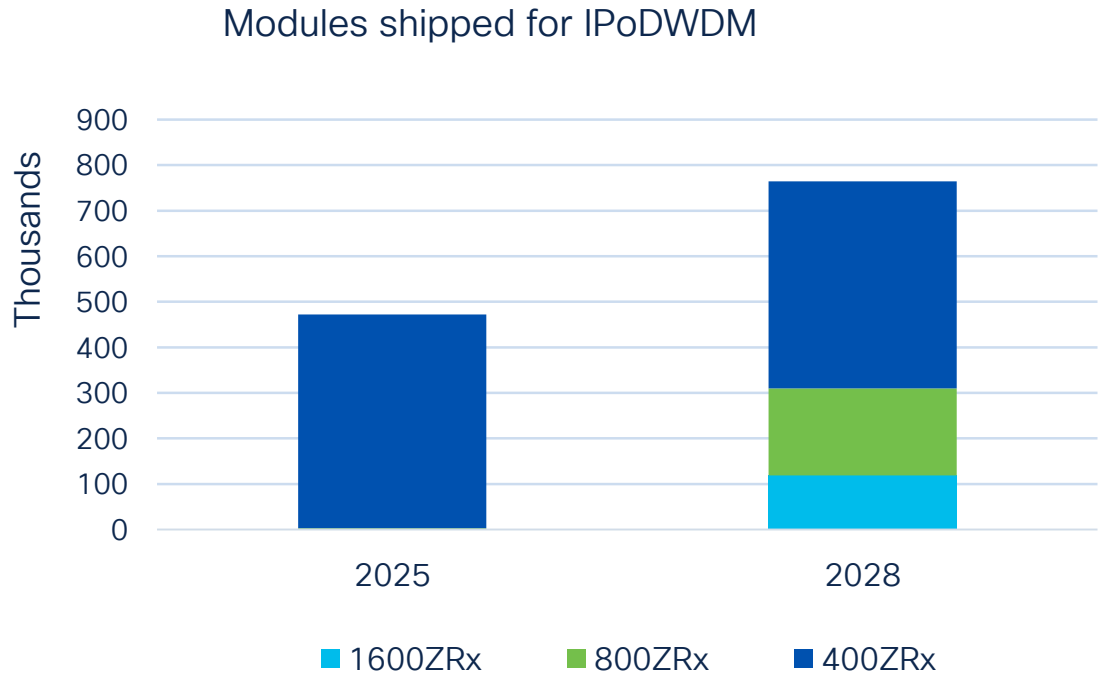

Source: Cignal AI, IP-over-DWDM Pluggables Forecast, Aug. 2024

Cisco/Acacia is pioneer and market leader for coherent pluggables

Vertical Integration is Key to Lowest Cost and Best Performance

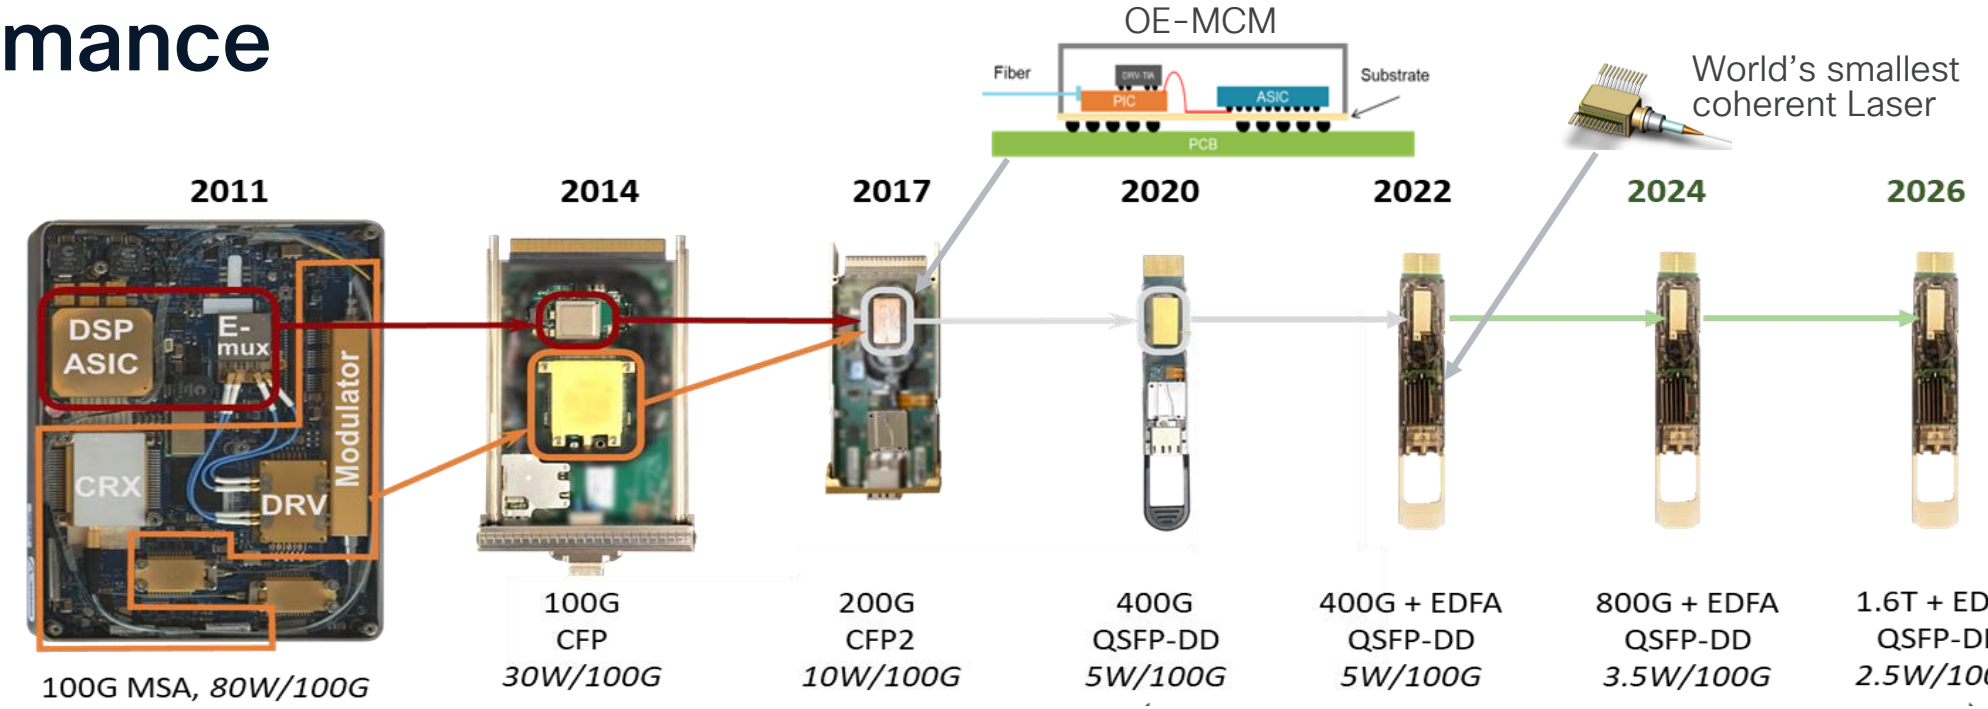

In-House:	DSP	+ SiPh Optics	+ OE-MCM	+ ASIC	+ Laser	100% Vertical
QAM:	QPSK	QPSK	16QAM	16QAM	PCS	PCS
CMOS:	40nm	28nm	16nm	7nm	4/3nm	2nm
Optics:	Discrete, InP, LiNbO3	SiPh	SiPh	SiPh	SiPh	SiPh/TFLN

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

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

400G DCO's: Native DWDM Interfaces in Routers

400G-ZR

- Access to Metro Networks
- 400GE Client | 400G Trunk
- Up to 120km
- TX Power -10dBm | C-FEC

400G-ZR+

- Access to Long Haul Networks
- 100GE, 200GE, 400GE, n x100GE MXP Mode | 100G, 200G, 300G, 400G Trunk Rates
- Up to 1,400km at 400G
  Longer distances at lower trunk rates
- TX Power -10dBm | C-FEC and O-FEC | TX Power +0dBm

400G-ZR+ Bright

Cisco 400G-ULH QSFP-DD ZR+

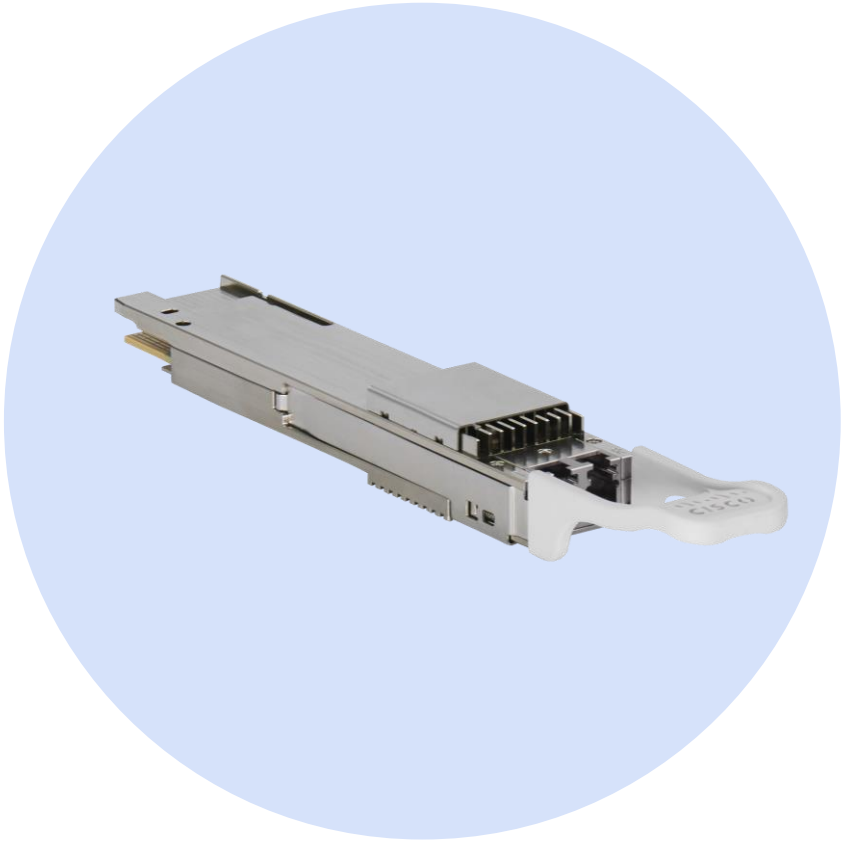

Product ID: DP04QSDD-ULH-A1= (C-Band)
DP04QSDD-LLH-A1= (L-Band)*

- 400G ZR+ QDD pluggable supporting Ultra Long-Haul Applications
  - Designed for existing QSFP-DD Ports and Platforms
  - Similar Power Consumption (<24W)
- Supports Multiple Baud Rates
  - 118GBaud QPSK – 3,000km
  - 98GBaud PCS – 2,500km
  - 87GBaud PCS - 2,000km
  - 75GBaud PCS – 1,600km
  - 66GBaud PCS – 1,200km
  - Existing 400G-ZR+ Modes for interoperability
- Easier provisioning with Appsel Codes
- Two different versions – C-Band and L-Band*

*Disclaimer: L-Band Version is pending Release.
Please confirm availability with Cisco.

400G ULH QSFP-DD Field Trials

Arelion: 2,253km Link

- Denver-Chicago, w/ margin
- 112.5 gigahertz spectrum
- 400G ULH pluggables
- Cisco 8000 series routers

Arelion

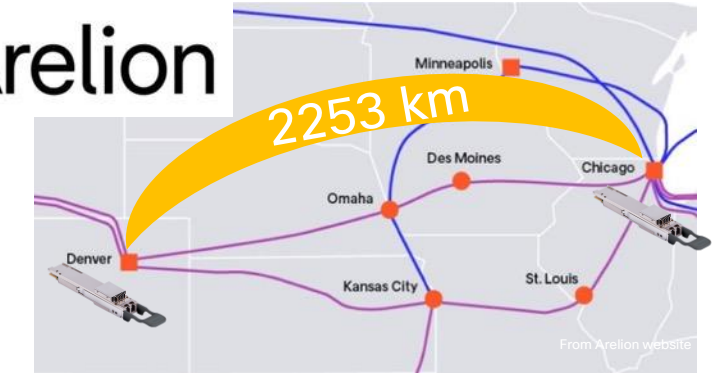

Reduces Costs	35% CAPEX	84% OPEX

Internet2: 3,000km Link

- Albany-Indianapolis
- 400G ULH pluggables
- Cisco 8000 series routers
- Third-party OLS

INTERNET2

Reduces Power	68%*

* Enabled by optical and routing advancements

400G QSFP-DD Digital Coherent Optics Portfolio & Use Cases

ER1

Point to Point
Intra-data center, campus interconnect, core-to-edge router

Lowest Cost
Based on fixed laser with simple point-to-point connectivity

Short Reach
Up to 45KM for unamplified at 13dB

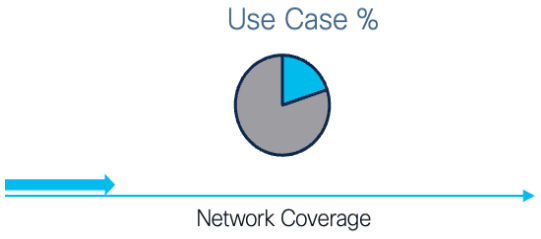

ZR

Point to Point
Web, Data Center Interconnect, Non-SP/SP router interconnect

Low Cost
Lowest cost 400G DCO option for very simple designs

Short Reach
Up to 120KM for P2P amplified links

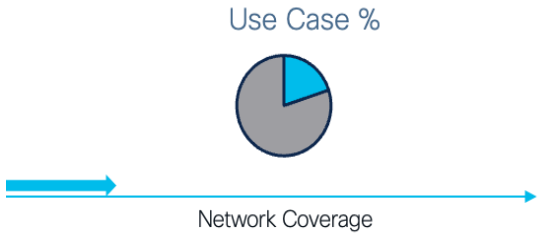

Bright ZR+

Transport Centric
Service Providers, Routed Optical Networking

High TX Power
+1dbm for difficult spans; interop with brownfield transponder & legacy line systems, ~1400km

Advanced Features
TOF, OTN and L1 Encryption features; integrated optical amplifier

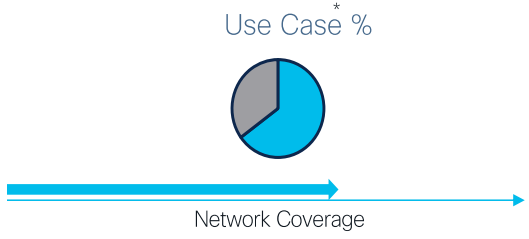

ULH

Transport Centric
Service Providers, Routed Optical Networking

Long haul optimized
400G with **RON Anywhere** up to 3000km, high Tx power

Power optimized
Designed to fit in all routers support 400G Bright ZR+ today

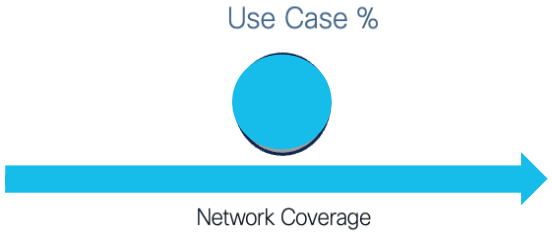

Welcome to the Next-Gen Coherent Pluggables !

- 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

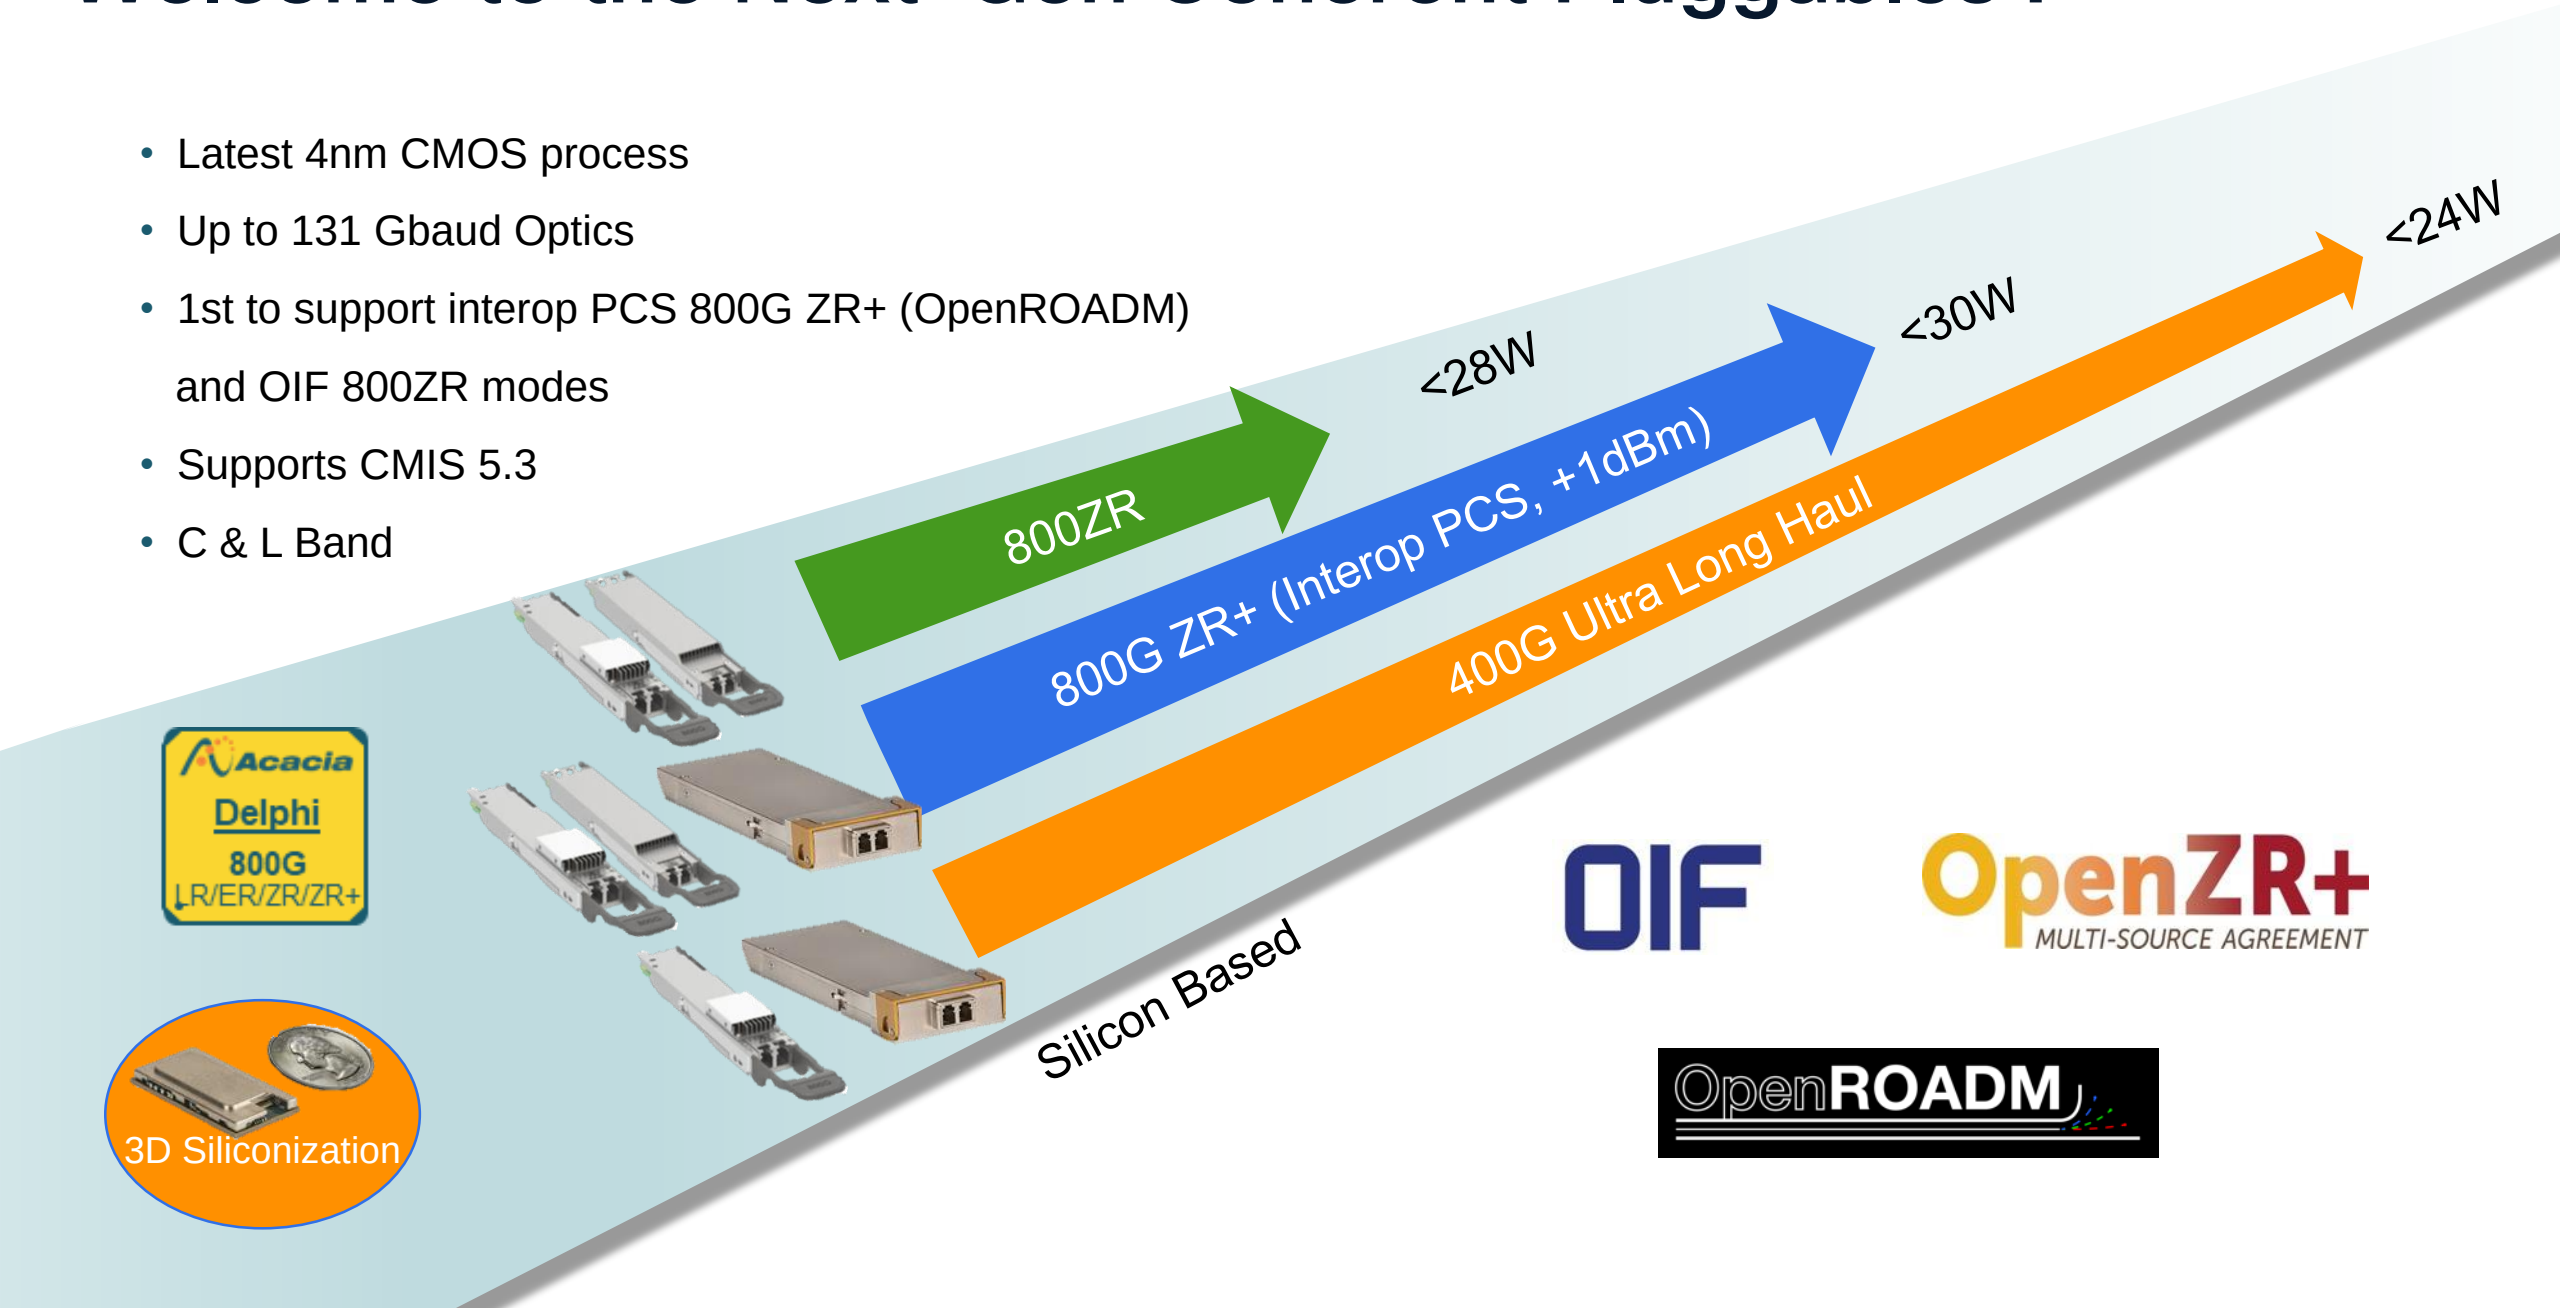

800G ZR+ Successful Service Provider Field Trials

Colt: First Ever 800G ZR+ Field Trial

- Frankfurt–Munich, 634km w/ margin
- Mix of G.652 and G.655 w/ multiple >100km spans, multiple ROADMs

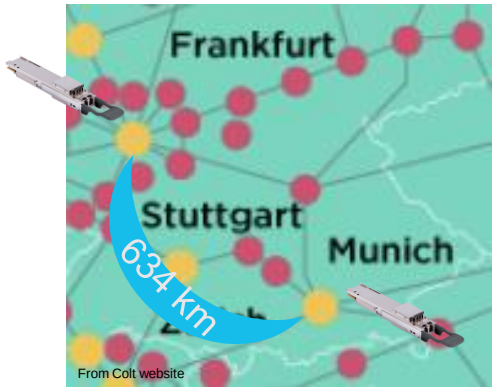

Reduces Power	33% compared to 400ZR+

Arelion: First >1000km 800G ZR+ Field Trial

- Stockholm–Hamburg, 1069km w/ margin
- Multiple >100km spans, multiple ROADMs

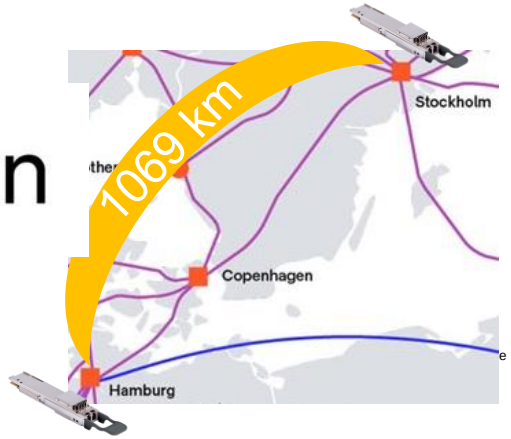

Reduces Costs	66% CAPEX	95% OPEX

Cisco 100G QSFP28 ZR DCO

Orderable!

Line Protocol	OTU4-SC, ITU-T G.709.2
Client interface	100GBASE-R CAUI-4
Case temp	0C to 70C (DP01QS28-E20) -20C to 85C (DP01QS28-E25)
Wavelength range	191.275THz to 196.125THz
Tx output power	-8dBm
Tx In-band OSNR	35dB
Tx Out-of-Band OSNR	23dB
Rx sensitivity	-30dBm
OSNR sensitivity	15.5dB/0.1nm
Minimum Rx power (OSNR limited)	-22dBm
PMD tolerance	>10ps
Dispersion tolerance	-1000ps/nm to 6000ps/nm
SOP change tolerance	50krad/s

Compatible with QSFP28 Router ports

- 100G Staircase FEC
- 80km dark fiber, 300km DWDM amplified
- C-band tuneable
- Commercial and Industrial Temp versions
- CMIS 5.2

Coherent Optics enables IP-Optical Convergence

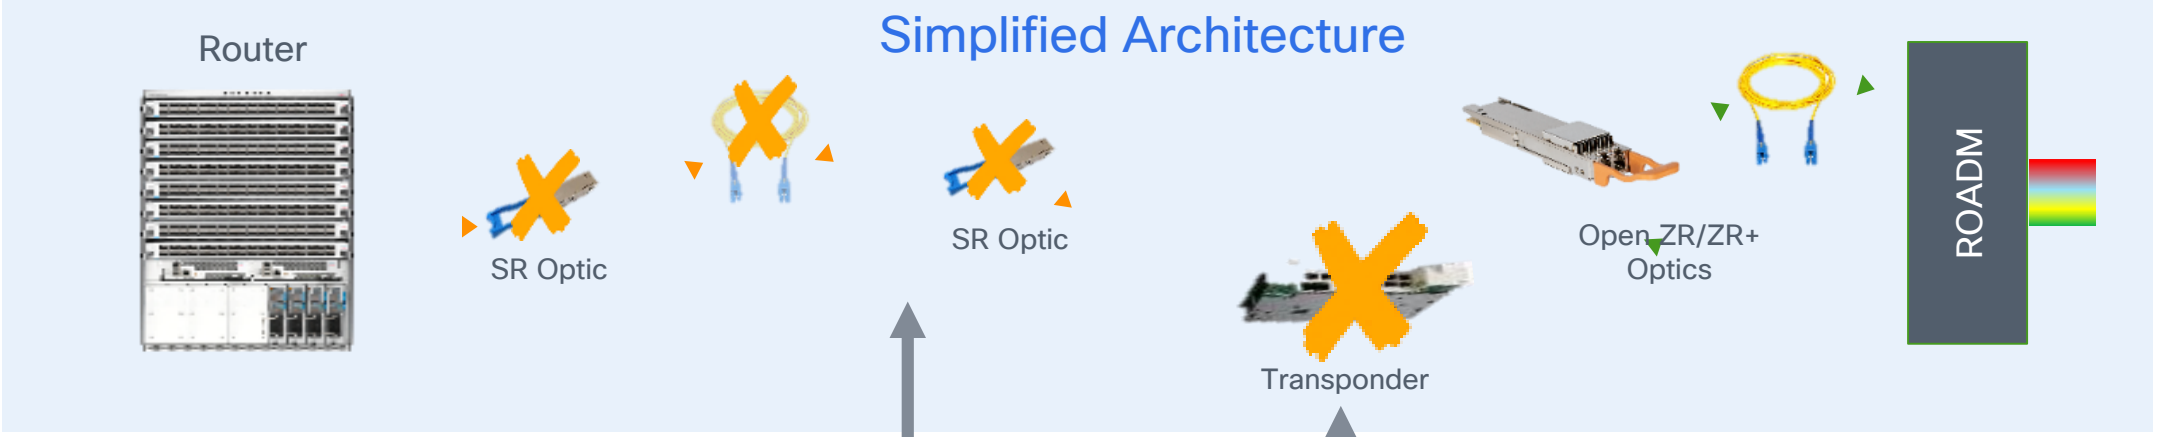

- Eliminate transponder and associated shelf/chassis
- Eliminate client optics

"Simplicity is the ultimate sophistication"
– Leonardo da Vinci

Routed Optical Network (RON)

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

Different use cases and 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

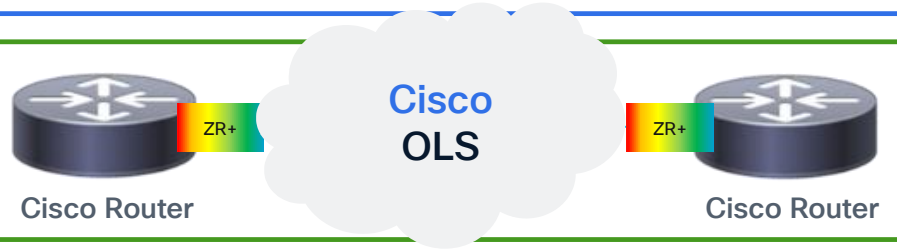

>75% of the RON deployments are over third-party line system! OLS: Open Line System

* ACG Research 2023

Coherent pluggable brings some important levers

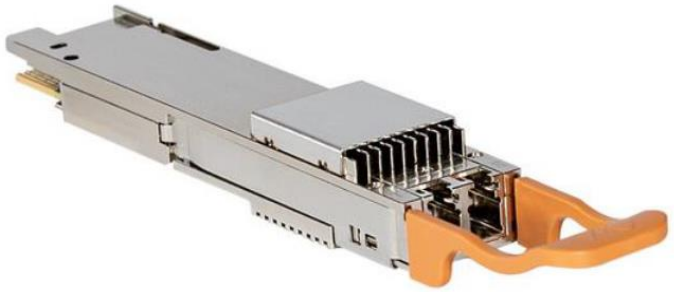

01

Network Costs & Efficiencies
Capex costs breakdown, $/100G
PAYG models etc.

02

Power, Thermal & 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

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 sustainability, operational efficiency & cost efficiency

Beyond 800G – how is AI impacting interconnect?

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

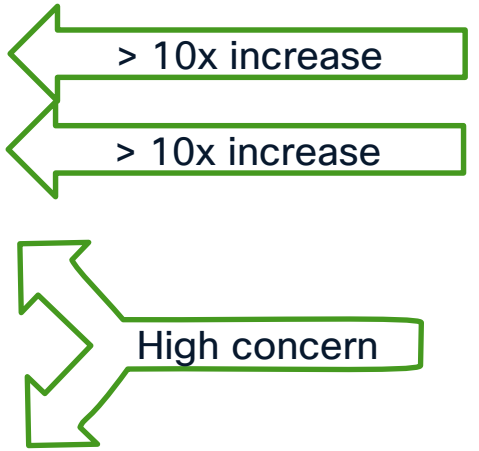

Priorities for Interconnect in an AI world

Minimize workload impacts
- Physical reliability
- Link Performance

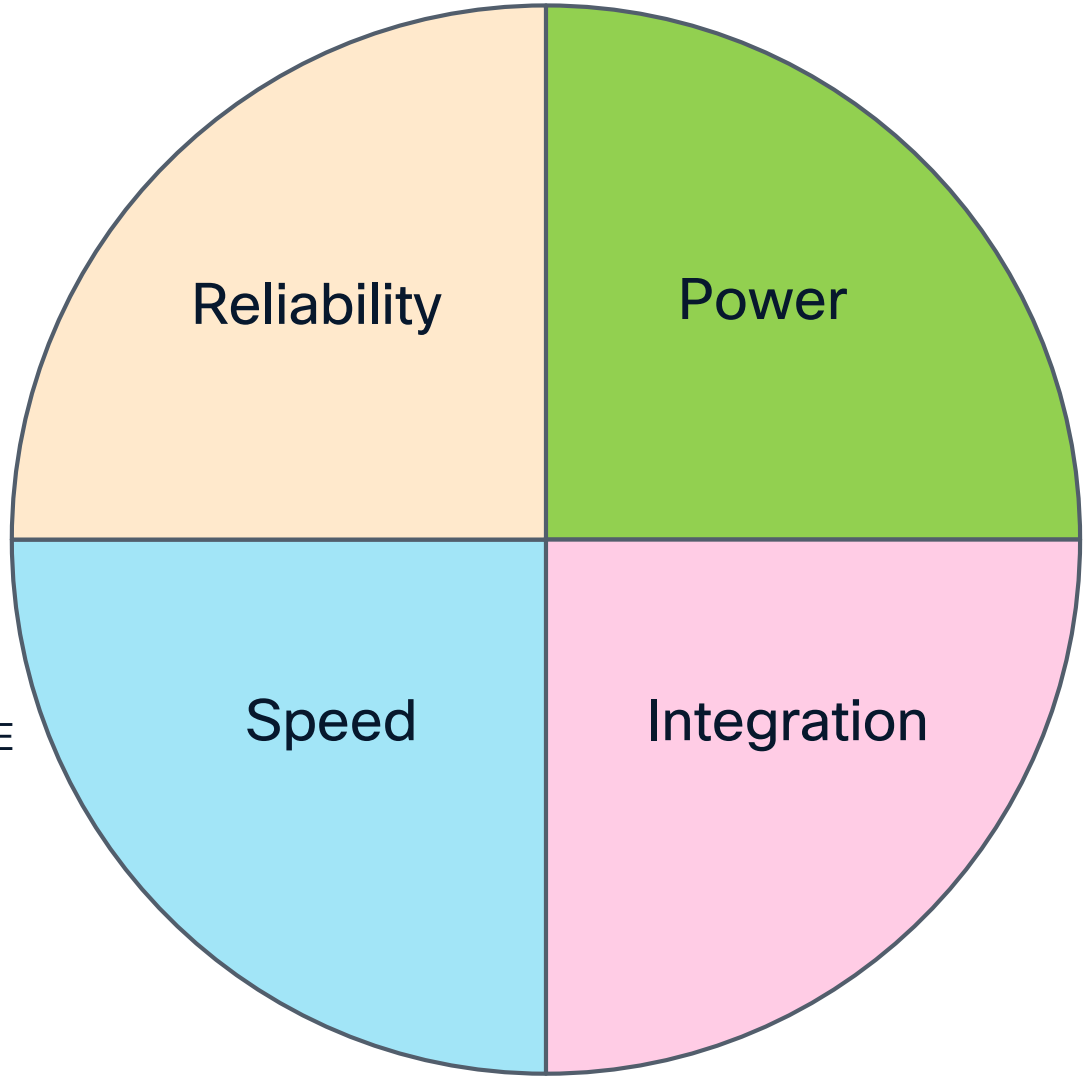

- Lower Watts per Gigabit
- Copper cables
- Linear Optics

Interfaces
800 GbE→1.6 TbE → 3.2 TbE

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

- Rack Density
- Switch Density

→ (Liquid cooling)

Reliability

Cisco is a Leader in Reliable Optics

Why Is reliability critical for AI

Failure of one Optic within a GPU cluster

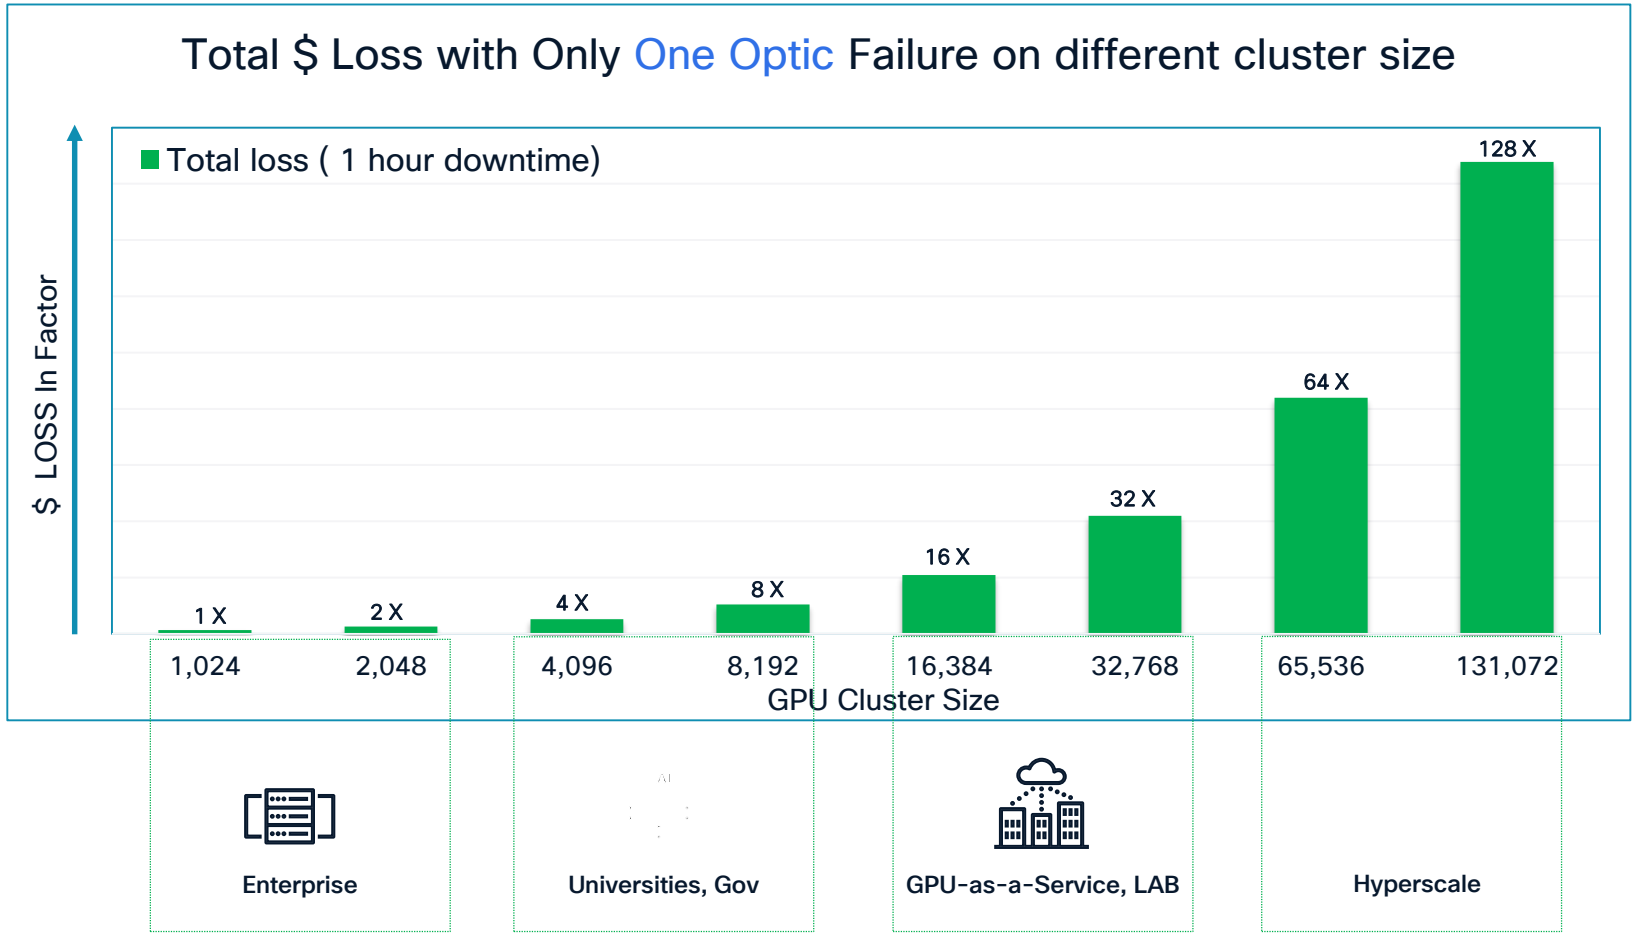

OUTCOMES

Significant Financial Impact – Higher TCO

Higher TCO due to replacements, maintenance, troubleshooting, and idle GPU

Cisco Silicon Photonics

A Foundation for Reliable AI Networks

Cisco's Leadership in Silicon Photonics

Compact Design	Integrates multiple functions into a smaller footprint, reducing complexity.
Fewer Components	Uses fewer lasers than traditional discrete designs, minimizing potential failure points and enhancing long-term reliability

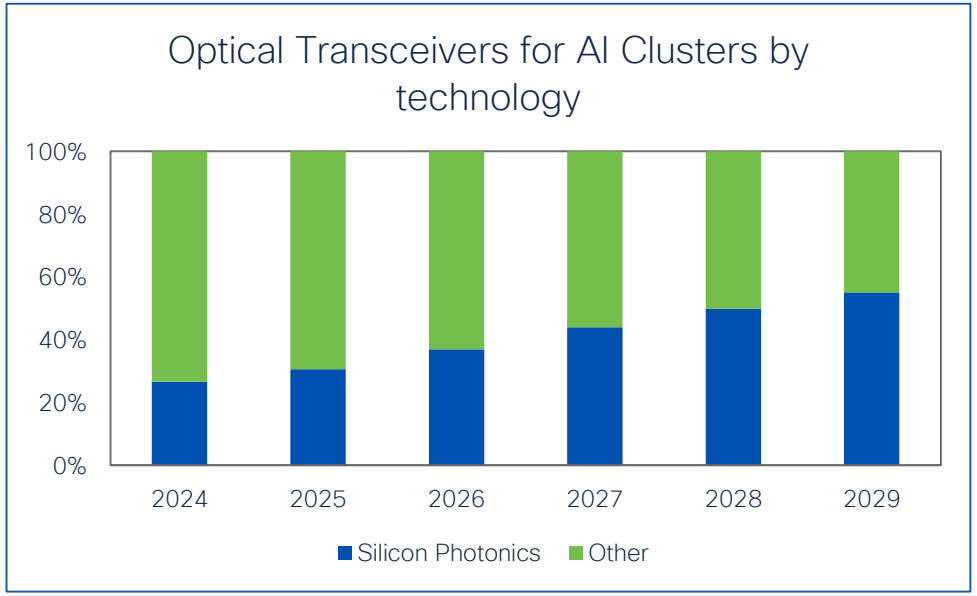

Source: Light Counting May 2024 Report:
Silicon Photonics, Linear Drive Pluggable and Co-Packaged Optics Report

Cisco brings years of high-volume silicon photonics manufacturing experience – building optics that are simpler and more robust

Maximum reliability in AI with 800G Silicon Photonics Optics

Cisco expertise in Silicon Photonics spans over 20 years and more than 5 million units

Silicon photonics technology allows to share laser sources, reducing the number of active components, and enhancing overall reliability compared to more discrete designs

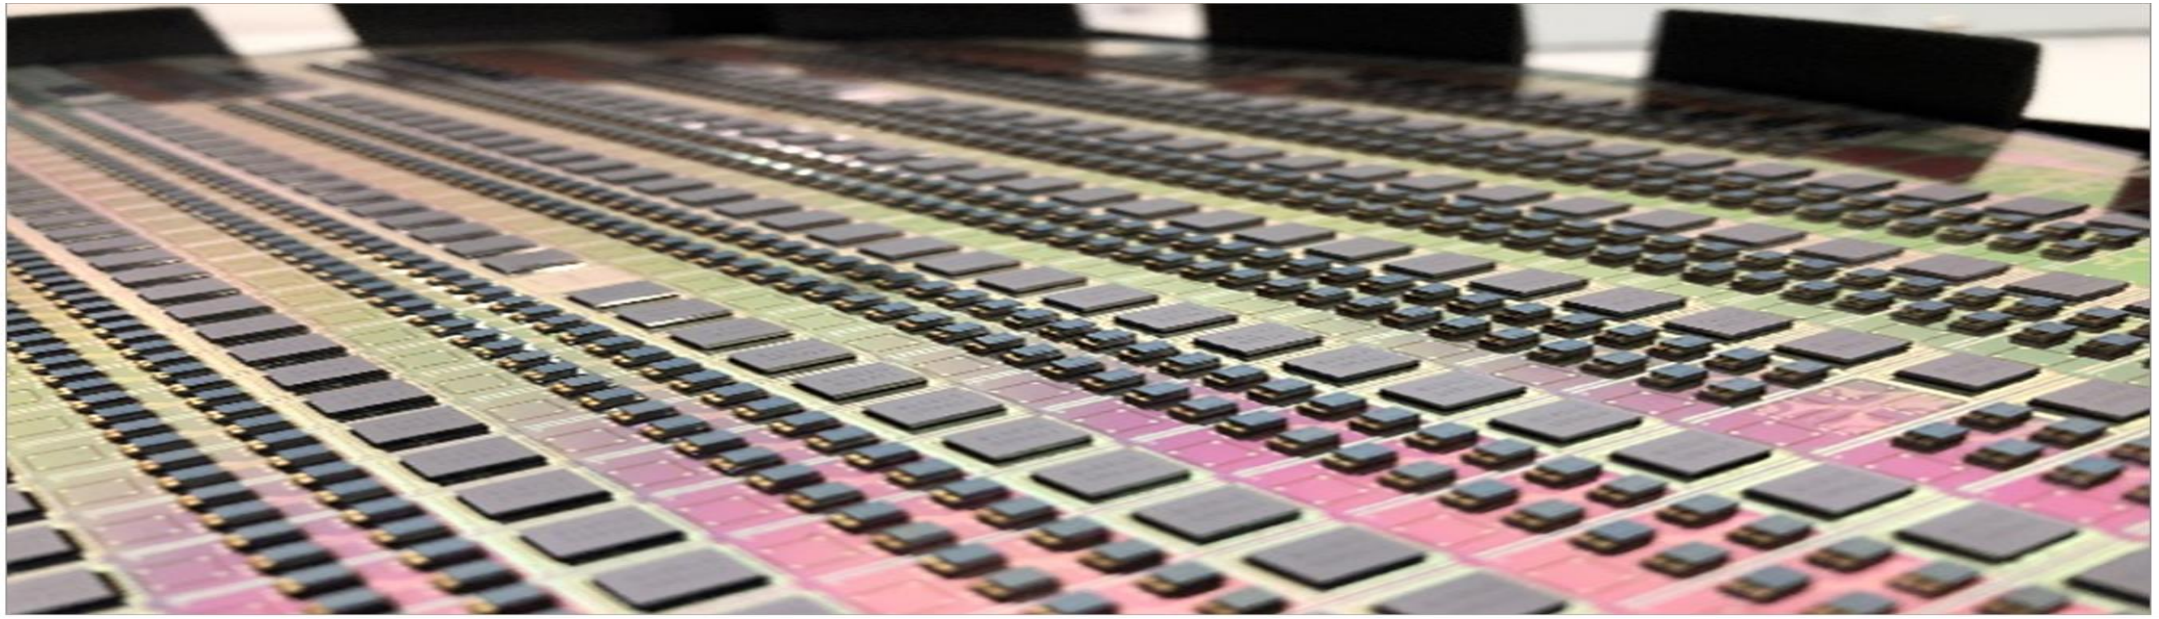

Resilient Optics for AI → Reliability/Performance, Links always on, Telemetry

Power & Integration

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

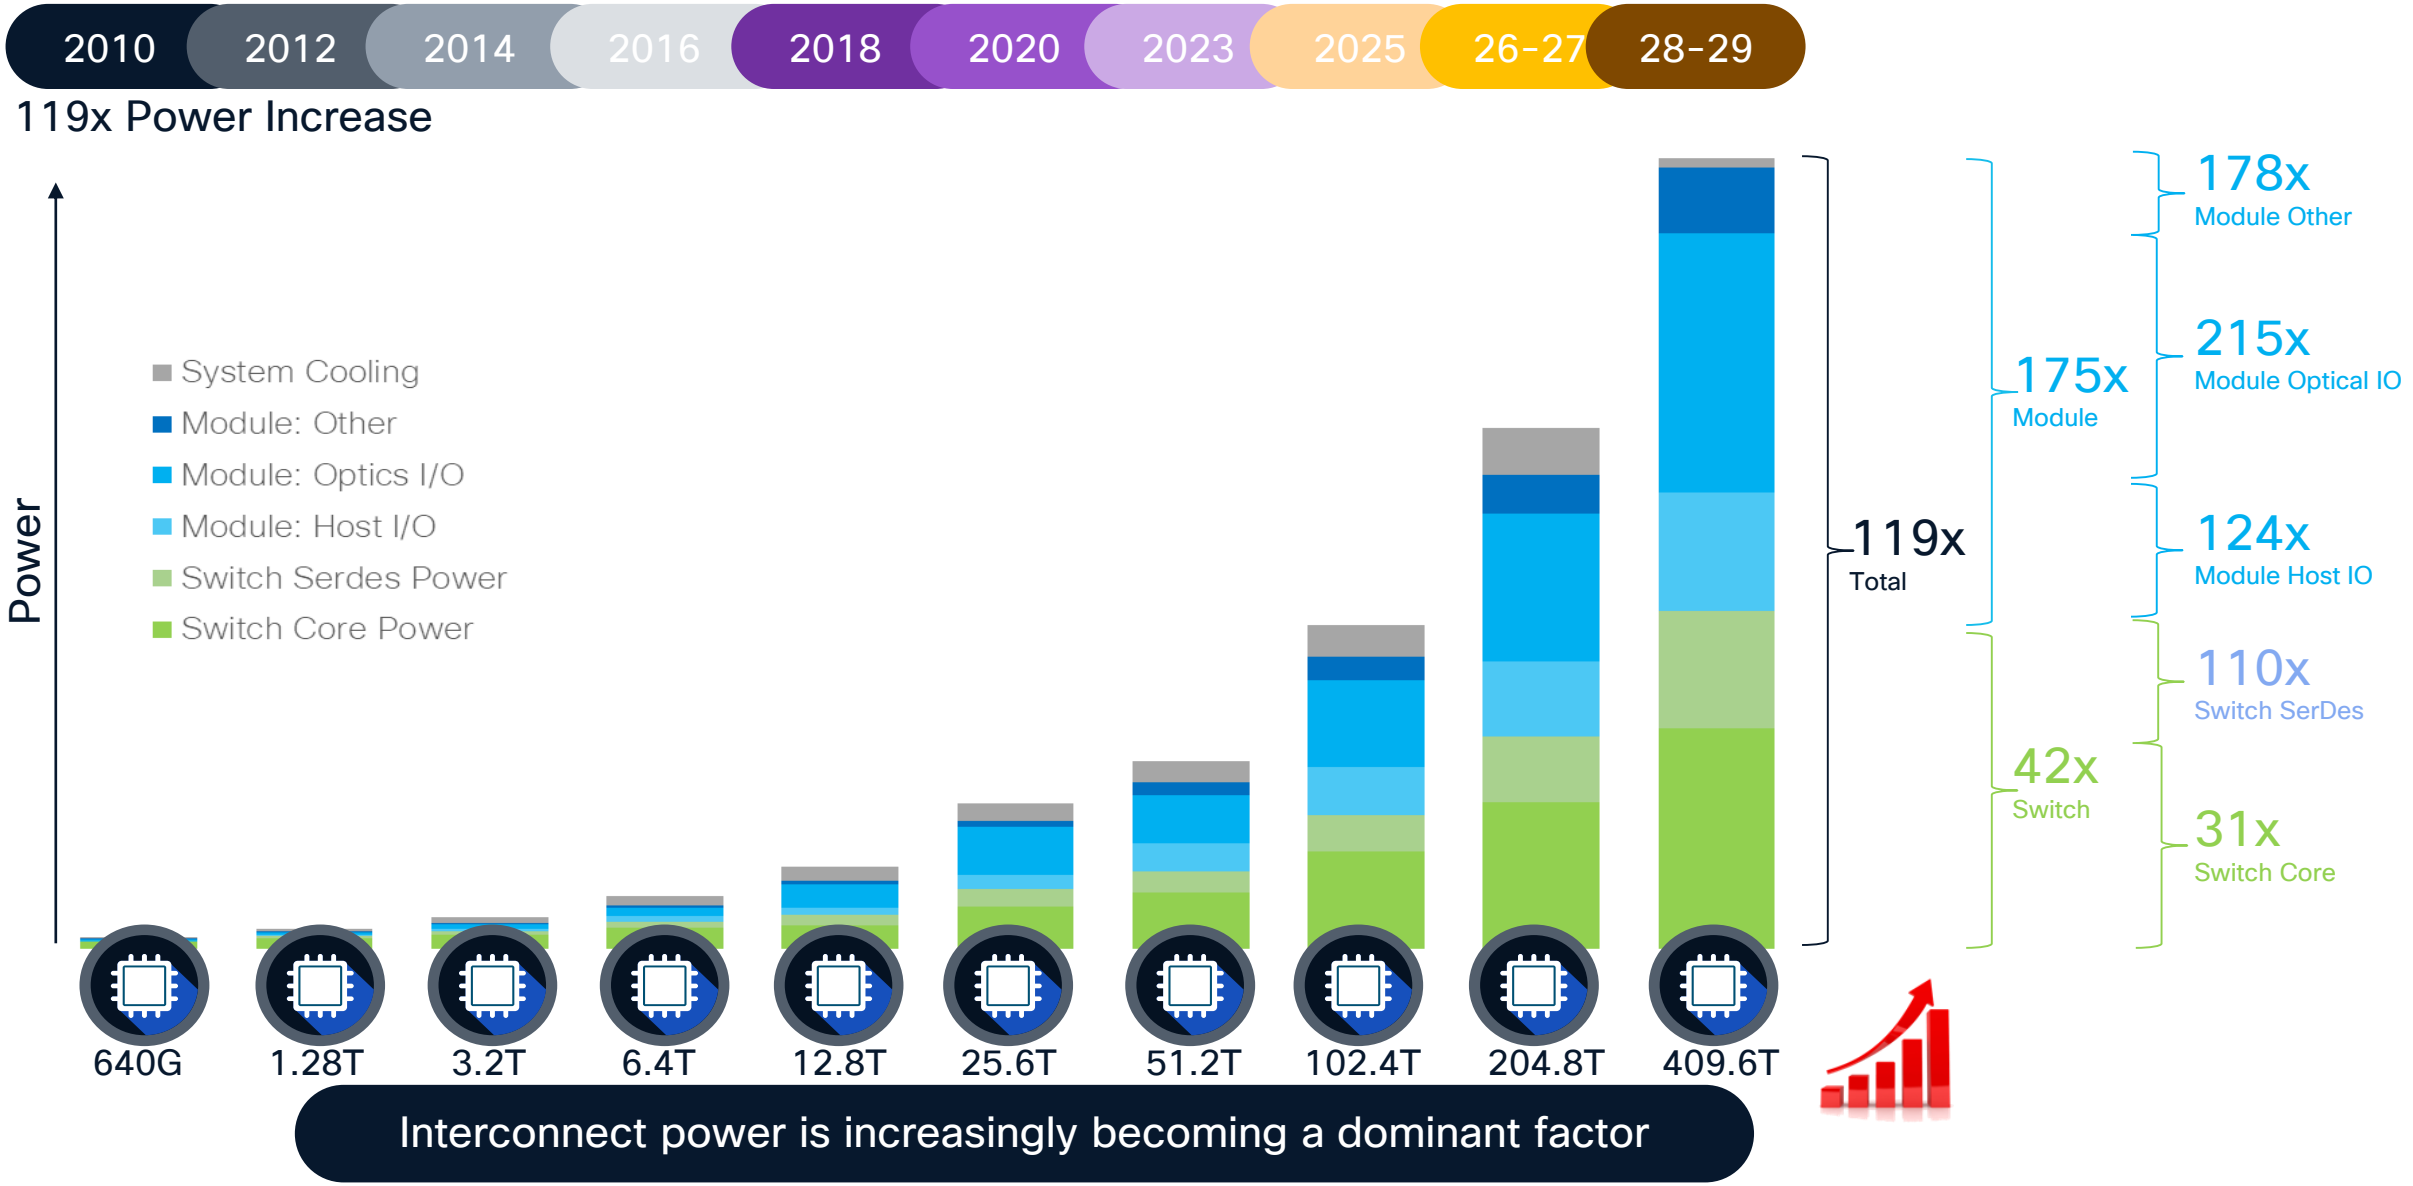

Lowering the power for interconnect

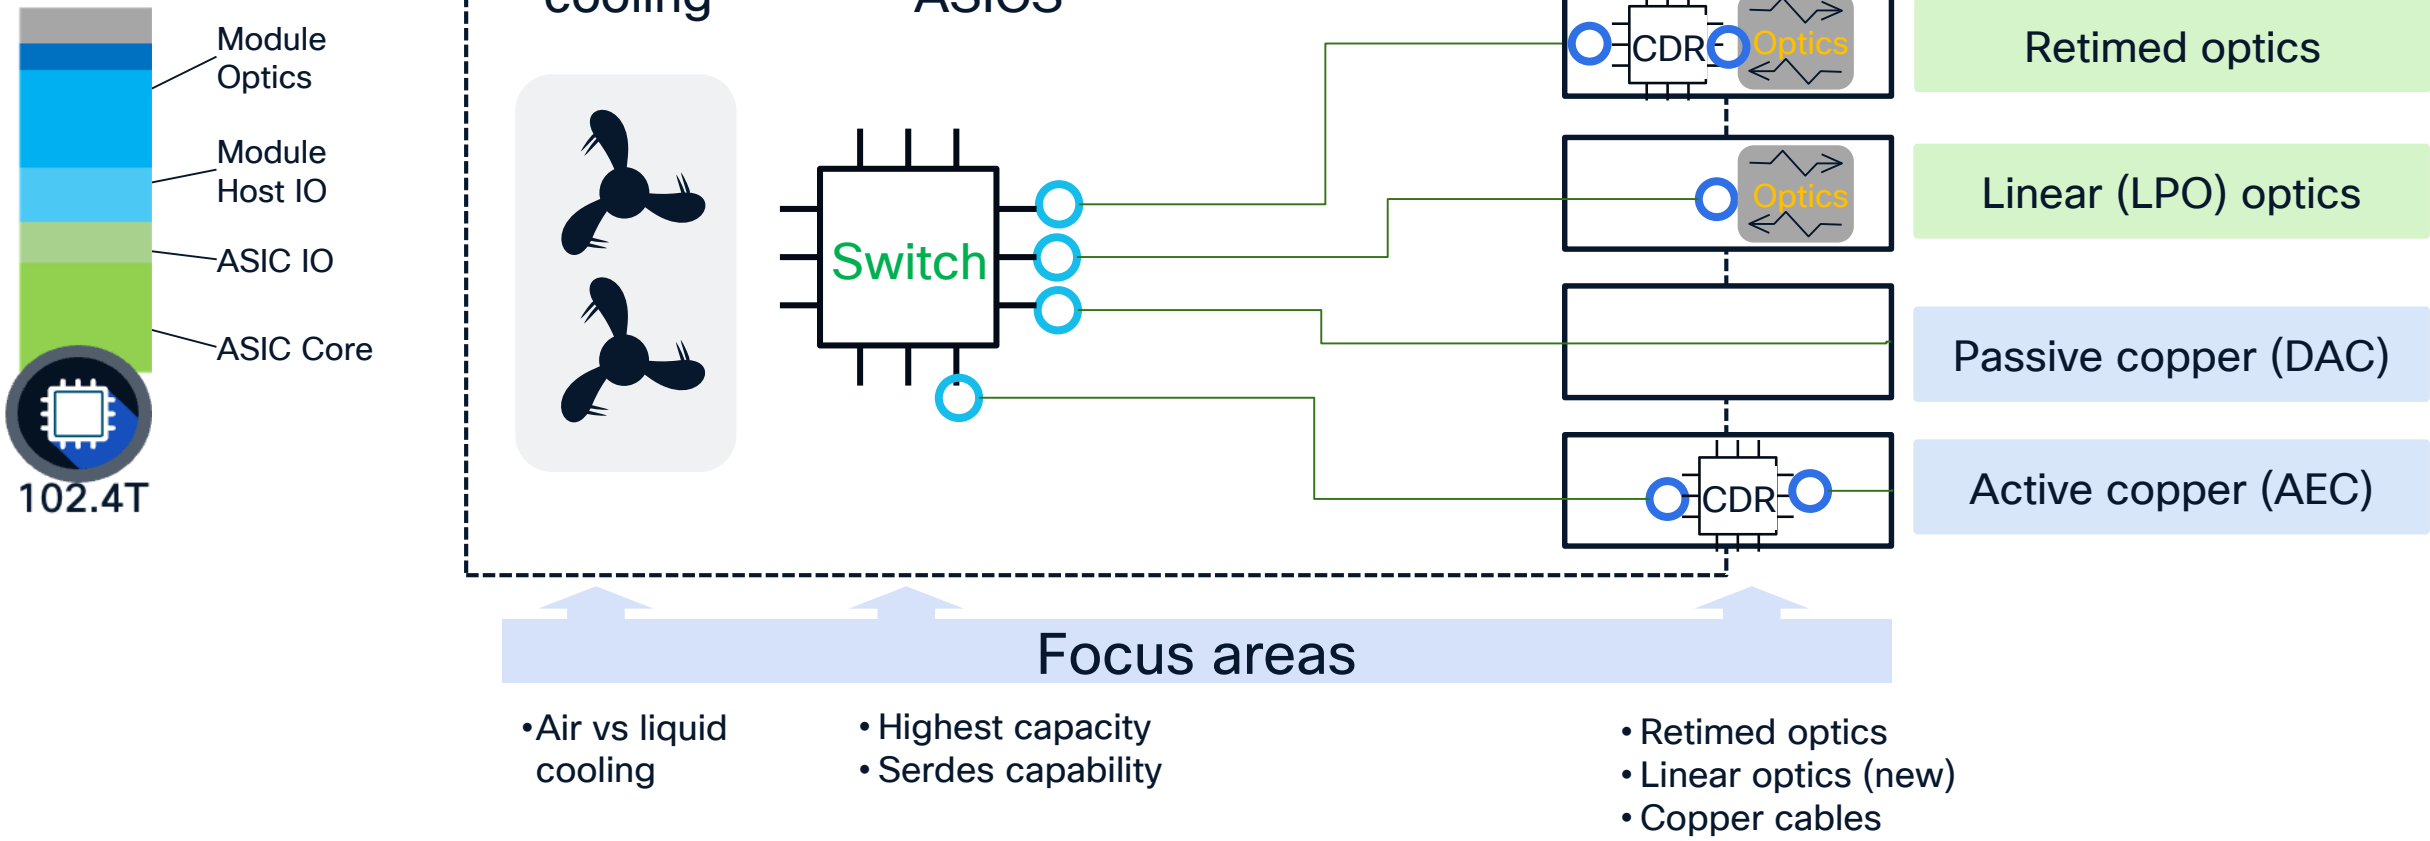

High-Density Interconnect options: Power vs. Reach Tradeoff

Power and Reach @ 800G

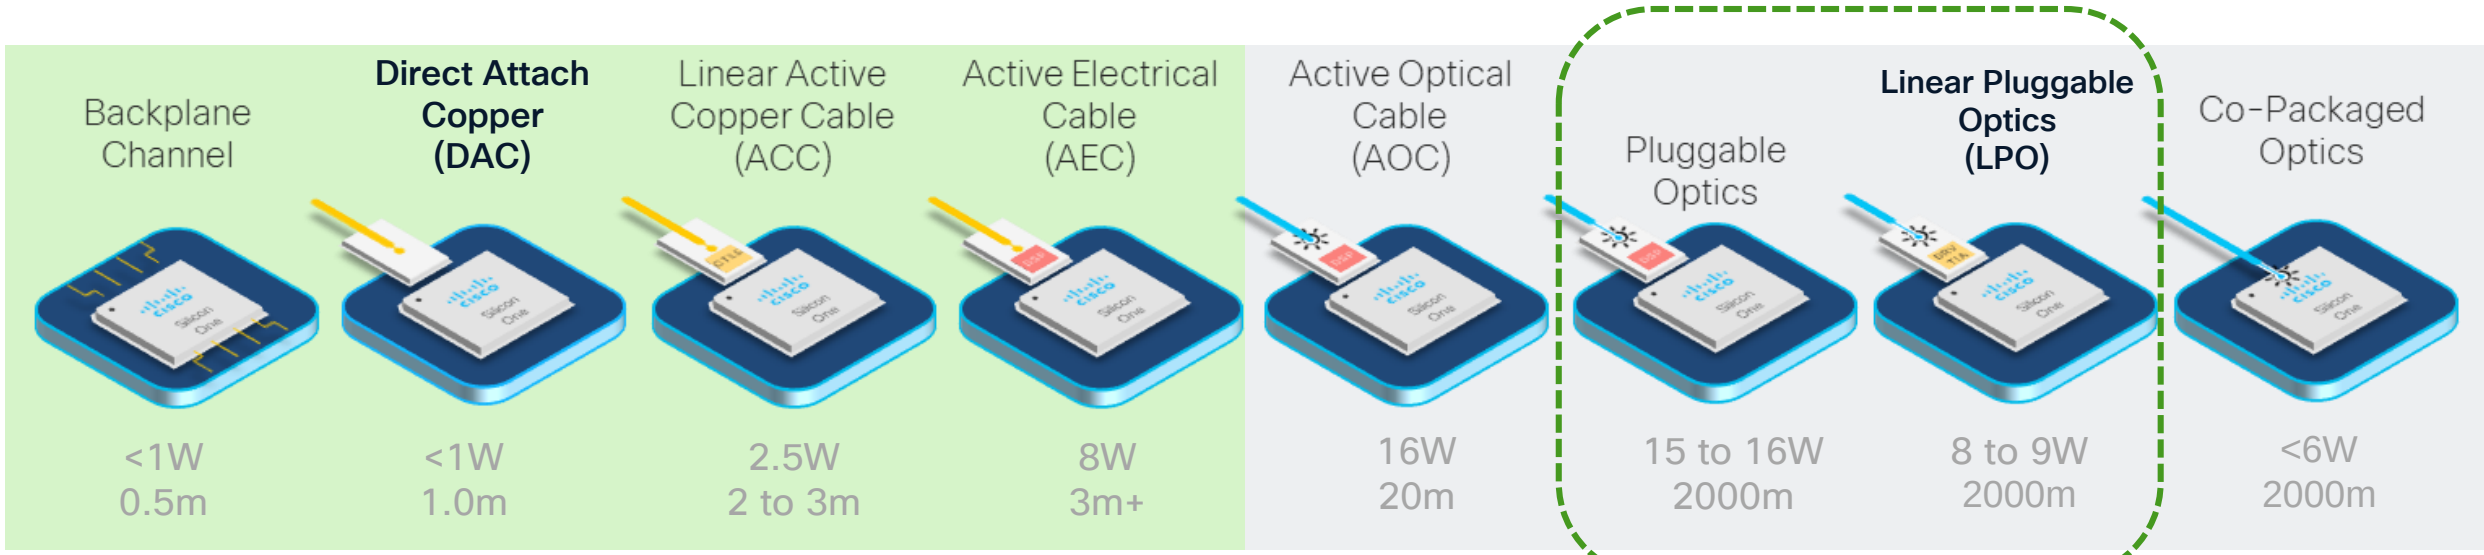

Maximize use within a rack

Use when inter-rack

LPO – Linear Pluggable Optics
- fully linear optics (Tx & Rx)

LRO – Linear Receive Optics
- DSP on transmit side

Linear Optics
- very attractive as a solution to help with the AI deployment power challenge
- But it needs an interoperable ecosystem – LPO MSA

What is Linear pluggable optics (LPO)?

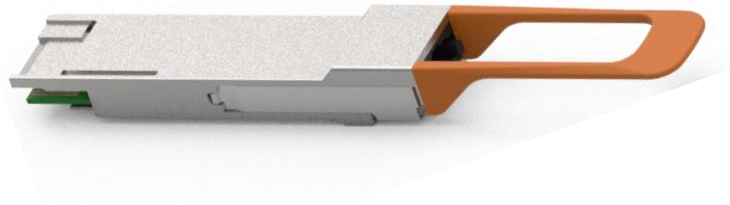

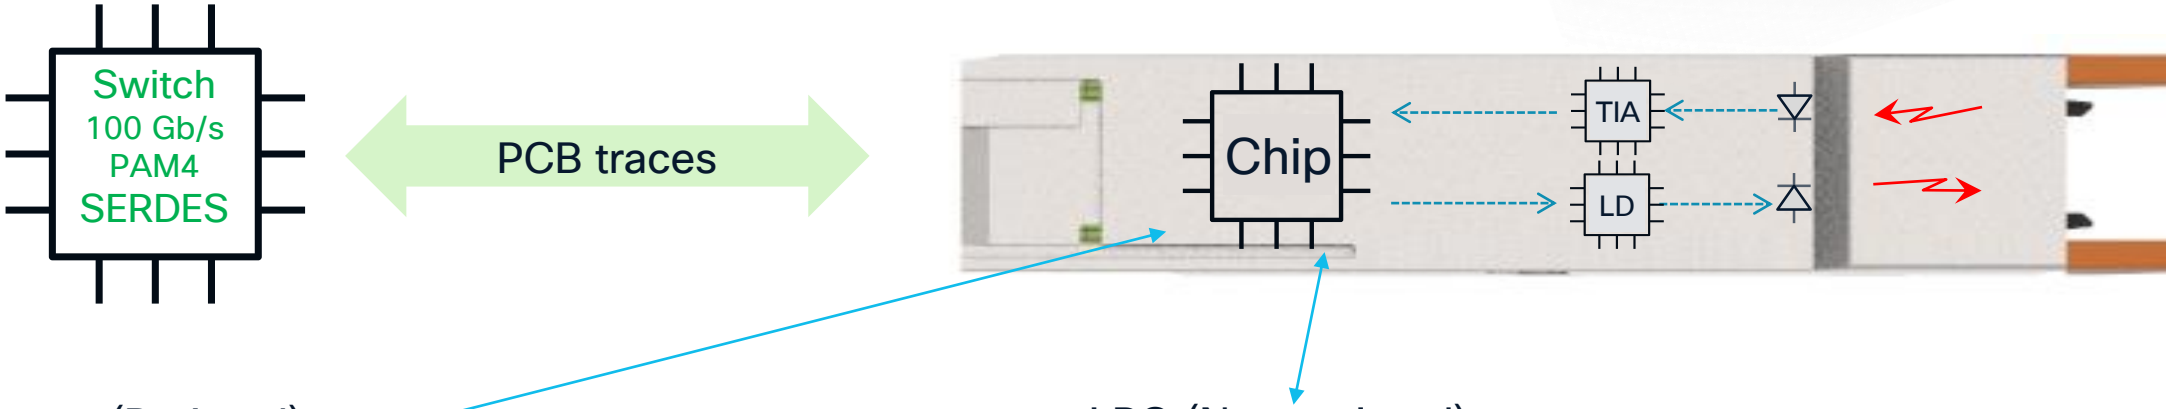

Current (Retimed)

- 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 (Non-retimed)

- **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?

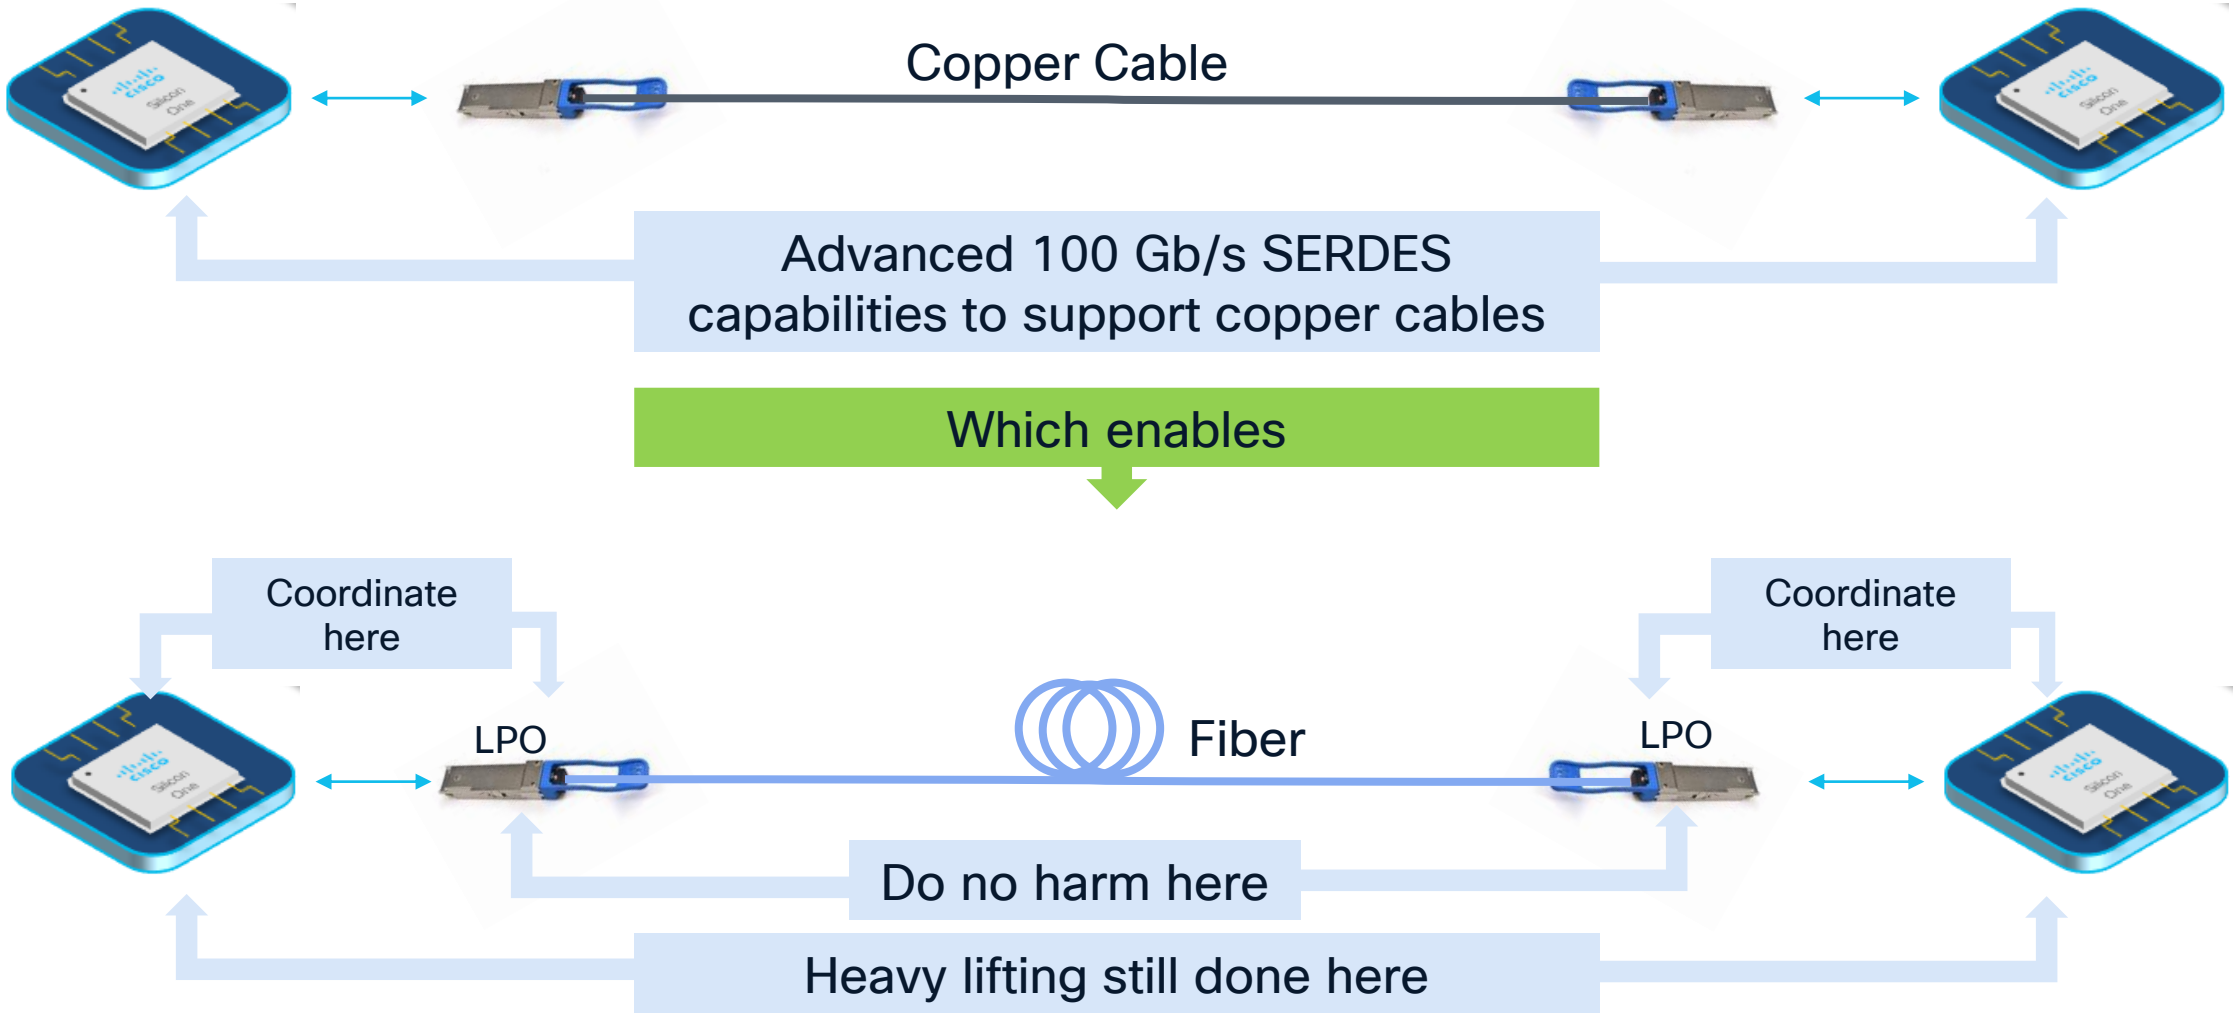

How does LPO work?

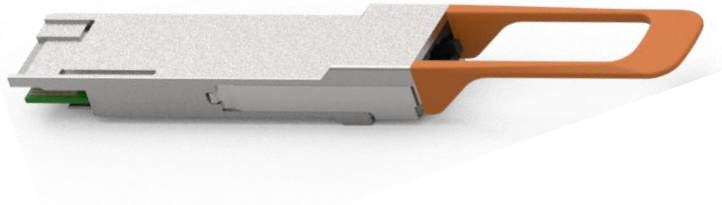

<u>ASIC Serdes are the key</u>

This is the enabling change!

- Very capable ASIC serdes Tx and Rx equalization capabilities
- Serdes are designed and capable to drive linear CR channel (copper cable)
- Cable (CR) channels are specified (IEEE) with lower host loss ports
- LPO modules need to work in every switch port, so they need to compensate for the extra host loss in those high loss ports
- LPO modules have limited linear gain and equalization that can compensate for host channel
- Host and module handshake to adjust equalizer settings
- <u>Linear</u> sub-components are key – <u>Silicon Photonics helps.</u>

Cisco’s Linear Pluggable Optics (LPO) Solution

Future-ready infrastructure for AI

Traditional Solution (with DSP)

Retimed Optics

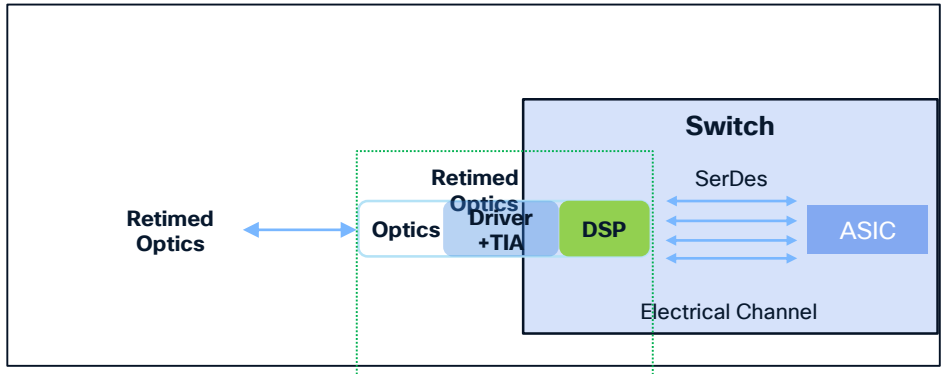

Linear Pluggable Optics (LPO)

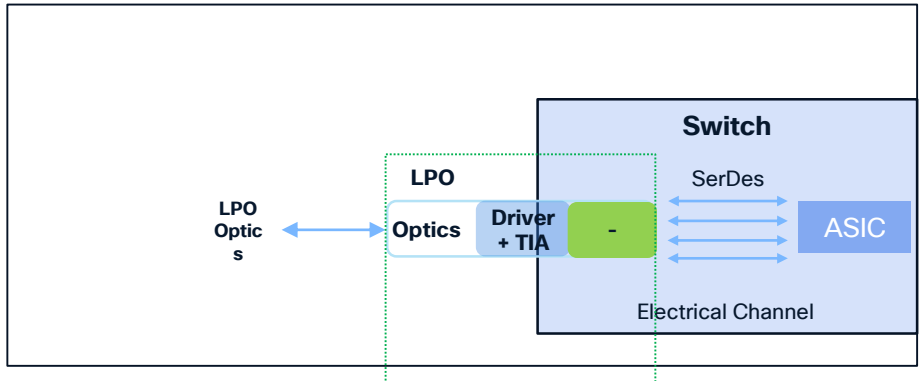

Feature	Retimed Optics	LPO Optics
DSP in Module	☑ Yes	✕ No
Power Consumption	High (16 W)	Low (8W)
Latency	Moderate (93 ns)	Ultra-low (4 ns)
Signal Conditioning	In-module DSP	Host ASIC

Cisco Silicon One G200: Retimed vs. LPO System Power Comparison

Cisco demonstrated LPO operation at OFC2024

<u>Set-up:</u>
Two identical Silicon One 51.2Tb 64-port G200-based switches:
- 100% retimed optics
- 100% LPO optics

<u>Result:</u>
Both switches ran full traffic on all ports
Overall power reduction: ~ 700W

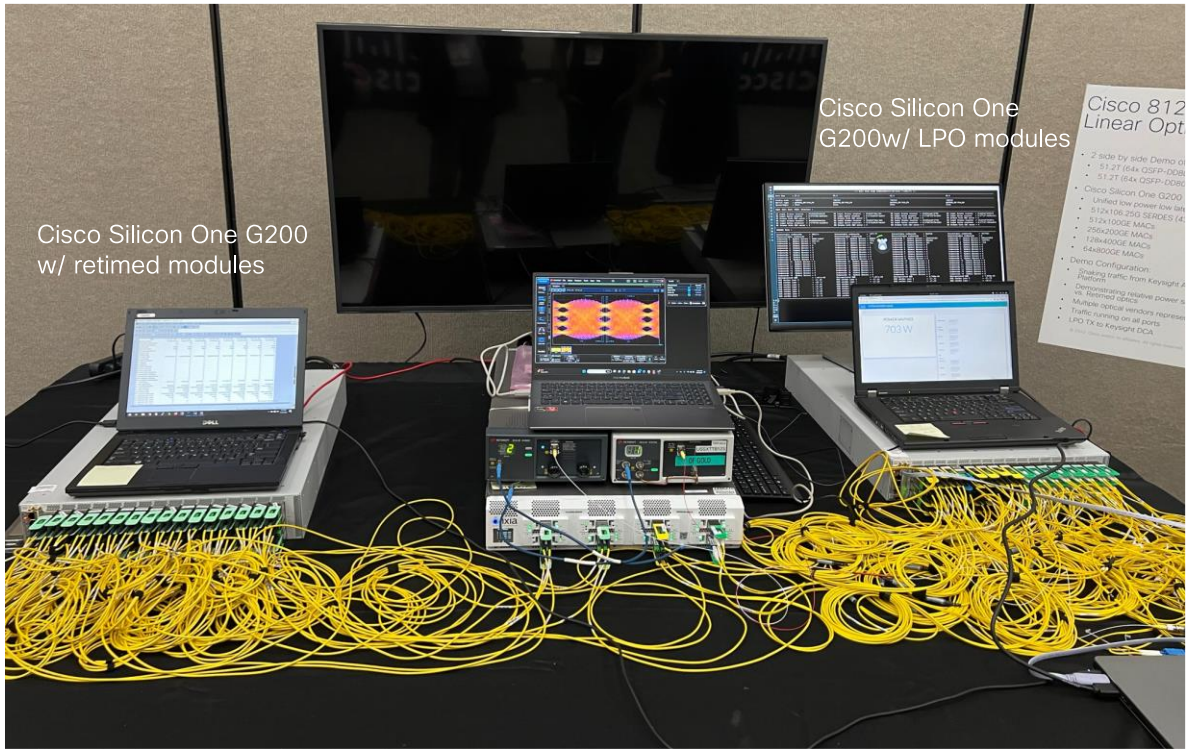

Evolution of pluggables

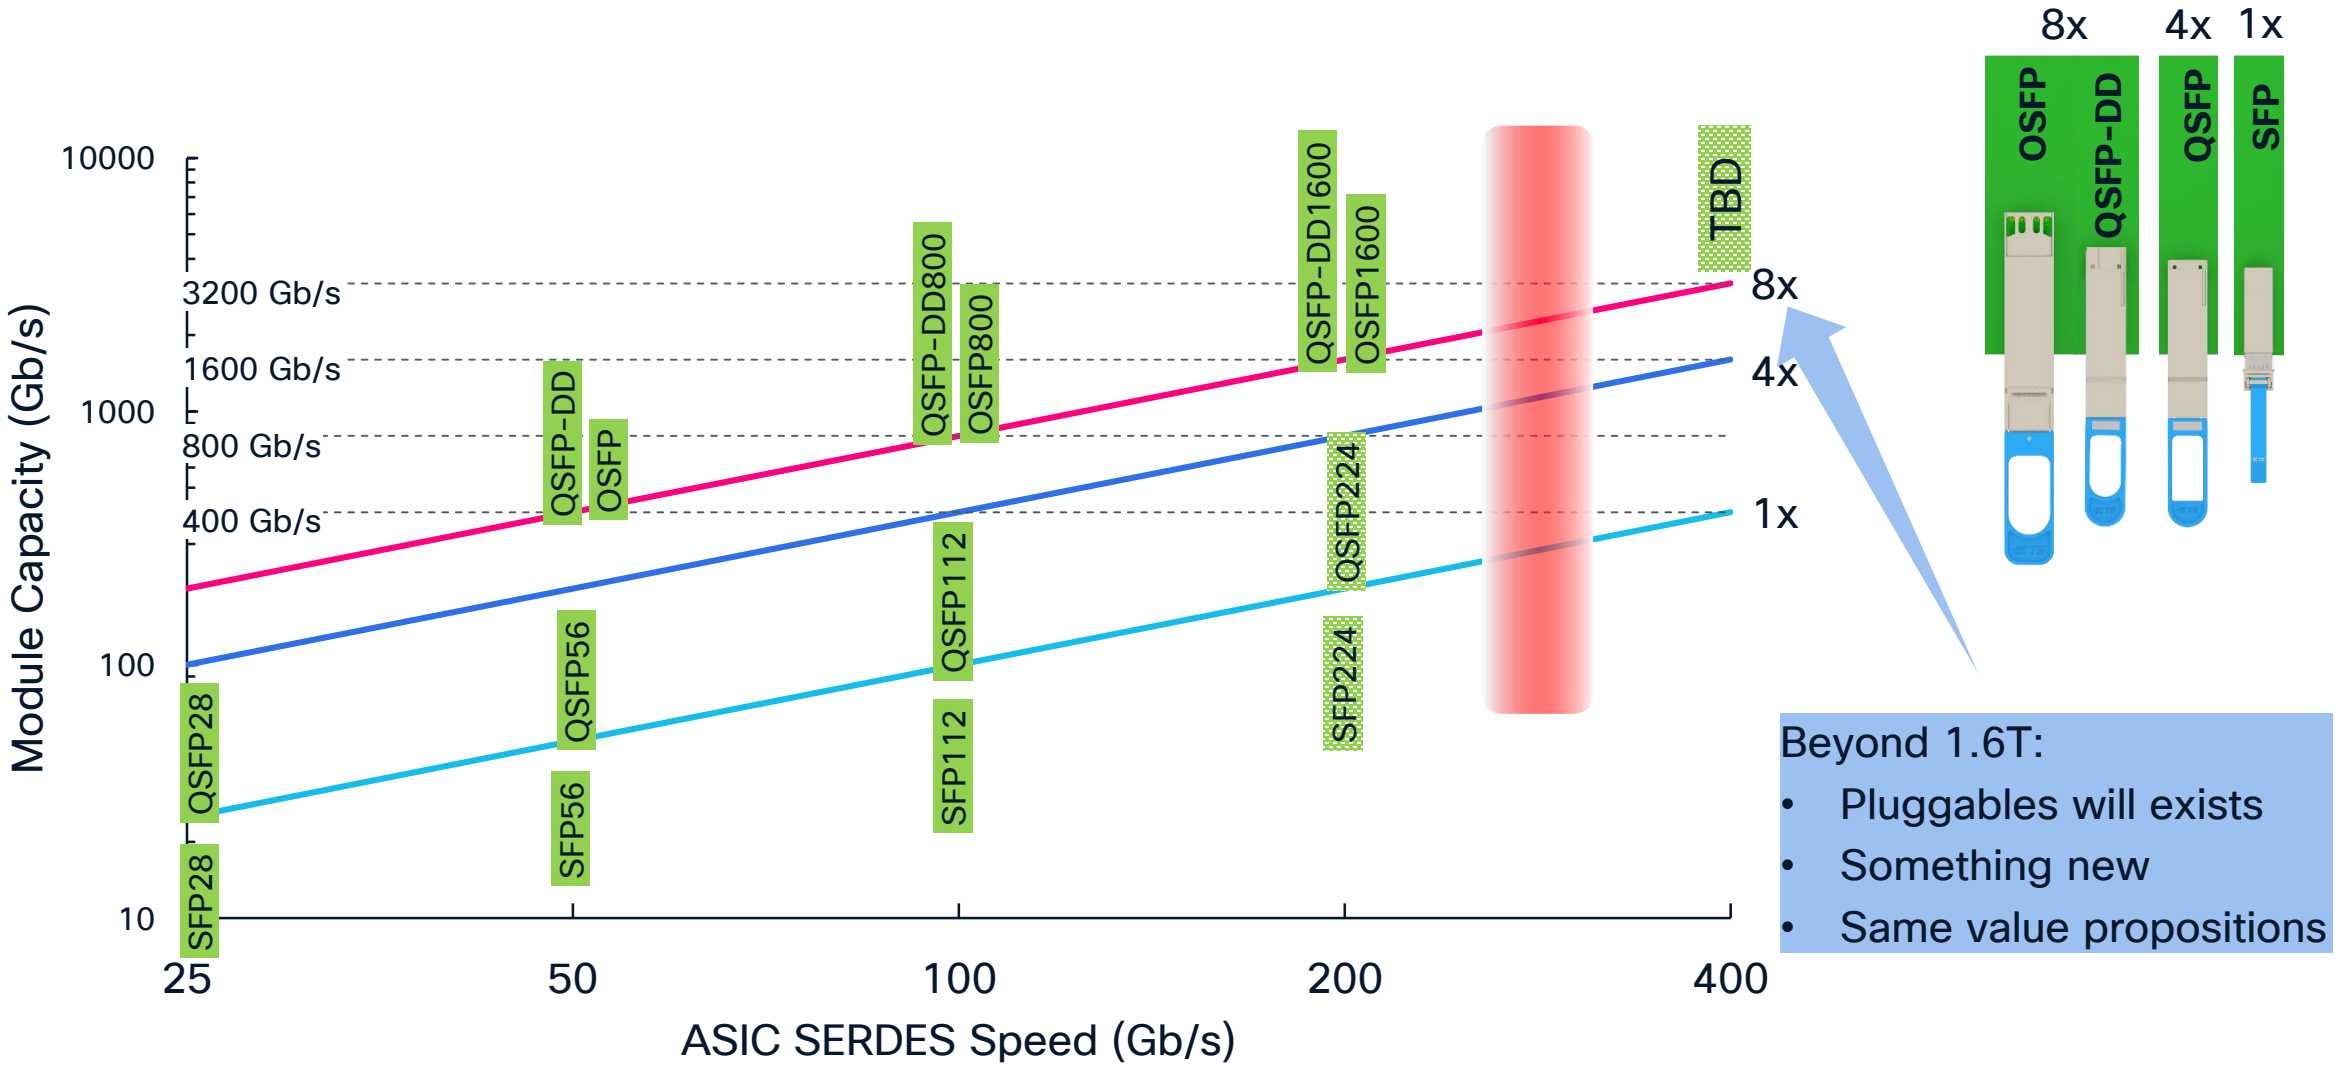

Beyond 1.6T:

- 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

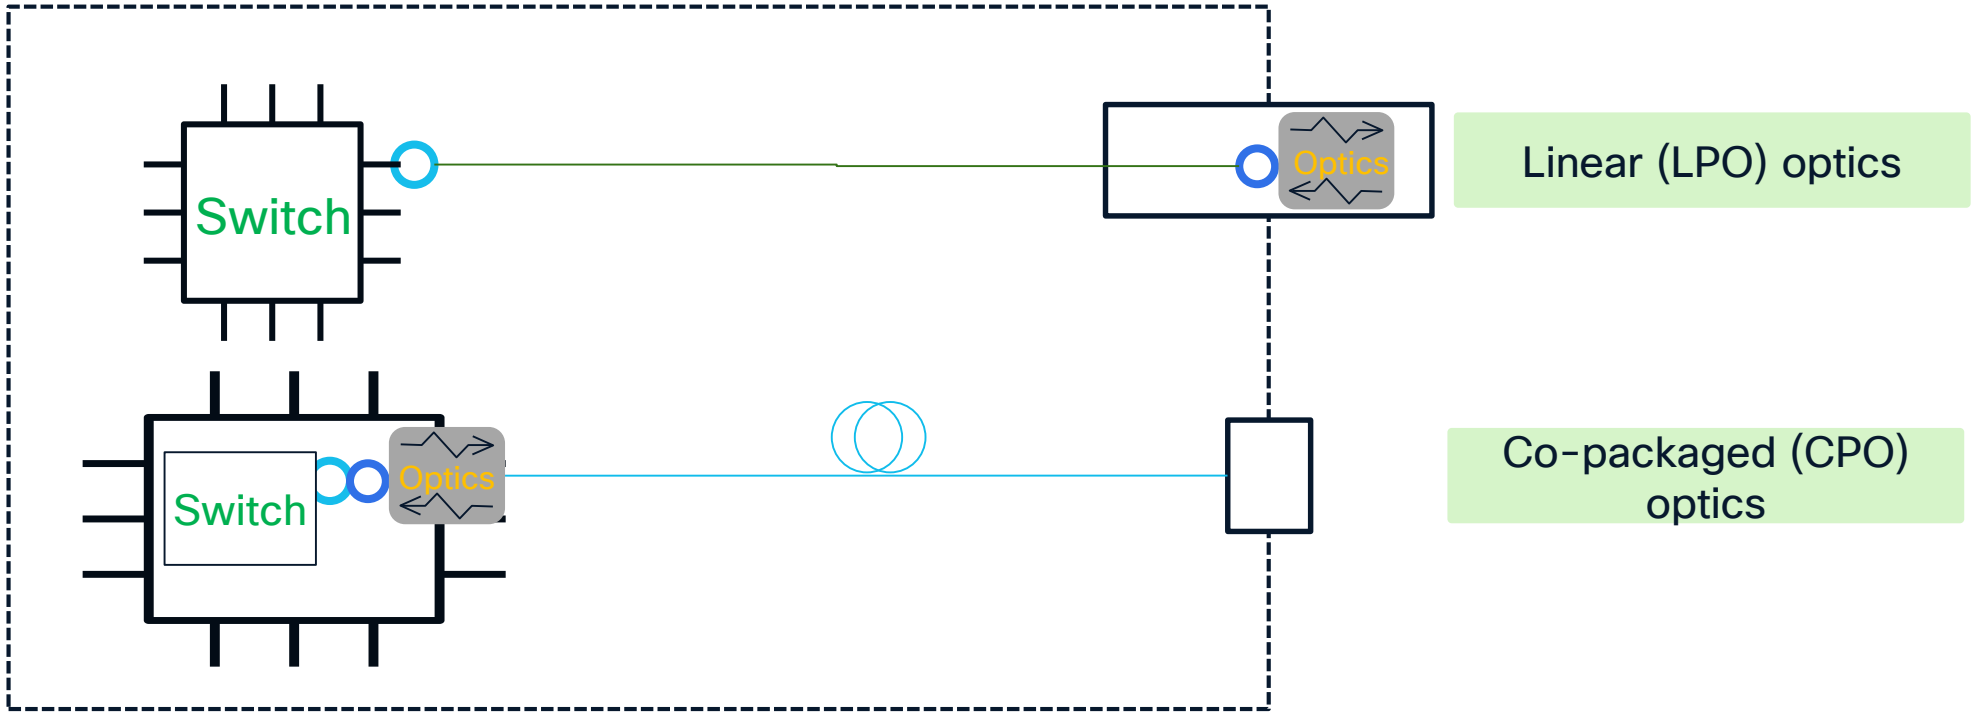

Cisco Silicon One: OFC 2023 Demo – 25.6T
Co-packaged Optics vs Retimed

CPO system achieved ~ 22% total system power reduction

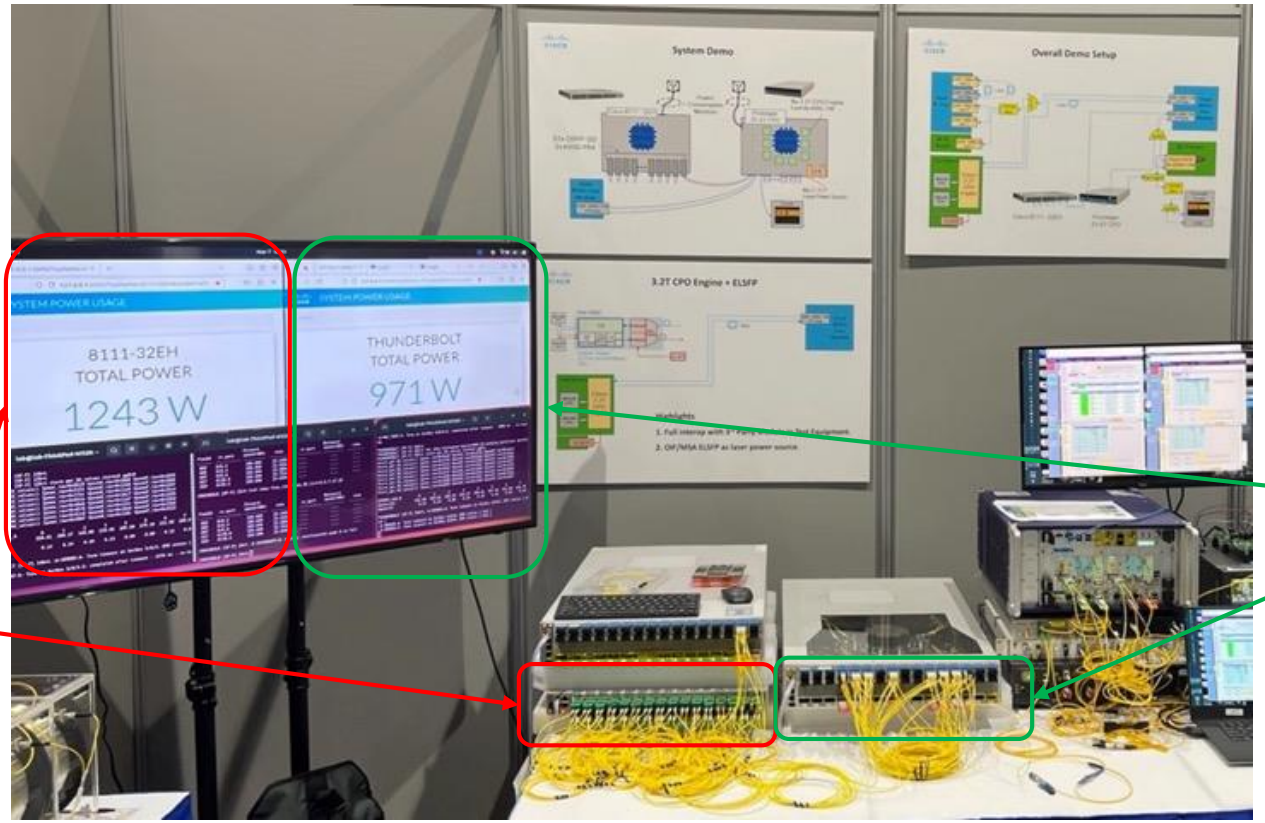

25.6T Ethernet Switch
64 ports @ 400G QSFP-DD

25.6T Ethernet Switch
Enabled by CPO and using ELSFP external light sources

Pluggables and co-packaged optics with co-exist

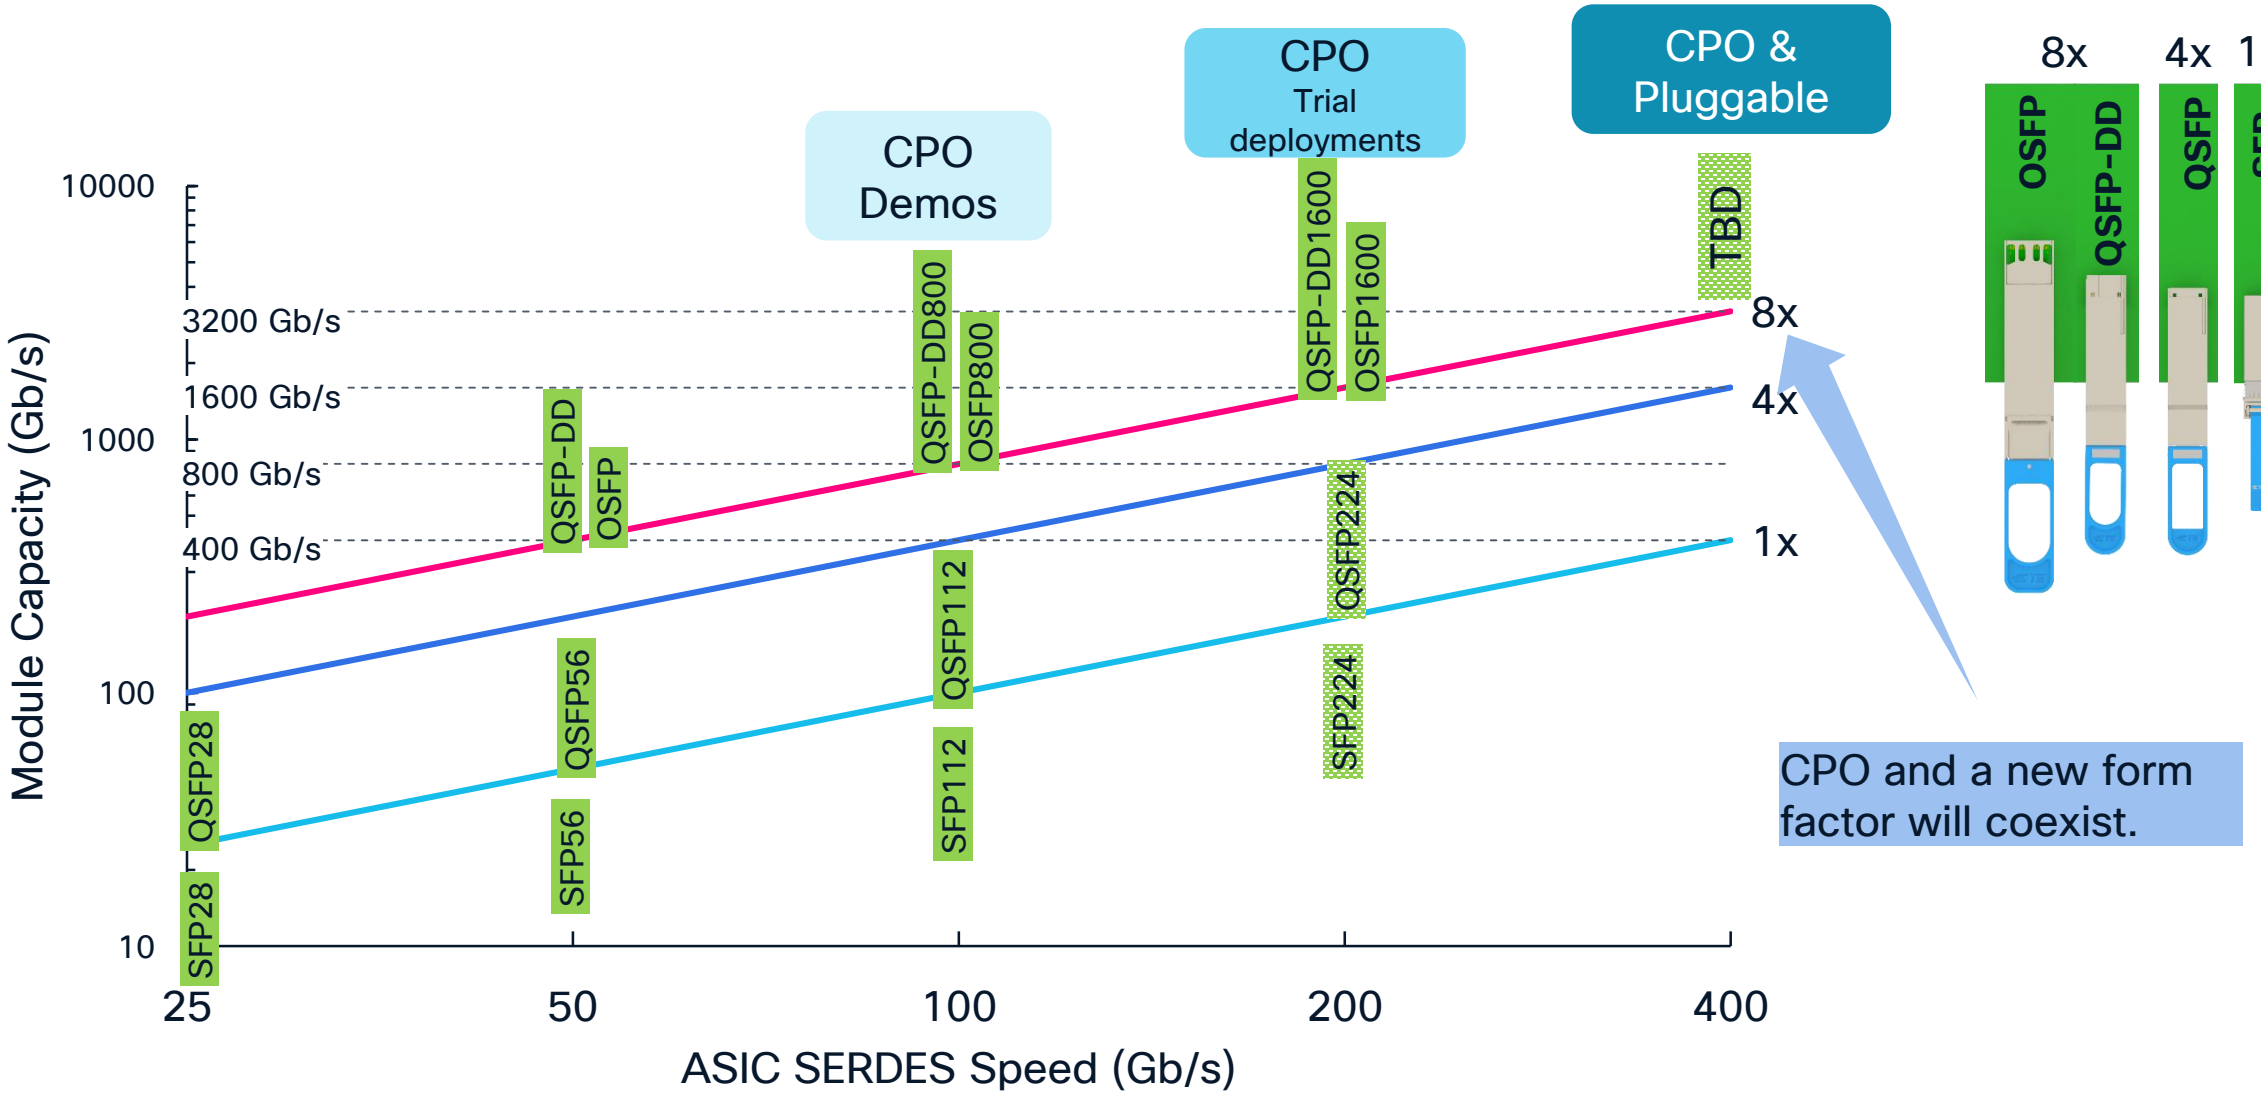

Speed

Relentless ASIC Advancement

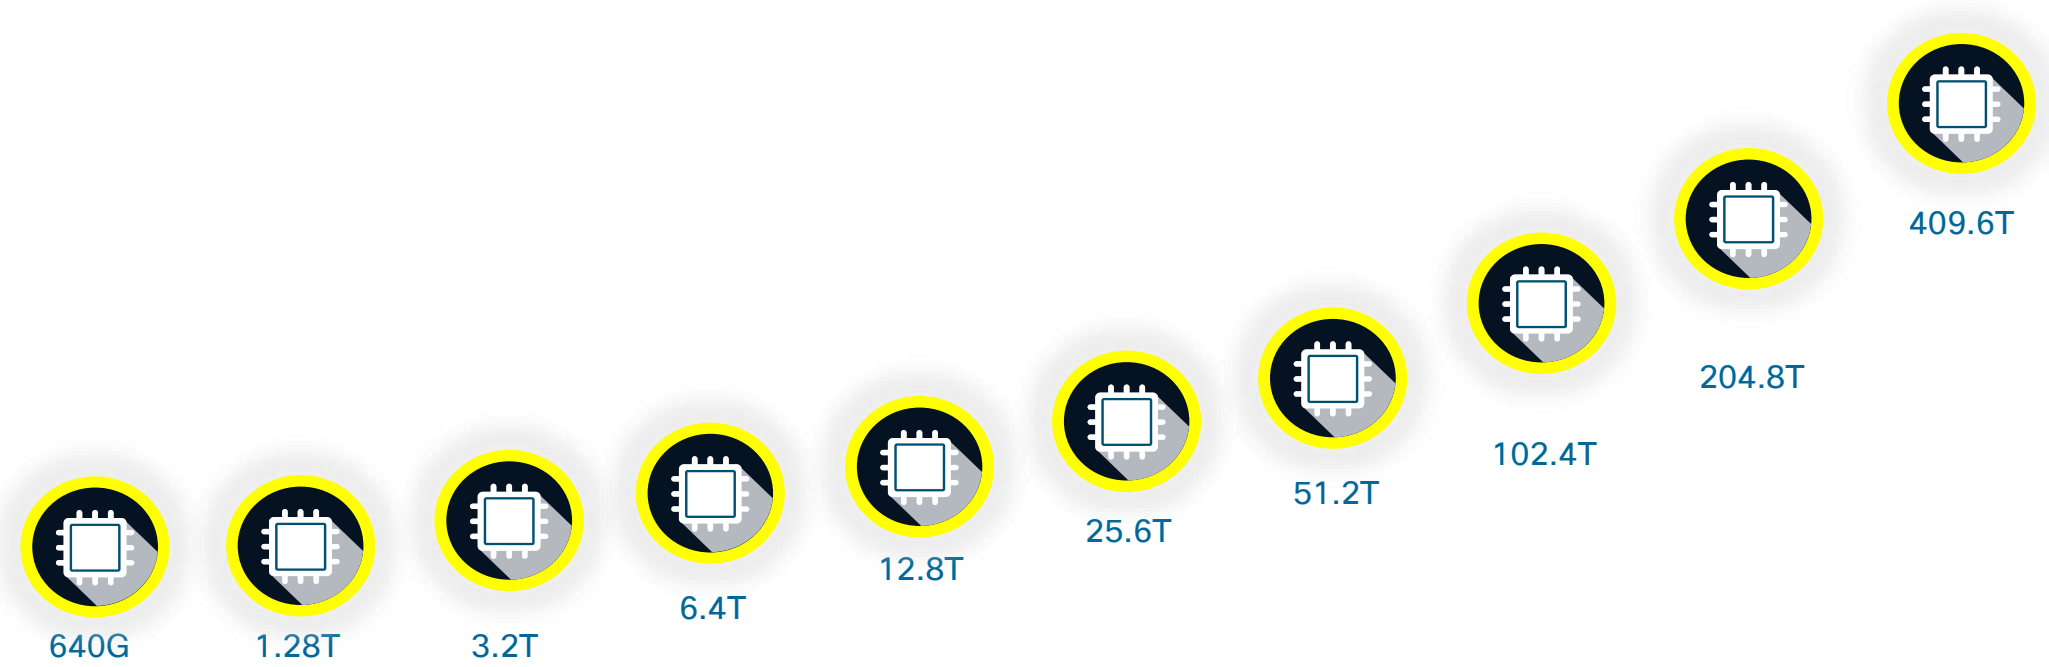

ASIC # SerDes	64			128			256			512					1024						
Serdes Gbps	10G NRZ				25G NRZ				50G NRZ				100G PAM4			200G PAM4				400G	
Port Gbps	40G				100G				400G				800G			1.6T				3.2T	
Ports/Silicon	16			32			64	32		64			64		64			128			

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

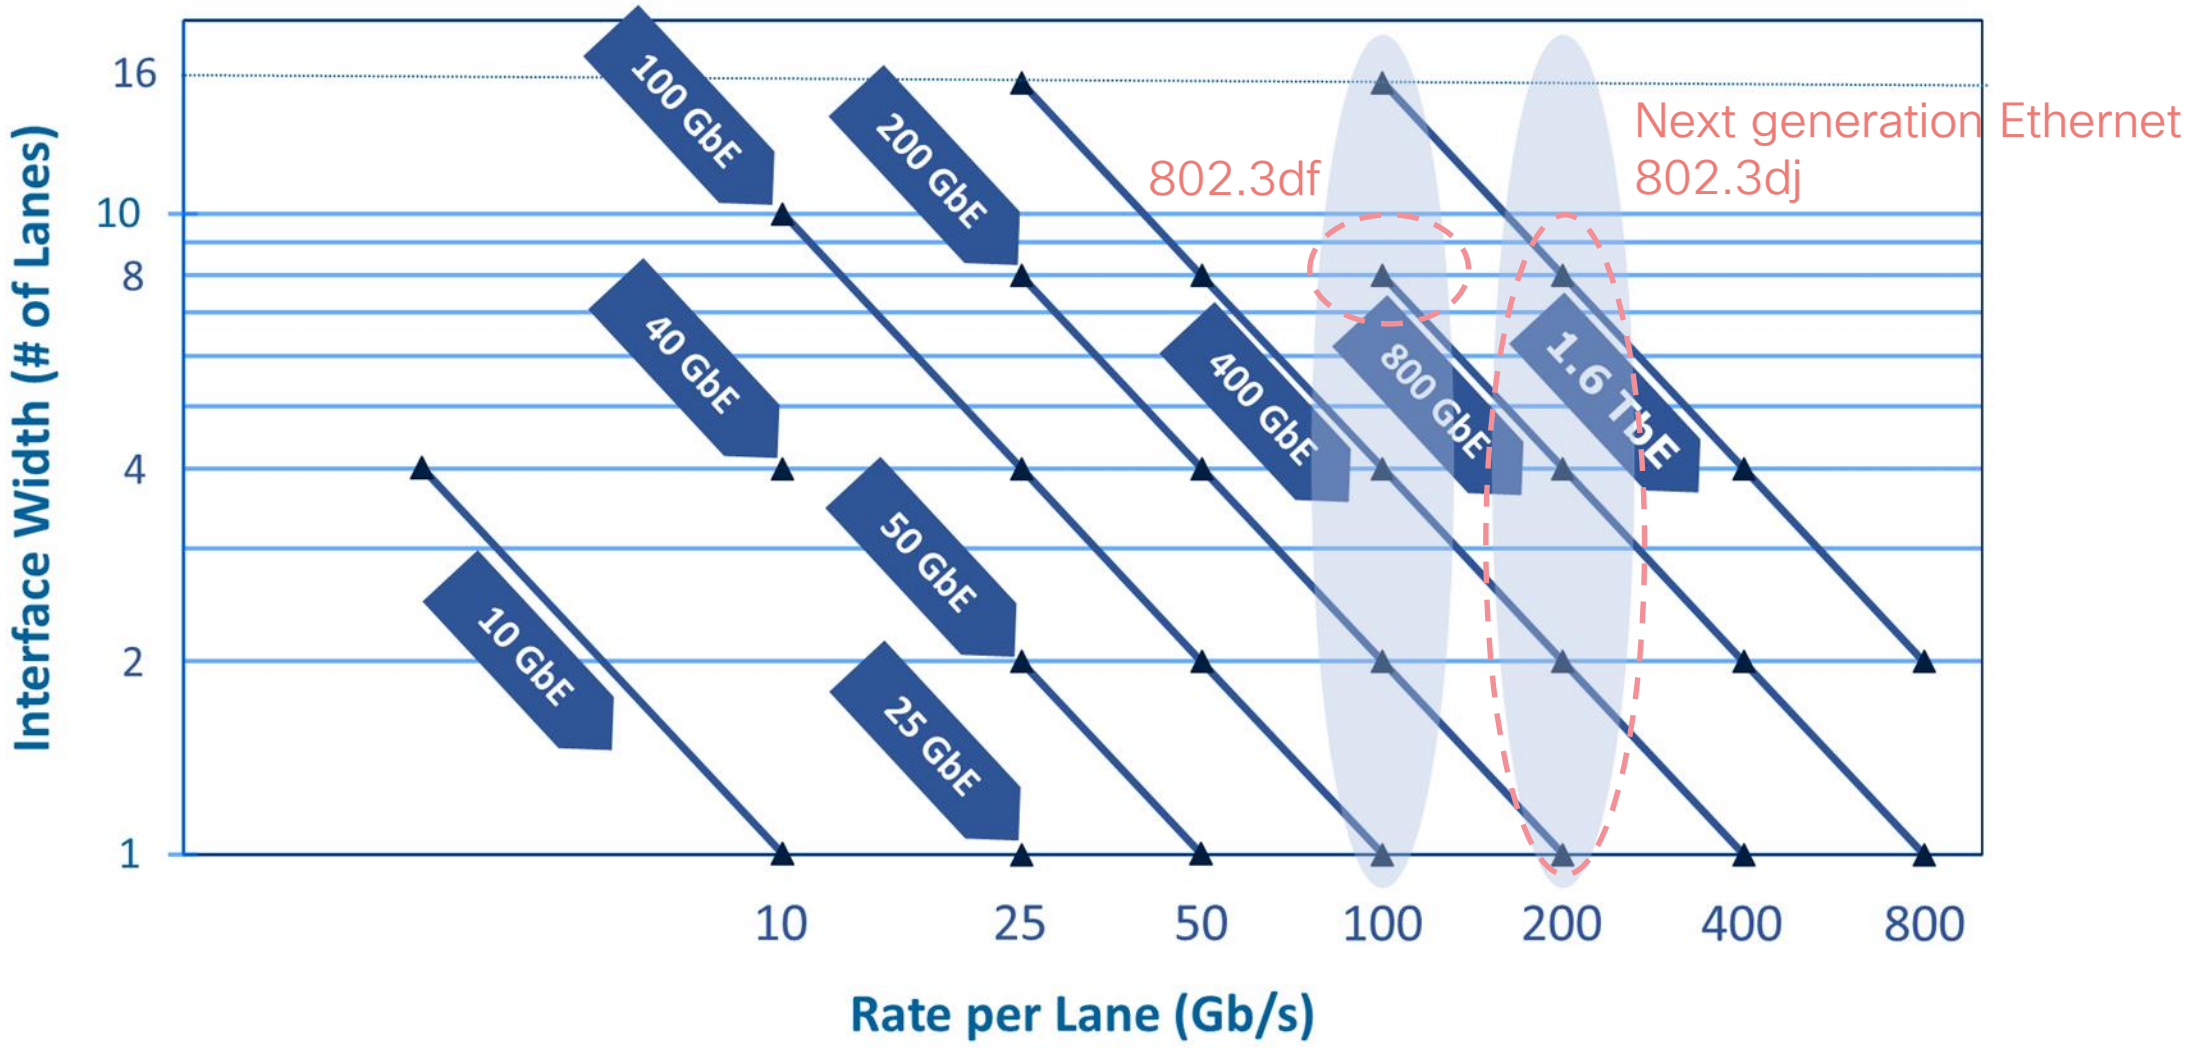

Ethernet's next generation: IEEE 802.3 P802.3dj → 200 Gb/s per lane for 1.6 TbE and below

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?

Continued flexibility of implementations

- Copper cables
- Multimode Fiber – 100m (future project)
- 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

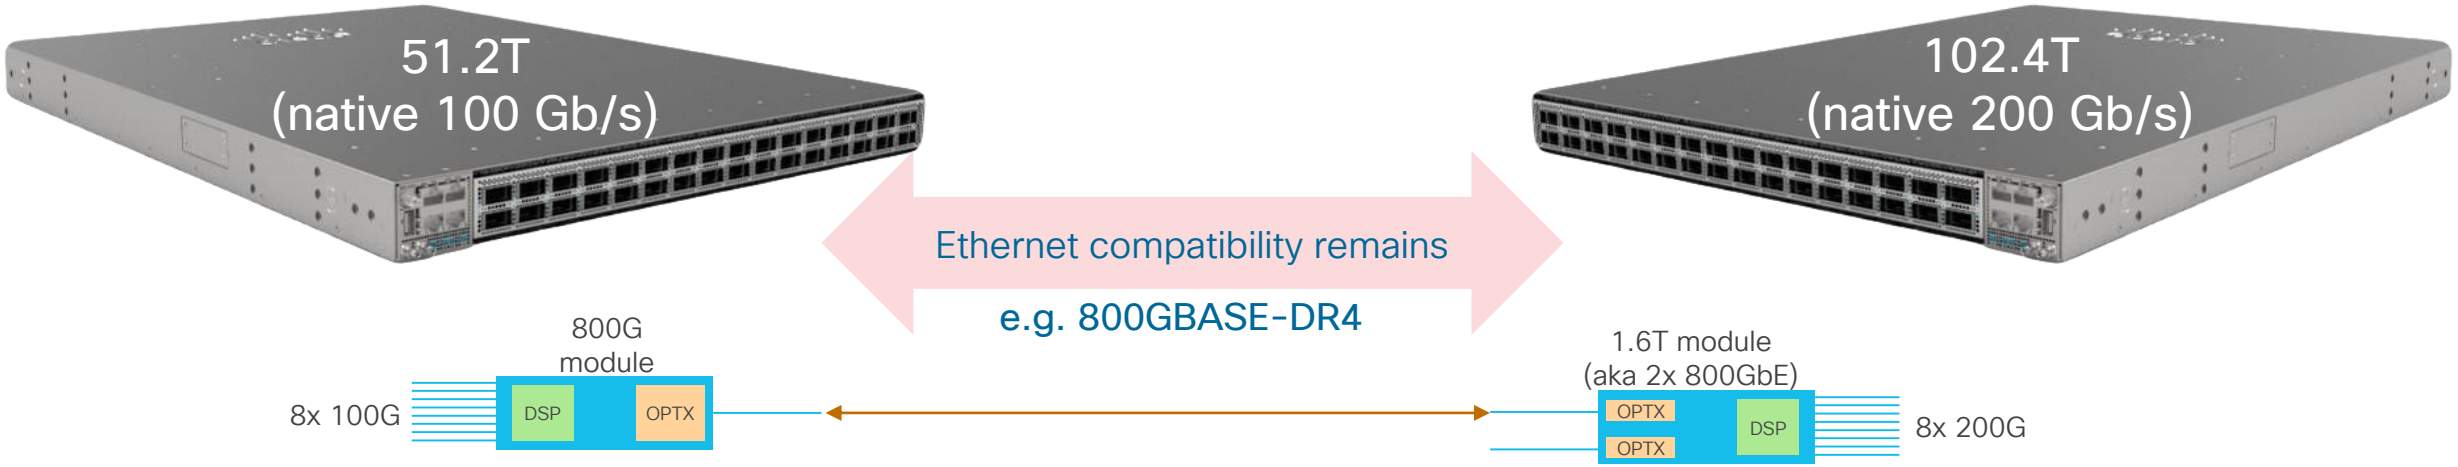

IEEE 802.3 P802.3dj's logic and architecture ensures compatibility with "legacy" systems and implementations.

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

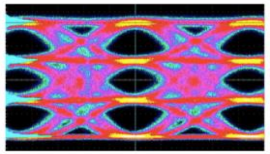

Compatible with current CMOS, SiPhotonics, EML, Cable, Fiber technology

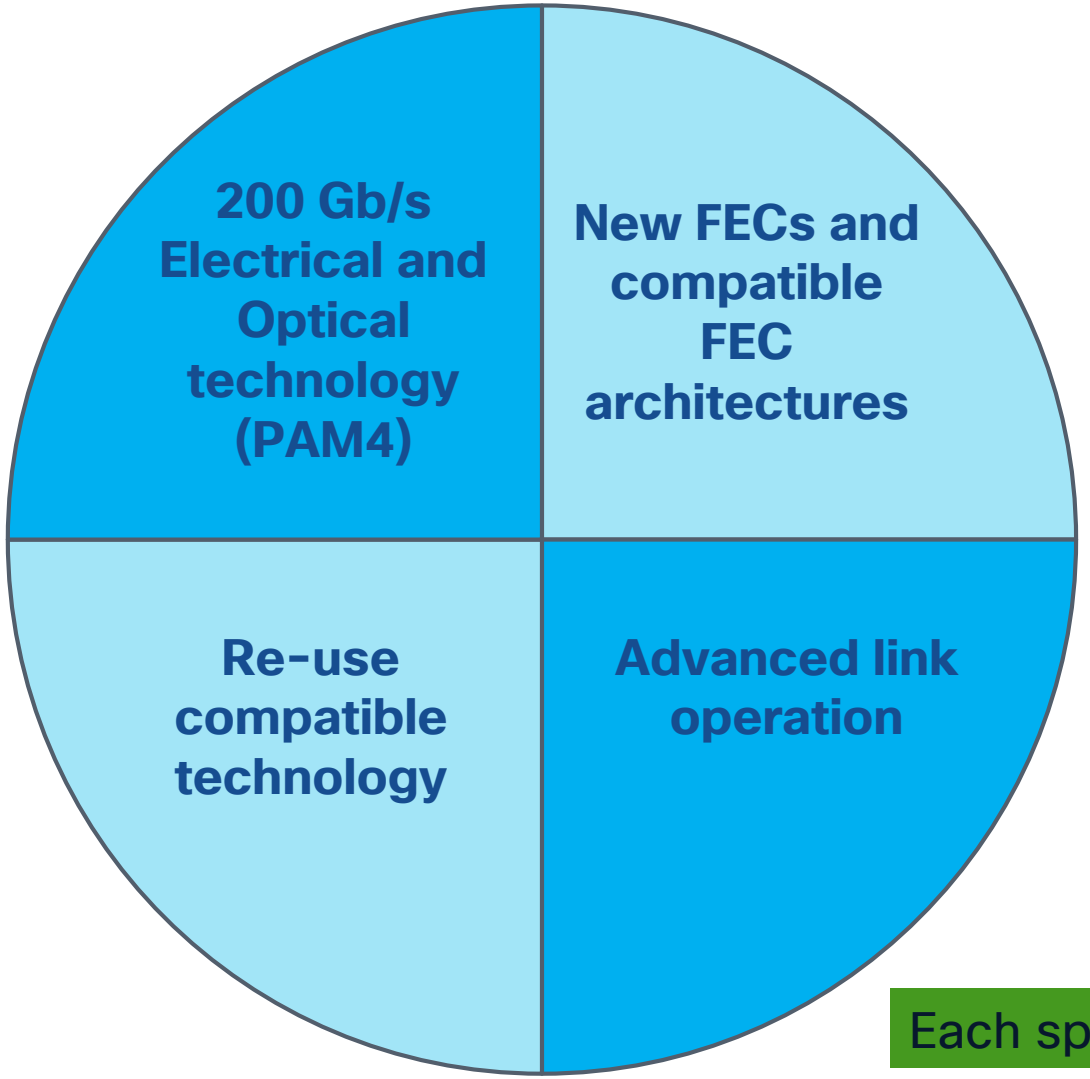

Multiple FEC schemes for different reaches. Common architecture to enable compatible implementations

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

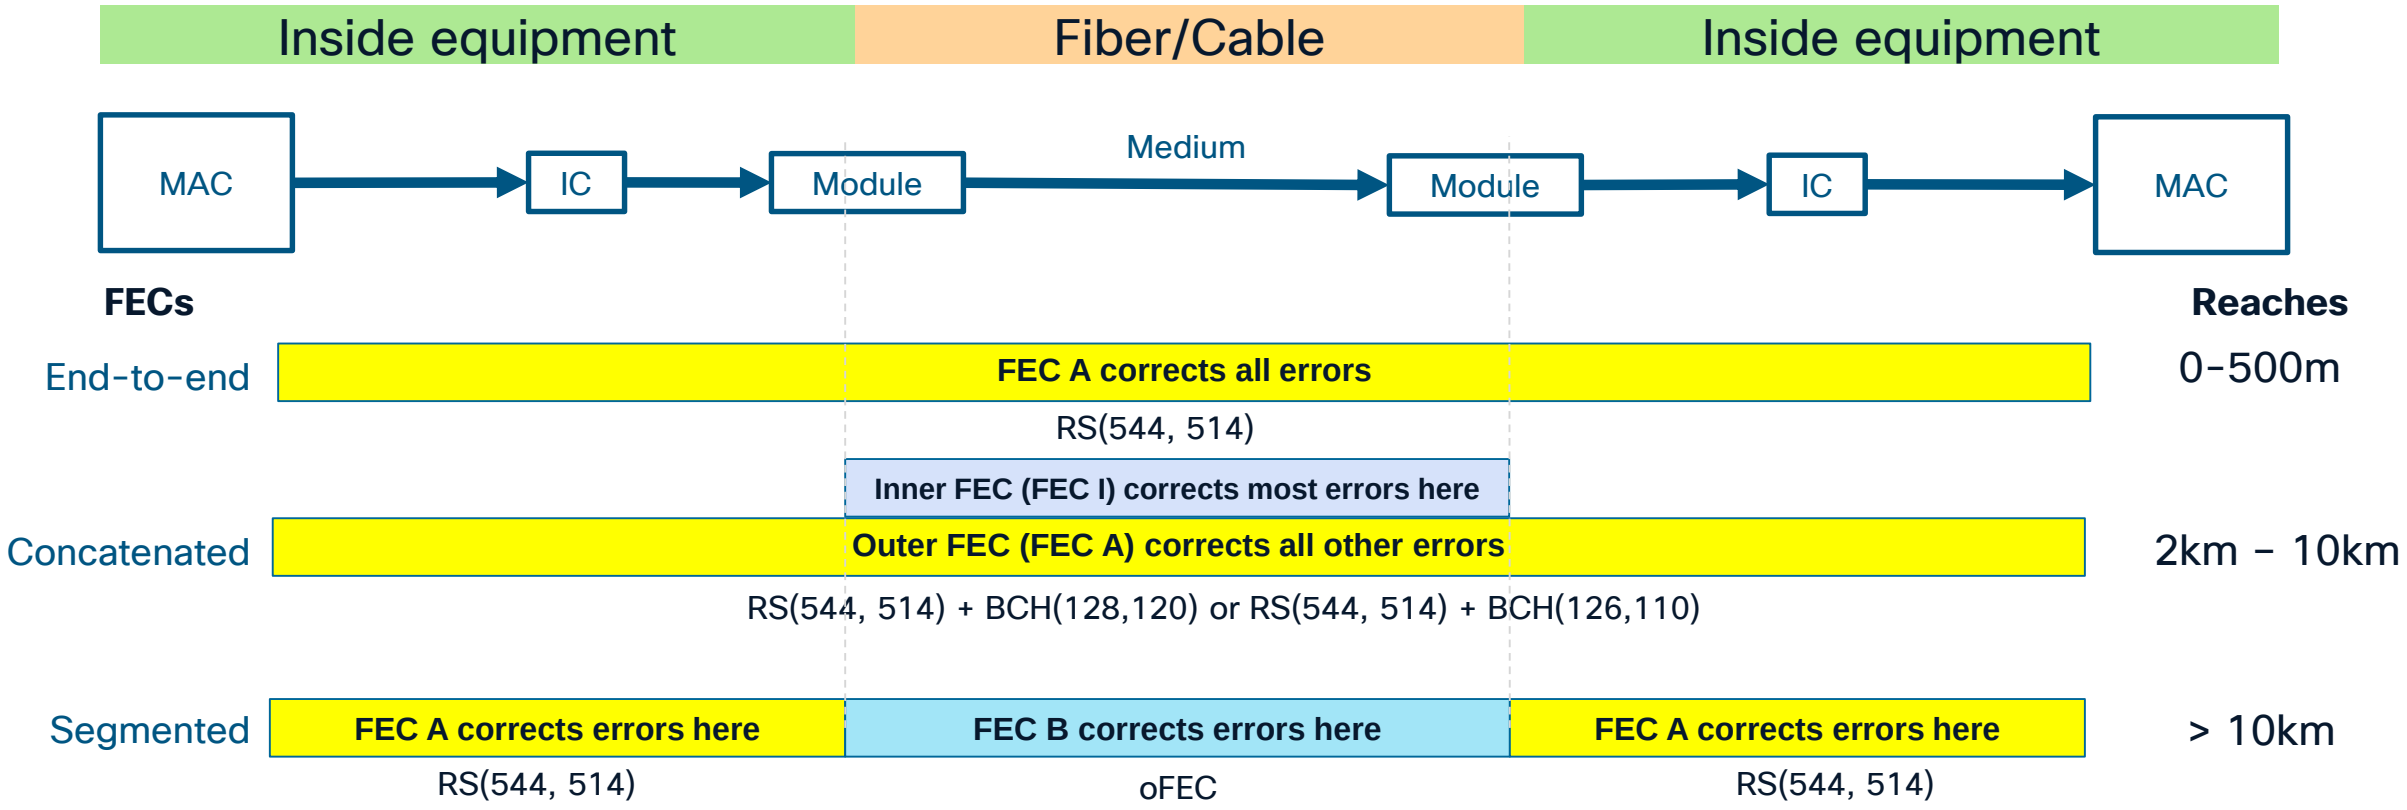

They key point:

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

The final slide: liquid cooling

If rack capacities grow from ~10kW → >100kW something needs to remove the heat. Two solutions under development:

- Cold plate liquid cooling
  - Liquid circulated through pipes and cold plates which are attached to key components (ASIC, optics) to efficiently remove heat
- Immersion liquid cooling
  - Equipment is immersed in special non-conductive liquids to very efficiently remove heat

Beyond 800G: recap

Requirements around AI are driving interconnect technology

- Faster speeds:
  - 100 Gb/s → 200 Gb/s → ...
  - 800 GbE → 1.6TbE → ...
- Lowering power is key
  - Copper cables still live @ 200 Gb/s
  - Novel implementations like LPO / CPO
  - Faster speeds yield better energy efficiency (W/Gb)
- Advanced cooling strategies are being considered
  - Liquid cooling and Immersion cooling
- AI networking will dominate the future optics design priorities for data centers

Wrap up

Key takeaways – thank you for your attention

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

- Flexible
- Scalable
- Performance

Pluggable optics are essential for AI era today. The industry is actively exploring alternative solutions for further optimization for AI's unique demands:

- Co-packaged optics
- Linear pluggable optics
- Silicon Photonics

The future will likely see a mix of these technologies as AI infrastructure continues to evolve.

Cisco continues to lead the optics evolution

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

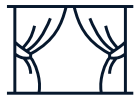

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

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

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

Contact me at: https://ciscolive.ciscoevents.com/clsd25/**BRKOPT-2699**

Thank you

cisco Live !


CISCO