

# Powering High-Performance Data Centers Through Advanced Connectivity

**Christian Ramírez Montilla,**  
**RCDD, DCDC, CCRE**  
TSG Manager LATAM  
**Siemon**



# From Concepts to Connectivity – Powering High-Performance Data Centers Through Advanced Connectivity

Christian Ramírez Montilla, RCDD, DCDC, CCRE

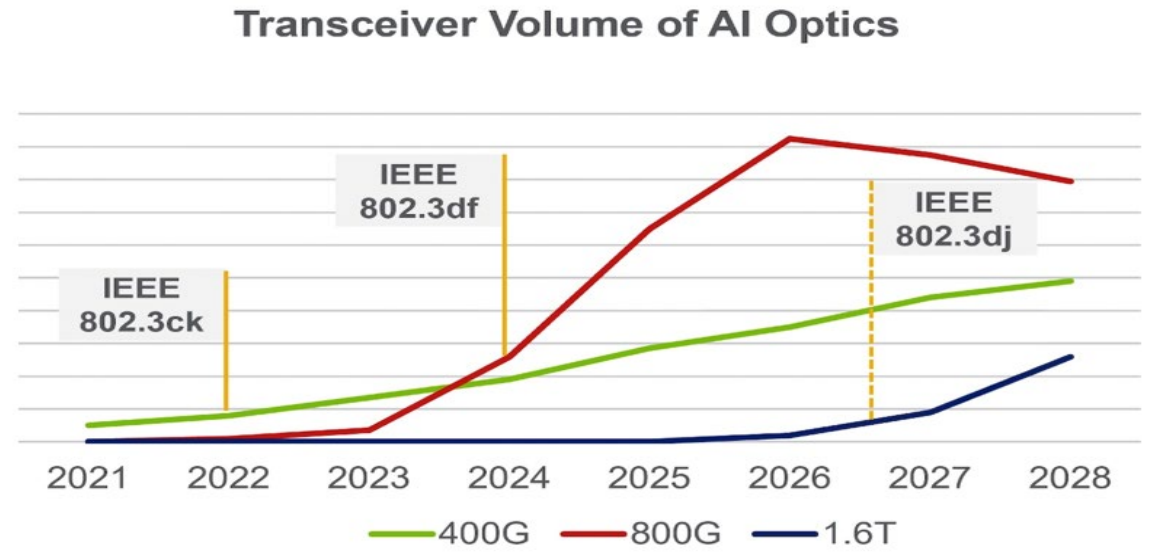
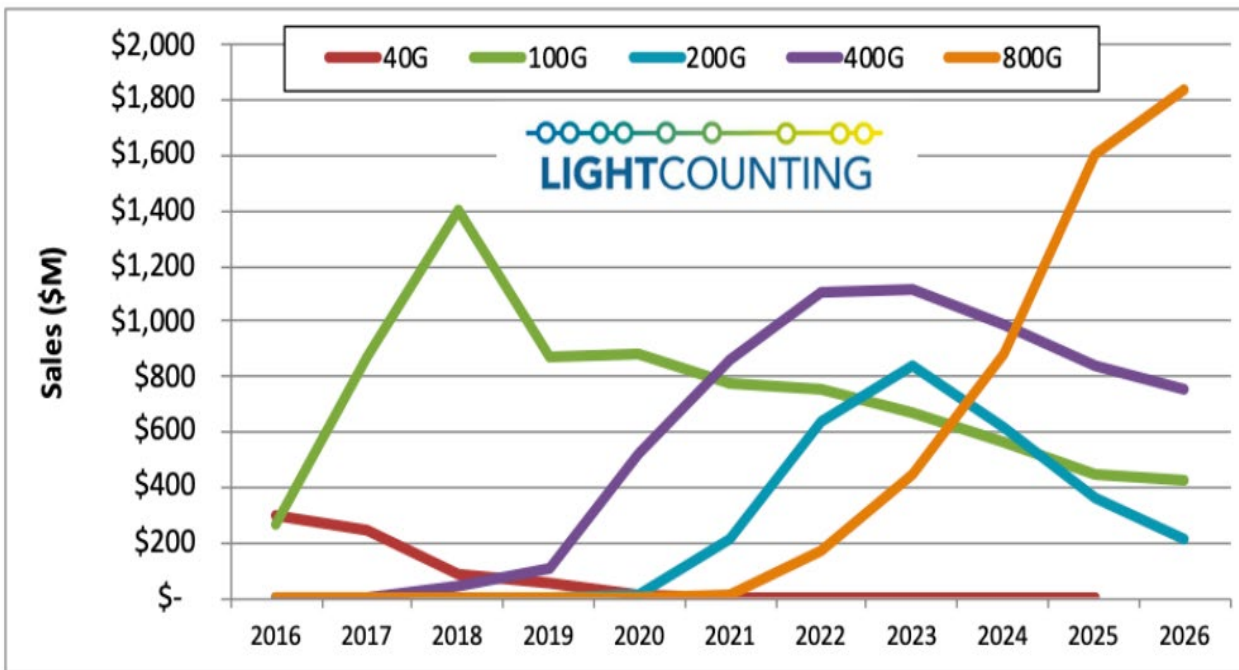
Technical Manager LATAM

[christian\\_ramirez@siemon.com](mailto:christian_ramirez@siemon.com)

Phone: + 57 3185359152



# Transceptores - Capacidades

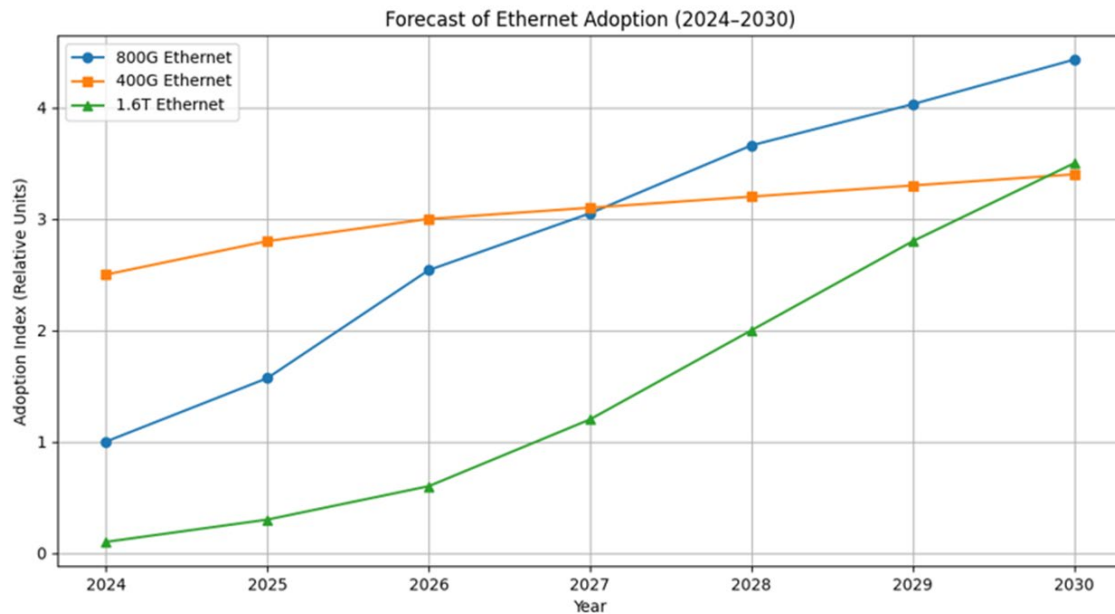


Source: LightCounting

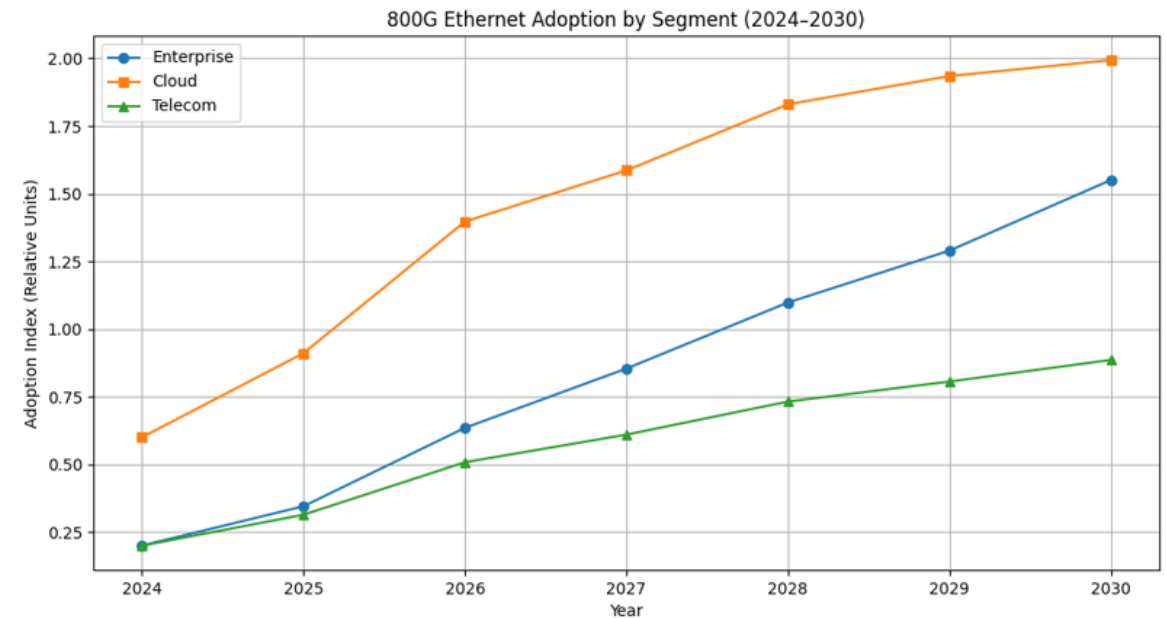
Source: Lightcounting,

# 800GbE Adoption Forecast – Today till 2030

## Ethernet Switch Speed Adoption



## 800GbE Market Segment Adoption



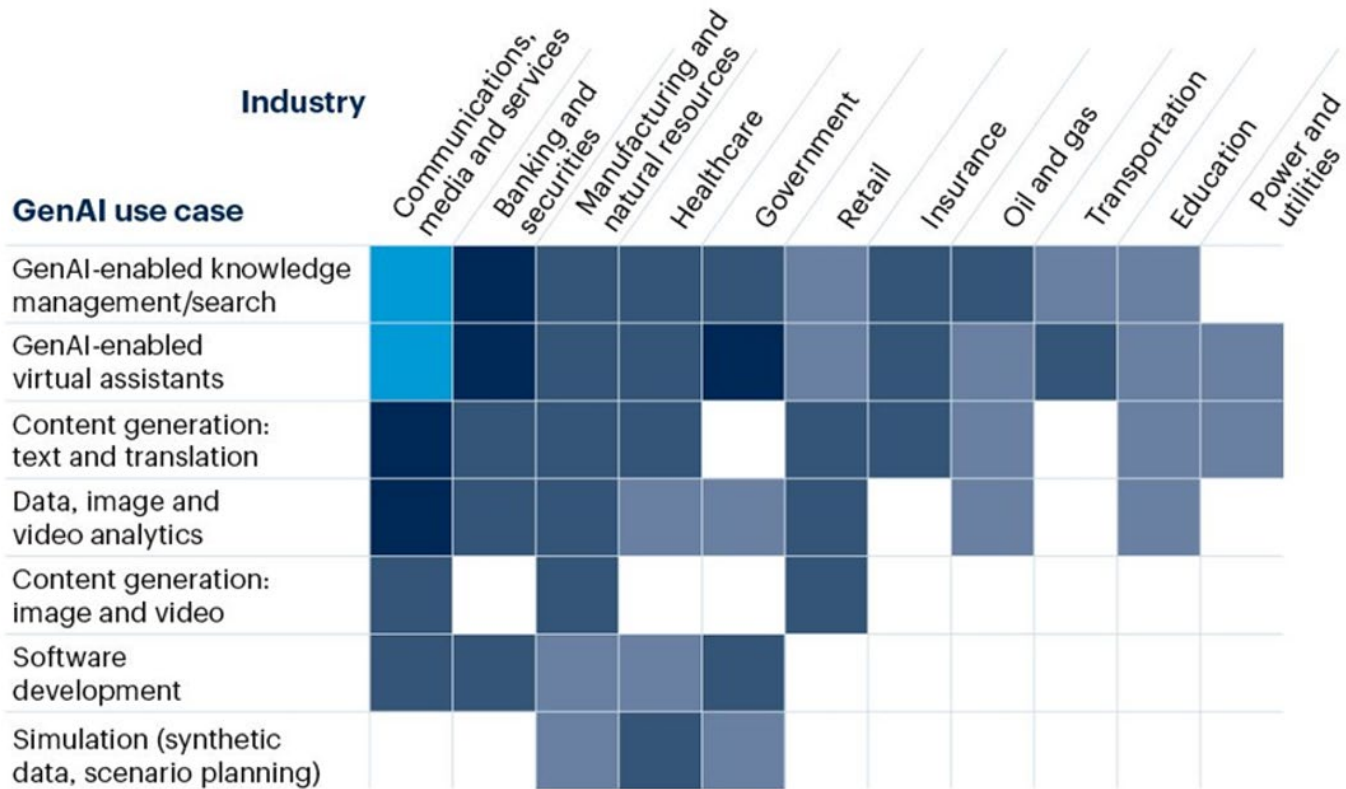
Multiple Sources: LightCounting, Crehan Research, Arista Networks, Keysight Technologies

# IEEE 802.3dj -2026

Ethernet Rate	Signaling Rate	DAC/ACC/AEC	SMF 500m	SMF 2km	SMF 10km	SMF 20km	SMF 40km
200 Gb/s	200 Gb/s	200GBASE-CR1	200GBASE-DR1	200GBASE-FR1			
400 Gb/s	200 Gb/s	400GBASE-CR2	400GBASE-DR2	400GBASE-DR2-2			
800 Gb/s	200 Gb/s	800GBASE-CR4	800GBASE-DR4 800GBASE-FR4-500	800GBASE-DR4-2 800GBASE-FR4	800GBASE-LR4		
800 Gb/s	800 Gb/s				800GBASE-LR	800GBASE-ER1-20	800GBASE-ER1
1.6 Tb/s	200 Gb/s	1.6TBASE-CR8	1.6TBASE-DR8	1.6TBASE-DR8-2			

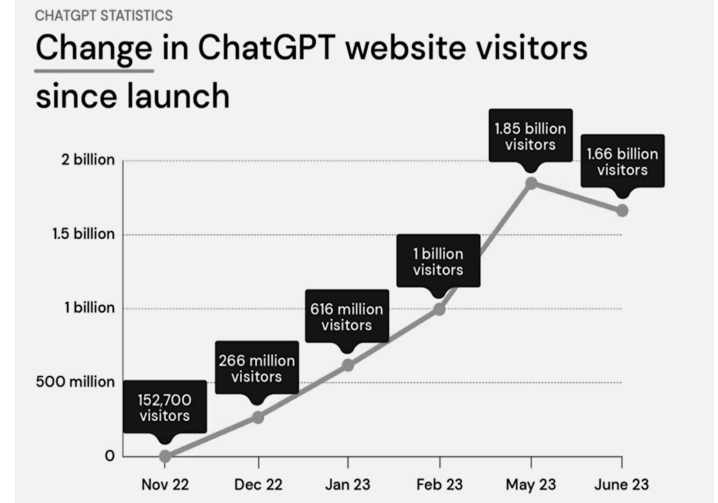
# AI Adoption

Correlation key: ■ Very strong ■ Strong ■ Medium ■ Weak □ No correlation/use case was not discovered



Source: Gartner  
© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. GTS\_3903150

Gartner



# AI – From Concepts to Connectivity

- “Can machines think?” Alan Turing - 1950
- “Success in creating AI would be the biggest event in human history.” Stephen Hawking - 2016
- “AI is one of the most important things humanity is working on. It is more profound than, I dunno, electricity or fire,” Pichai - 2018
- “AI is as revolutionary as mobile phones and the Internet” Bill Gates - 2023
- “It’ll be 10 times bigger than the Industrial Revolution – and maybe 10 times faster.” Demis Hassabis - 2025
- “AI is the new electricity” Andrew NG - 2025
- “AI will reach human levels by around” Ray Kurzweil 2025”



# AI Factory?

“Engines that generate intelligence, not just places where data is stored. They are fundamentally different from traditional data centers.” **Jensen Huang**

Booz Allen Hamilton CTO **Bill Vass** is on board with that particular definition. “An AI factory is a data center that is designed specifically for AI-specific hardware”

Deloitte defines an AI factory as a combination of hardware, software, and services that enable the AI lifecycle. “The primary product is intelligence, measured by token throughput,” said **Nicholas Merizzi**, a principal at Deloitte Consulting and the firm’s AI infrastructure US offering lead..

For others, the AI factory concept is more about the software



# High-performance computing components

In a modern data center, AI, HPC, and accelerated computing are integrated into a single architecture called the "AI Factory."

## xPUs:

Acelerador	OEM / Fabricante principal
GPU	NVIDIA, AMD, Intel [ <a href="https://www.geekflare.com">geekflare.com</a> ]
TPU	Google (Alphabet) – ASIC propietario de Google Cloud [ <a href="https://www.techshark.io">techshark.io</a> ]
NPU	Apple (Neural Engine), Google (Tensor), Samsung (Exynos), Huawei (Kirin), Qualcomm [ <a href="https://www.geekflare.com">geekflare.com</a> ]
IPU	Graphcore [ <a href="https://www.techshark.io">techshark.io</a> ]

## AI Node:

AI server with xPUs = xPUs node

## Node's NIC:


Pluggable modules or Direct attach cables



# High-performance computing Networks

## 4 Separate Network Connections

EFFICIENCY FOCUS



**GPU Network**  
400G/800G Ethernet (or)  
NDR-400G/800G InfiniBand



**Storage Network**  
100G/200G/400G Ethernet

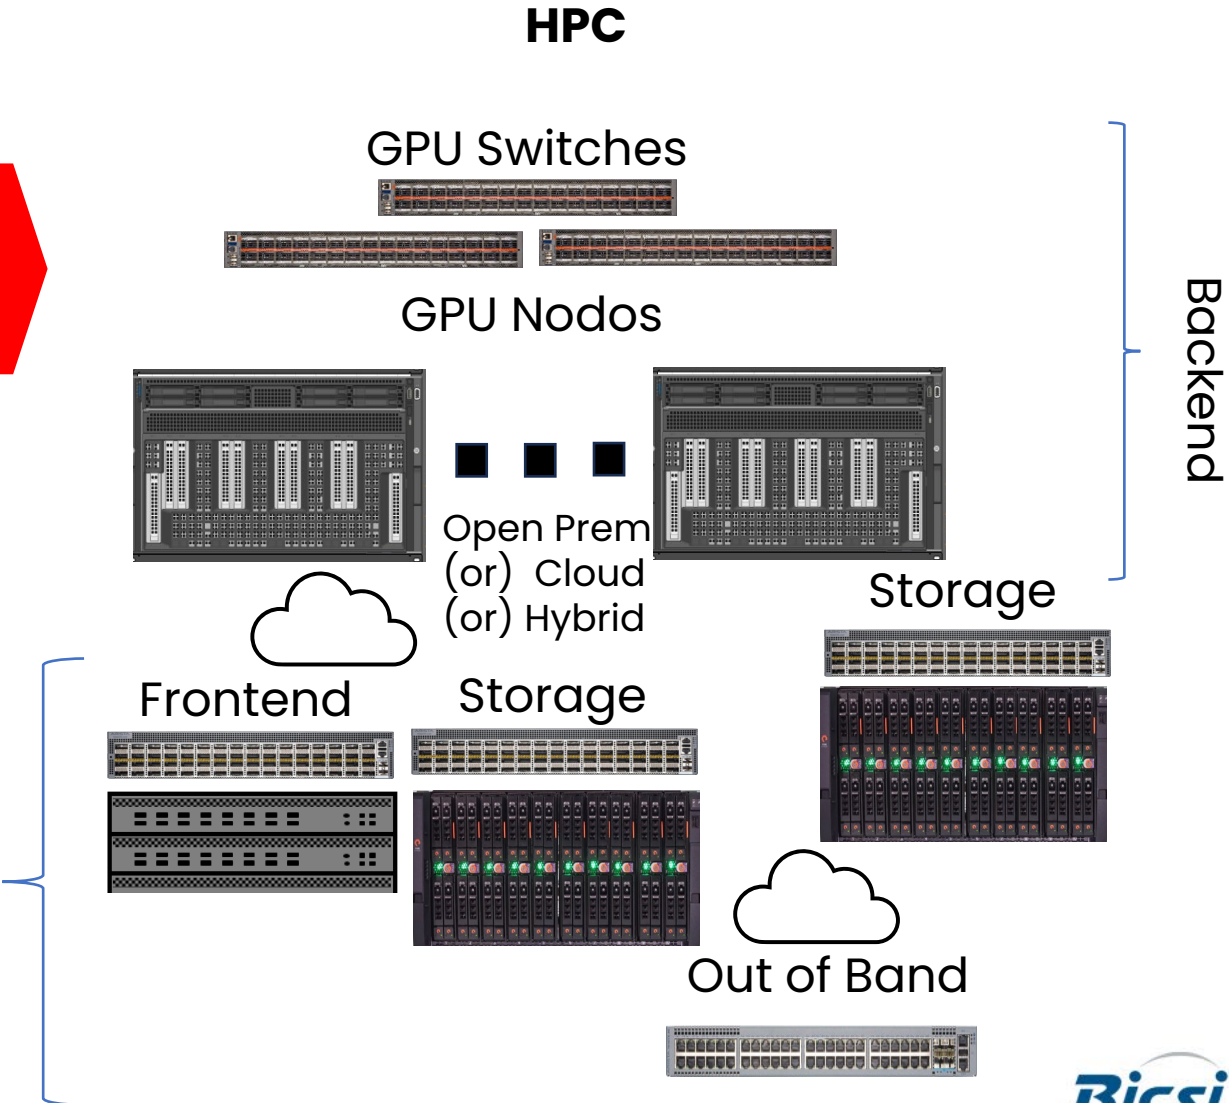
EFFICIENCY FOCUS



**Frontend Network**  
100G/200G/400G Ethernet



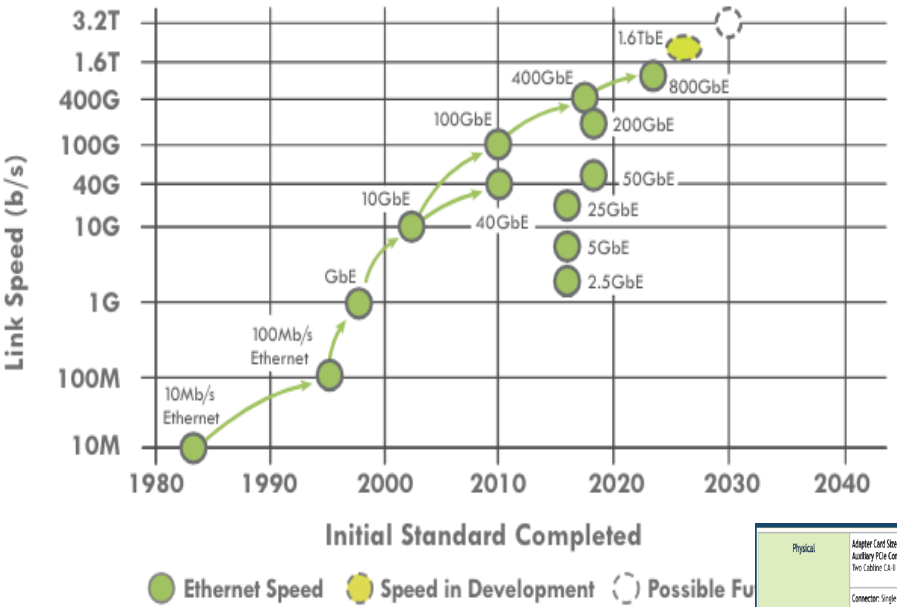
**Out of Band**  
1G/10G Ethernet



# DC Speeds



## ETHERNET SPEEDS



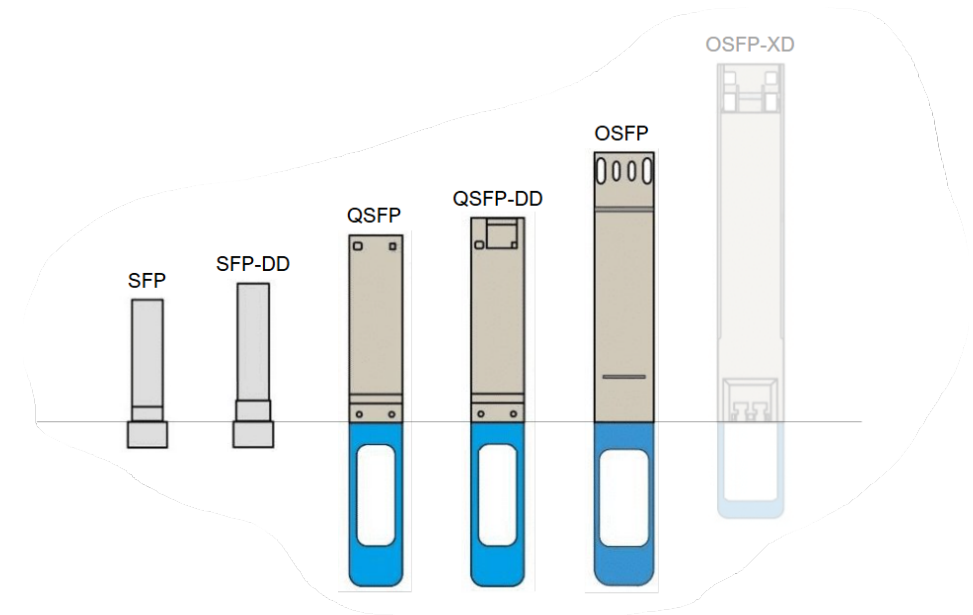
Application Name	Date Rate Name	Lane Speed	Link Speed (4X)
Single Data Rate	SDR	2.5 Gb/s	10 Gb/s
Double Data Rate	DDR	5 Gb/s	20 Gb/s
Quad Data Rate	QDR	10 Gb/s	40 Gb/s
Fourteen Data Rate	FDR	14 Gb/s	56 Gb/s
Enhanced Data Rate	EDR	25 Gb/s	100 Gb/s
Higher Data Rate	HDR	50 Gb/s	200 Gb/s
Next Data Rate	NDR	100 Gb/s	400 Gb/s
Extended Data Rate	XDR	200 Gb/s	800 Gb/s
Global Data Rate	GDR	400 Gb/s	1.6 Tb/s

Physical	Adapter Card Size: 1.6 in. x 2.71 in. (167.65mm x 69.10mm) Auxiliary PCB Connector Card Size: 3.59 in. x 2.22 in. (91.30mm x 56.10mm) Two Cable Clamps (with and without) Length: 35cm		
Connector	Single QDF56 InfiniBand and Ethernet (copper and optical)		
Protocol Support	InfiniBand: IB5 v1.4.0 Auto-Negotiate: 15.12/14.5/13.5/12.5/11.5/10.5/9.5/8.5/7.5/6.5/5.5/4.5/3.5/2.5/1.5/0.5 Gb/s per lane, QDR (10Gb/s per lane), FDR10 (10.24Gbps per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane port), HDR100 (100 Gb/s per lane), HDR (200Gb/s per lane port)		
Ethernet	200GbE-CM, 200GbE-RA, 200GbE-SM, 100GbE-CM, 100GbE-RA, 100GbE-SM, 50GbE-SM, 50GbE-RA, 50GbE-CA, 40GbE-RA, 40GbE-SM, 40GbE-CA, 25GbE-RA, 25GbE-SM, 25GbE-CA, 10GbE-RA, 10GbE-SM, 10GbE-CA, 10GbE-SR		
Data Rate	InfiniBand	SDR/DDR/FDR/HDR/XDR	
	Ethernet	1/10/25/40/50/100/200 Gb/s	
Gen3	SERDES @ 8.0GCT/s, x14 lanes @ 0.8 and 1.1 compatible		



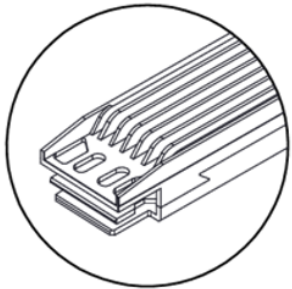
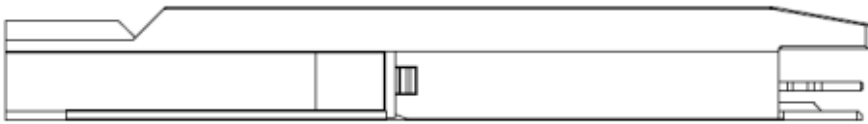
# Some Form Factor – Modules and HSI

- SFP (Small Form Pluggable)
- SFP-DD (Small Form Pluggable Double Density)
- QSFP (Quad Small Form Pluggable)
- QSFP-DD (Quad Small Form Pluggable Double Density)
- <http://www.qsfp-dd.com/>
- OSFP (Octal Small Form Pluggable)
- OSFP-XD Octal Small Form Pluggable eXtra Dense)
- <https://osfpmsa.org/>

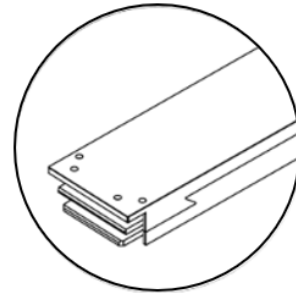


# Not all are the same- HSI

Switch End  
(IHS/Finned Top)




Server End  
(RHS/Flat Top)

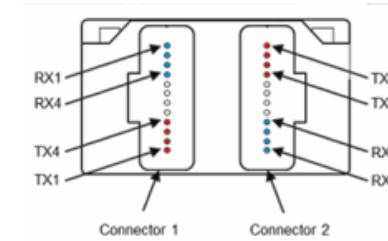


# Not all are the same - Modules

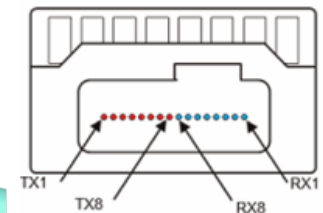

## 800G Dual MPO-12



- Above industry standard performance
- Multimode and Singlemode
- 200G, 400G and 800G
- Ethernet and Infiniband (HDR and NDR)
- Supports dual Base-8 or single Base-16 jumpers




## 800G Single MPO-16



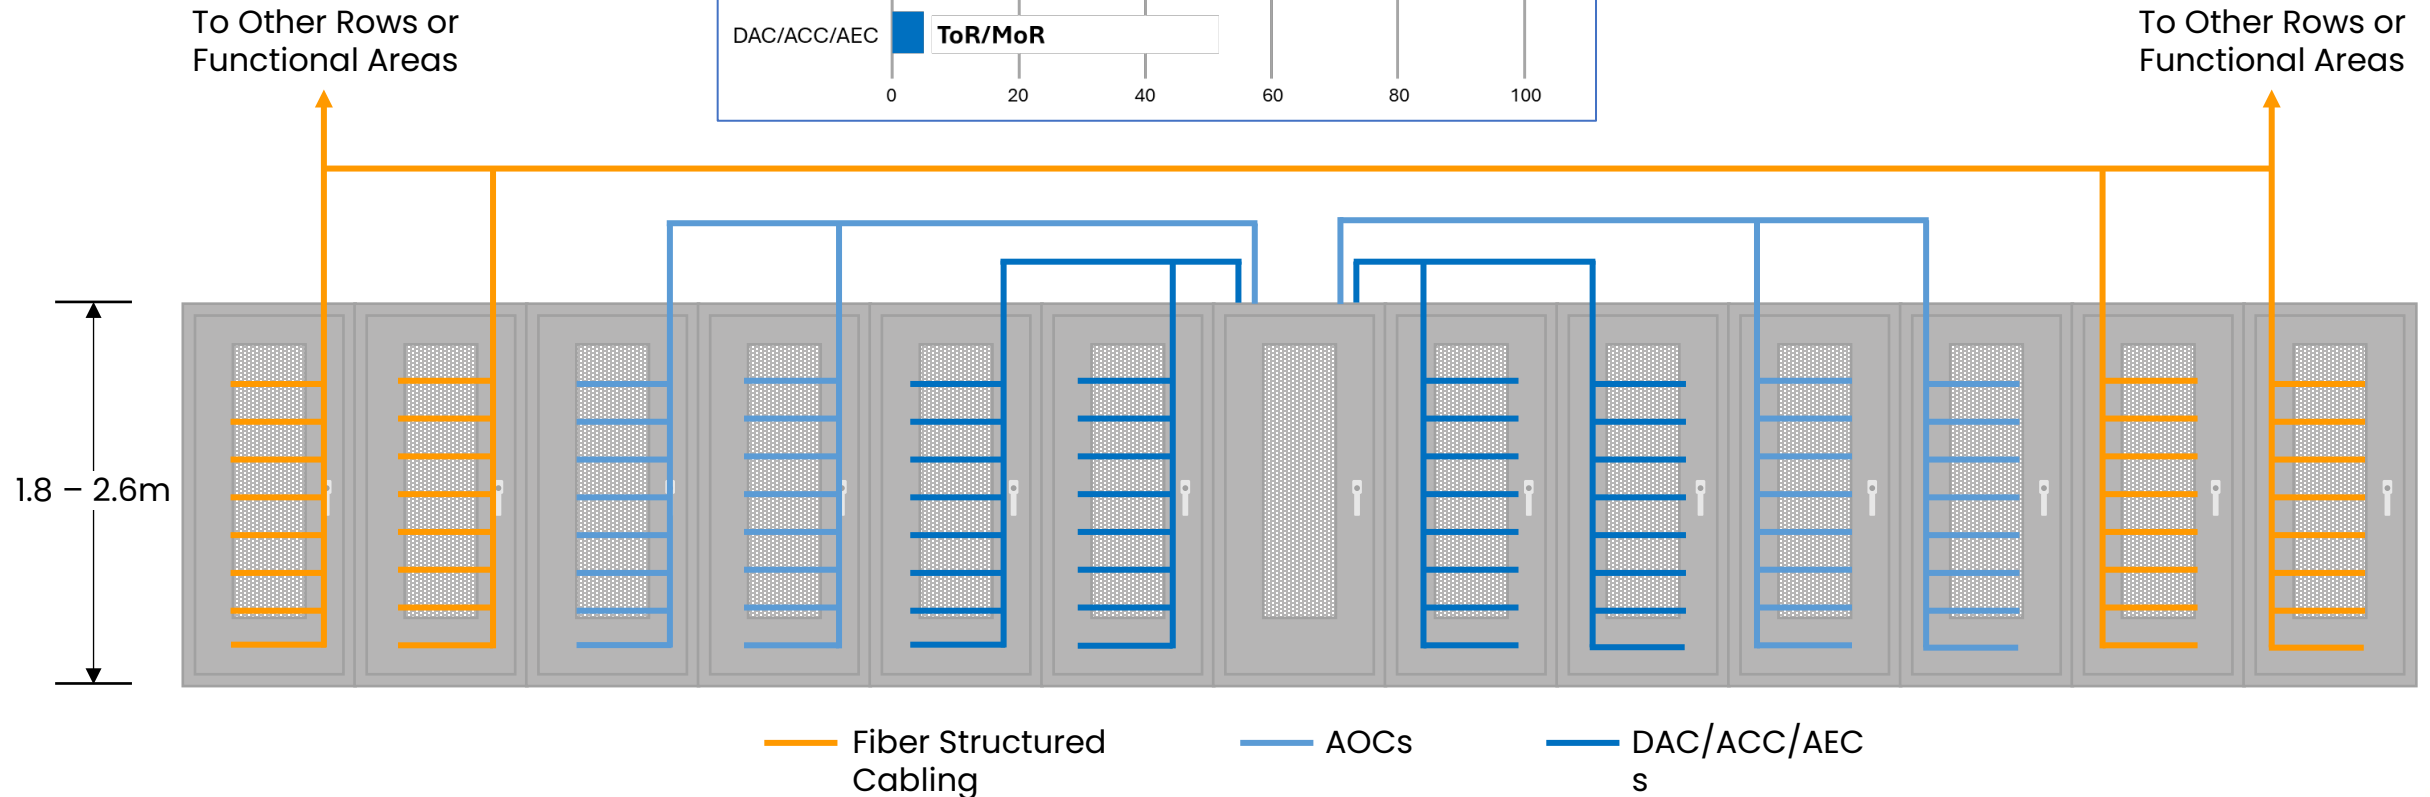
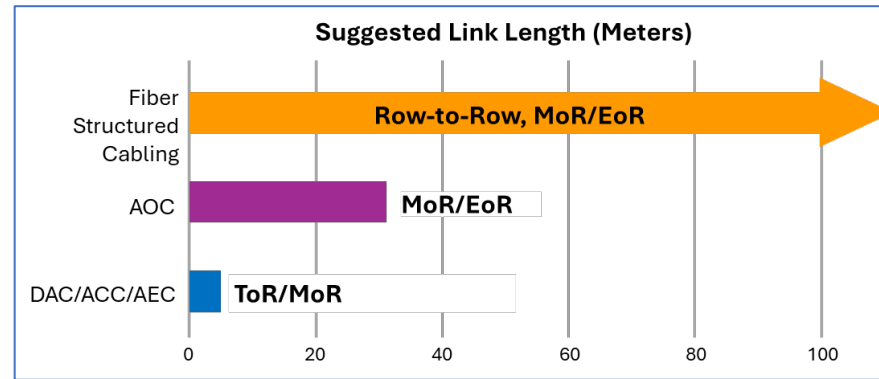
Singlemode  
Or Multimode Options

# High-performance computing components

- Linear Pluggable Optics (LPO), a technology that uses pluggable optical modules without digital signal processing (DSP) chips to reduce power consumption, cost, and latency in data center interconnections.
- Co-packaged optics" (CPO) is an advanced technology in high-speed networking and computing where optical communication components are integrated into the same physical package as the electronic processing chip



# Direct Attached Cabling Vs Structured Cabling



# Direct Attached/Point to Point Cabling Vs Structured Cabling

Direct attach cabling: Cabling that provides a connection between equipment without any intervening connections.

The use of direct attach cabling as an **alternative** to structured cabling.

The reference architecture from NVIDIA and other key players has changed everything.

## DACs, ACCs, AOCs, and Transceiver

[Is this page helpful?](#)

### Interconnects

There are two main ways to link switches and adapters by using either copper wires or optics. Copper has a length or reach limitation of less than 5 meters and two different optical technologies enable using different technologies for the least cost to fit the application. Multimode short reach optics typically have 50m-100m maximum reaches. Single-mode, mid to long reach optics are typically 100m, 500m, 2km, 10km, and 40km maximum reaches.

There are many different technology combinations of optical connector, plugs, optical connectors, electronics, and optics. This document concentrates on high-volume products offered by specifically NVIDIA for accelerated AI data center-oriented cables and transceivers.

Definition	Photo
DAC (Direct Attach Copper) cable with QSFP connector	
DAC with SFP connector	
AOC (Active Optical Cable) with QSFP connector	
QSA (QSFP to SFP Adapter)	
QSFP transceiver QSFP28 Transceiver for 100G transmission QSFP56 Transceiver for 200G transmission QSFP112 Transceiver for 400G transmission	
QSFP-DD transceiver 8 lane 200/400G transceiver	
OSFP transceiver Single/Dual 8 lane 1/2 x 400G transceiver	
SFP transceivers 25G SFP28 Transceiver (~1 W)	

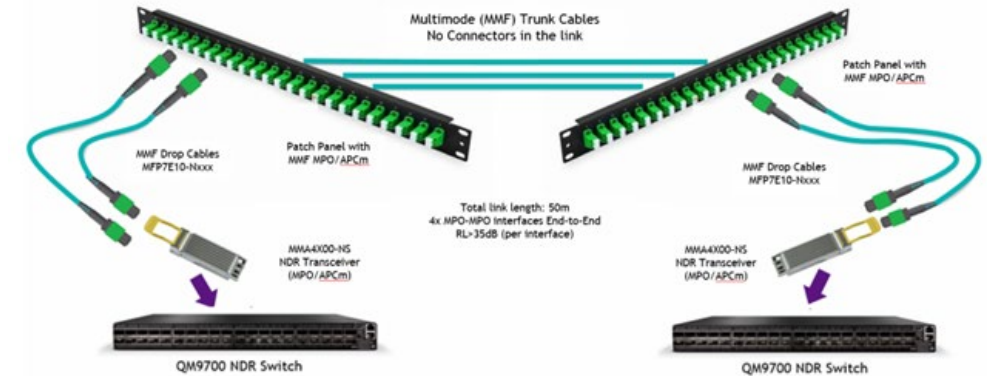
NVidia

# Structured Cabling

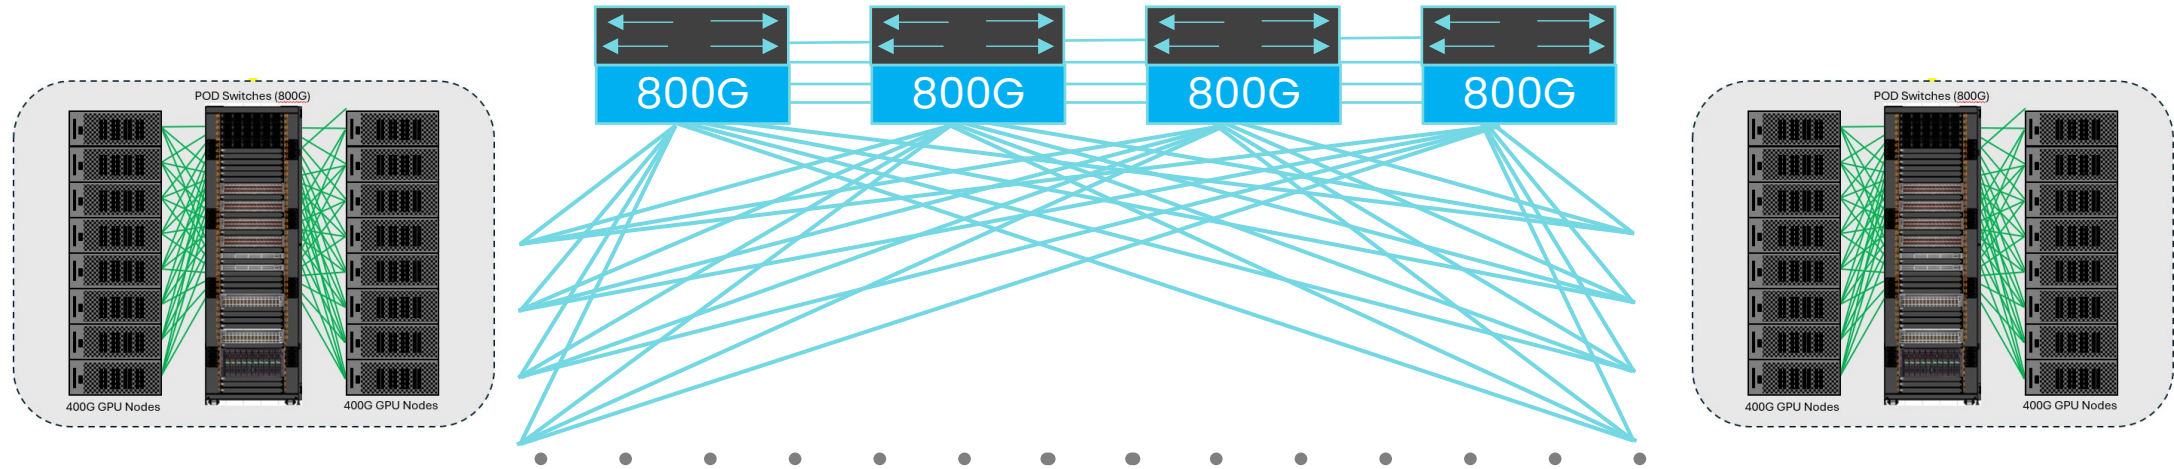
“Structured Cabling eases servicing of the optical transceivers in the data center as transceivers are detachable from the cable. This also allows easier installation since there are no ESD sensitive or bulky transceivers involved in the cabling.”

“Patch panels and trunk cables add further flexibility to the infrastructure.”

“The maximum reach specifications of NVIDIA IB NDR/XDR and 400/800G ETH transceivers assume two optical connector junctions or patch panels in the link.”



# Efficient Cabling Scale-Out – Leaf Switch to Spine Switch



## Structured Fiber Cabling

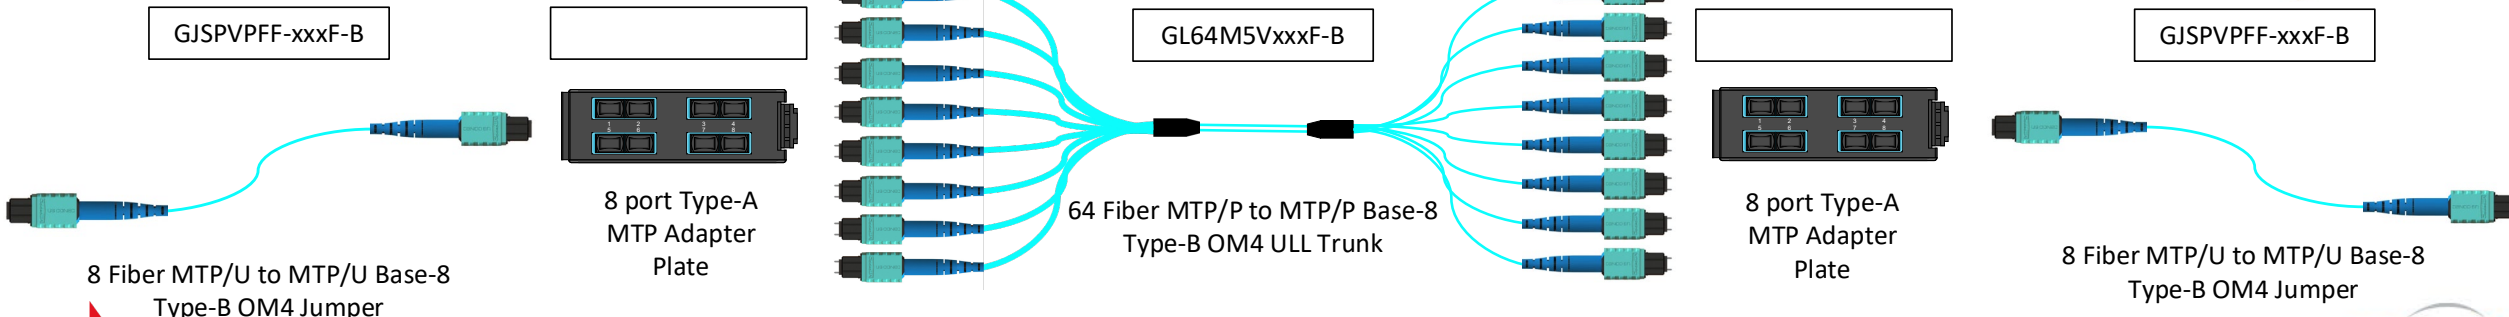

Best for across rows

## Bundled MTP APC Trunks

Single Mode Fiber (or) Multimode

## 800G LPO Optics

DR (or) SR



# What type of cable?

- MMF (Multi-Mode Fiber) – The type of fiber used for VCSEL (Vertical Cavity Surface Emitting Laser) based transmission, normally operating at 850 nm wavelength.
- OM2 – orange – used for data rates at 1-14 Gb/s, 62.5  $\mu\text{m}$  fiber core diameter
- OM3 – aqua – 70 m reach for 25/100 Gb/s transceivers, 50  $\mu\text{m}$  core diameter
- OM4 – aqua – 100 m reach for 25/100 Gb/s transceivers, 50  $\mu\text{m}$  core diameter
- OM5 – aqua green – not commonly used (2023)
  - <https://docs.nvidia.com/>
- SMF (Single-Mode Fiber) – The type of fiber operating at 1310 or 1550 nm wavelength. Single-mode fiber usually has a yellow jacket and can reach 100s of km. The tiny 7-9  $\mu\text{m}$  light carrying core makes building single-mode optics much more expensive than multi-mode optics.

Flat Top Transceiver



#### ✓ Note

Images are for illustration purposes only. Product labels, colors, and lengths may vary.

#### Key Features

- IB and ETH support
- 400G SR4 multimode
- 4-channels of 100G-PAM4 electrical and optical modulation
- Flat top OSFP connector shell
- 850nm VCSEL
- Maximum reach:
  - 30m using OM3 fiber
  - 50m using OM4 fiber
- Single MPO-12/APC optical connector
- Operates as a 200Gb/s NDR200 transceiver with 2-fiber splitter ends
- 8.5 Watts (max) using 4-channels
- 5.5 Watts (max) using 2-channels
- Single 3.3V power supply
- Class 1 laser safety
- Hot pluggable, RoHS compliant
- CMIS 4.0 compliant
- Case temperature range 0°C to +70°C

# NVIDIA Interconnect Cabling Requirements

- Requires ultra-low loss cabling.

[https://files.siemon.com/en/specsheet/siemon-base-8-plug-and-play-cable-assemblies\\_spec-sheet.pdf](https://files.siemon.com/en/specsheet/siemon-base-8-plug-and-play-cable-assemblies_spec-sheet.pdf)



Parameter		MPO MM Patch/Trunk	MPO SM Patch/Trunk	LC SM Patch/Trunk
Connect or	Type and Grade	MPO low loss or MTP Elite, Green Tab	MPO low loss or MTP Elite, Green Tab	LC low loass grade B per IEC 61753-1, blue housing.
	IL (Max)	0.35dB <b>0.30dB</b>	0.35dB <b>0.30 dB</b>	0.25dB <b>0.20dB</b>
	RL	-35dB <b>35dB</b>	-60dB <b>65 dB</b>	-55dB <b>60dB</b>
	Gender	Female for patch, Male for Trunk	Female for patch, Male for Trunk	N/A
	End face polish	APC	APC	UPC
	Spec compliance	IEC 61754-7 and EIA/TIA-604-5 (FOCIS-5)	IEC 61754-7 and EIA/TIA-604-5 (FOCIS-5)	IEC 61754-20, TIA 604-10-A (LC)
Fiber	Grade	OM3: L≤30m OM4: 30m<L≤50m	OS2	OS2
	Attenuation	3 dB/km	0.4 dB/km	0.4dB/Km
	Spec compliance	ITU-T G.651.1	ITU-T G.652D and G.657.A1	ITU-T G.652D and G.657.A1

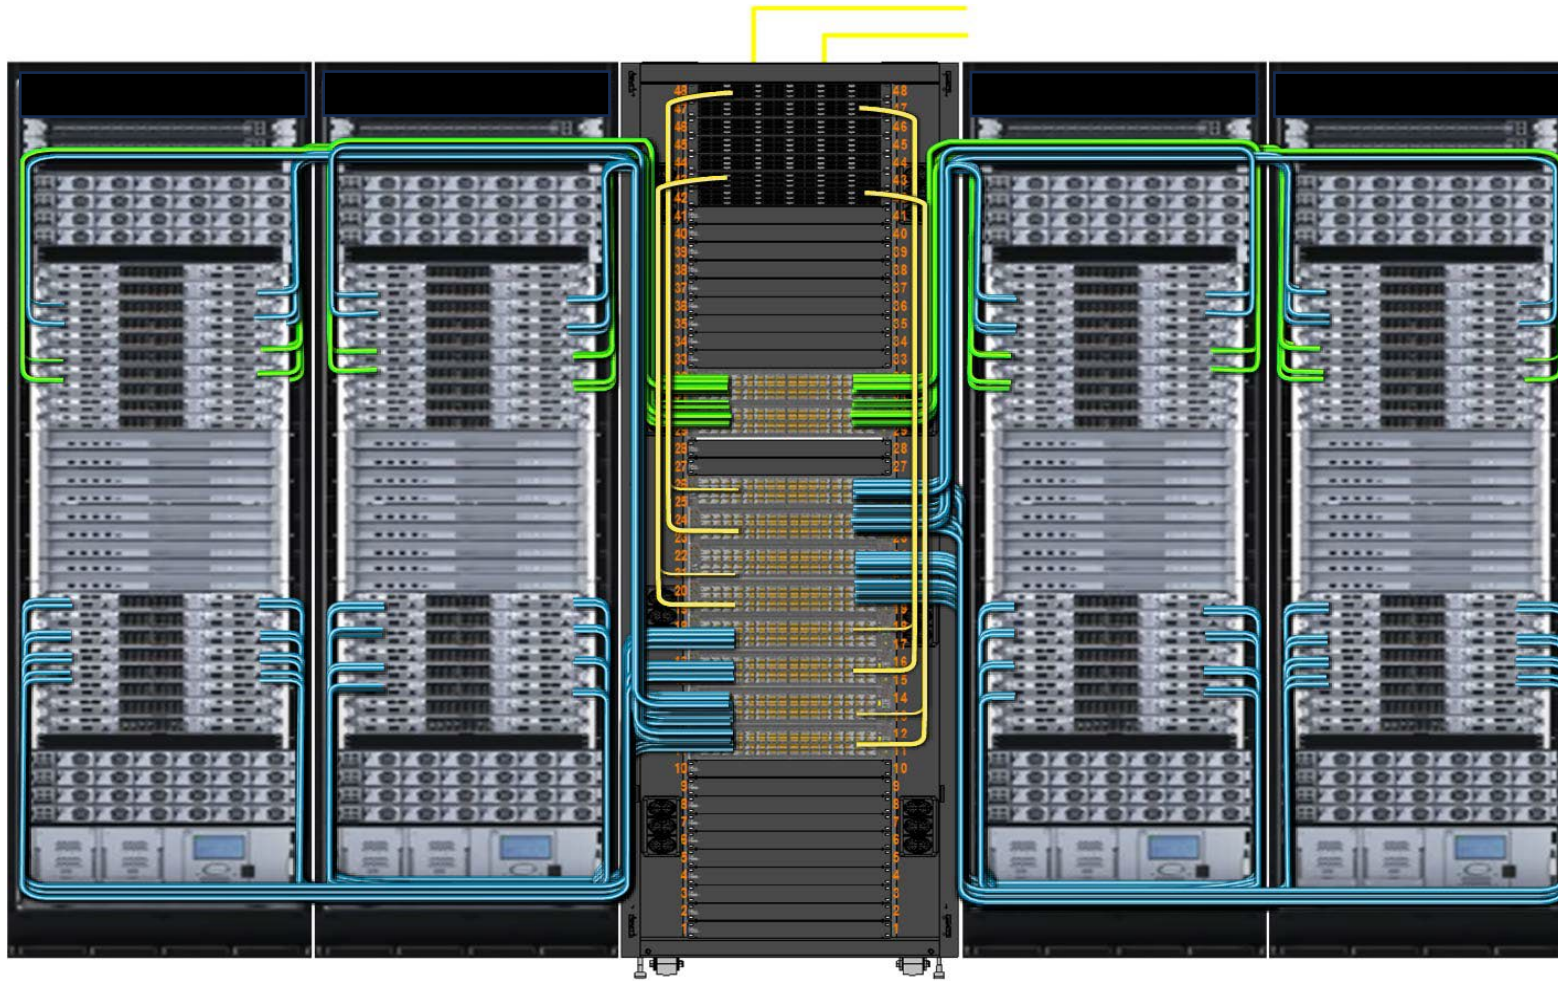
Source: NVIDIA Interconnect Cabling Requirements for InfiniBand NDR/XDR and 400/800G ETH Solutions, June 2025

# NVIDIA Interconnect Cabling Requirements

- NVIDIA specifies the maximum channel loss limit for NDR/XDR and 400/800G.
- New maximum channel limits for Siemon Ultra-Low-Loss provide headroom.

Parameter		MPO MM Patch/Trunk	MPO SM Patch/Trunk	LC SM Patch/Trunk
NVIDIA end-to-end requirements	Max Interfaces (connectors) throughout an optical link	4 <b>2 connections</b>	4 <b>2 connections</b>	4 <b>2 connections</b>
	MPO/MTP adapter type (patch panel)	TYPE A (KEY UP – KEY DOWN) Complies to GR-1435	TYPE A (KEY UP – KEY DOWN) Complies to GR-1435	A-B adapter, per TIA-568
	Max Length	50m	500m	2Km
	Max IL	2dB	2dB	2dB

# Direct Attached Cabling - HSI



# Cabling Considerations – Cable Flexibility and Cable Management - HSI

**FAST & FLEXIBLE CABLES**

When Your Rack Layout Throws you a Curve Ball...

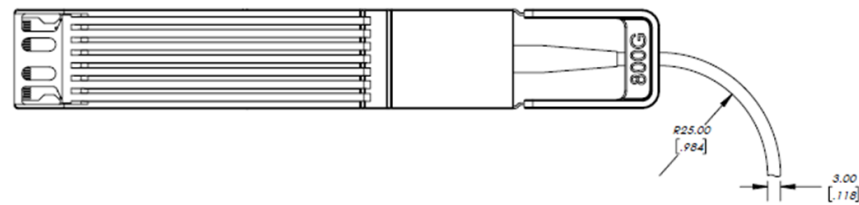


Get speeds up to 100G and a curve that will find the sweet spot of any server deployment.

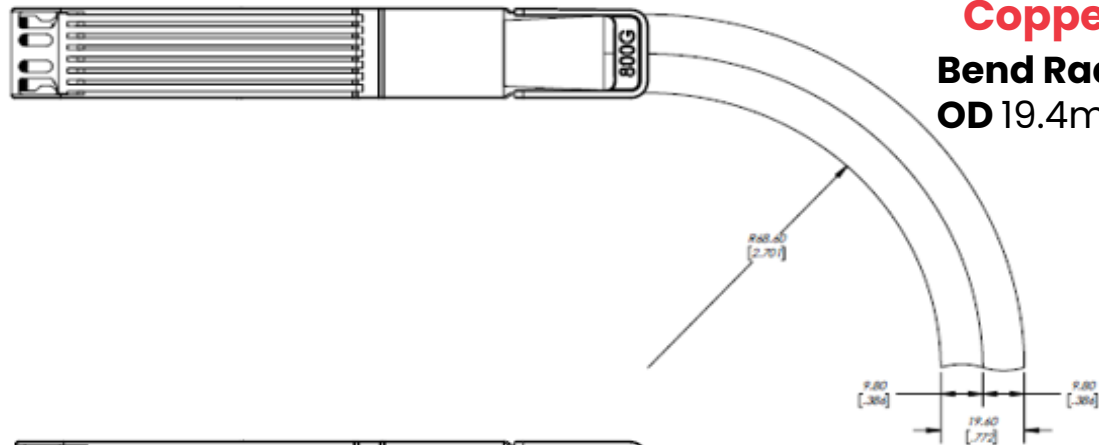
www.siemon.com



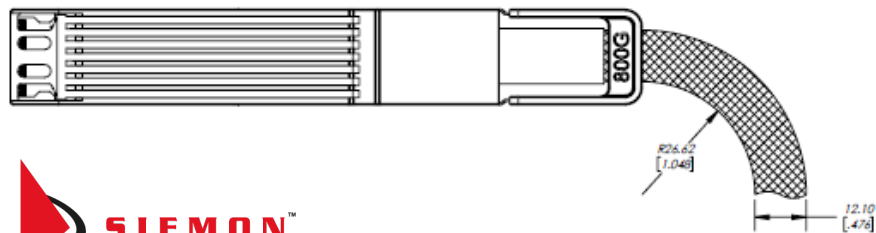
## Cable Assembly – Bend Radius



**Fiber/AOC - OM3**  
**Bend Radius** 25mm (.98 in.)  
**OD** 3mm (.118 in.)




**Copper - Dual Cable**  
**Bend Radius** 72mm (2.8 in.)  
**OD** 19.4mm (.772 in.)



**Copper - Mesh**  
**Bend Radius** 27mm (1 in.)  
**OD** 12.1mm (.476 in.)

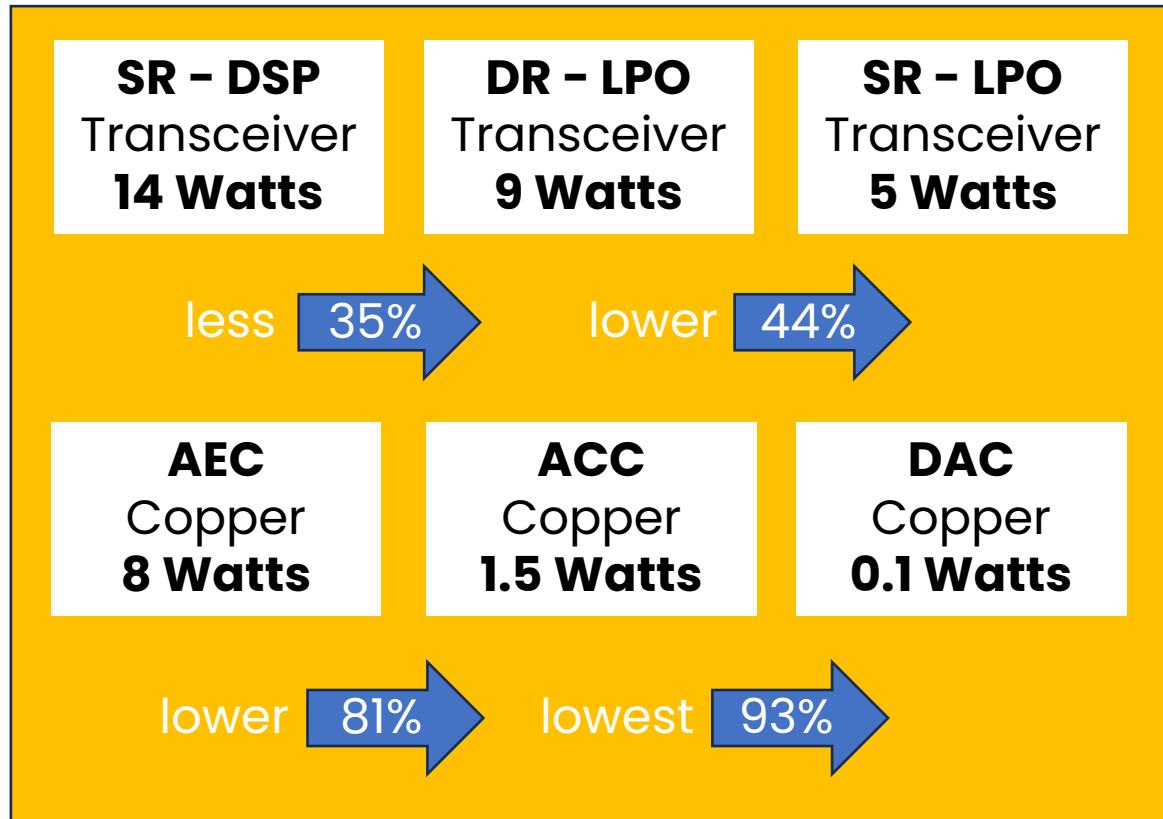
## Every Bend Counts



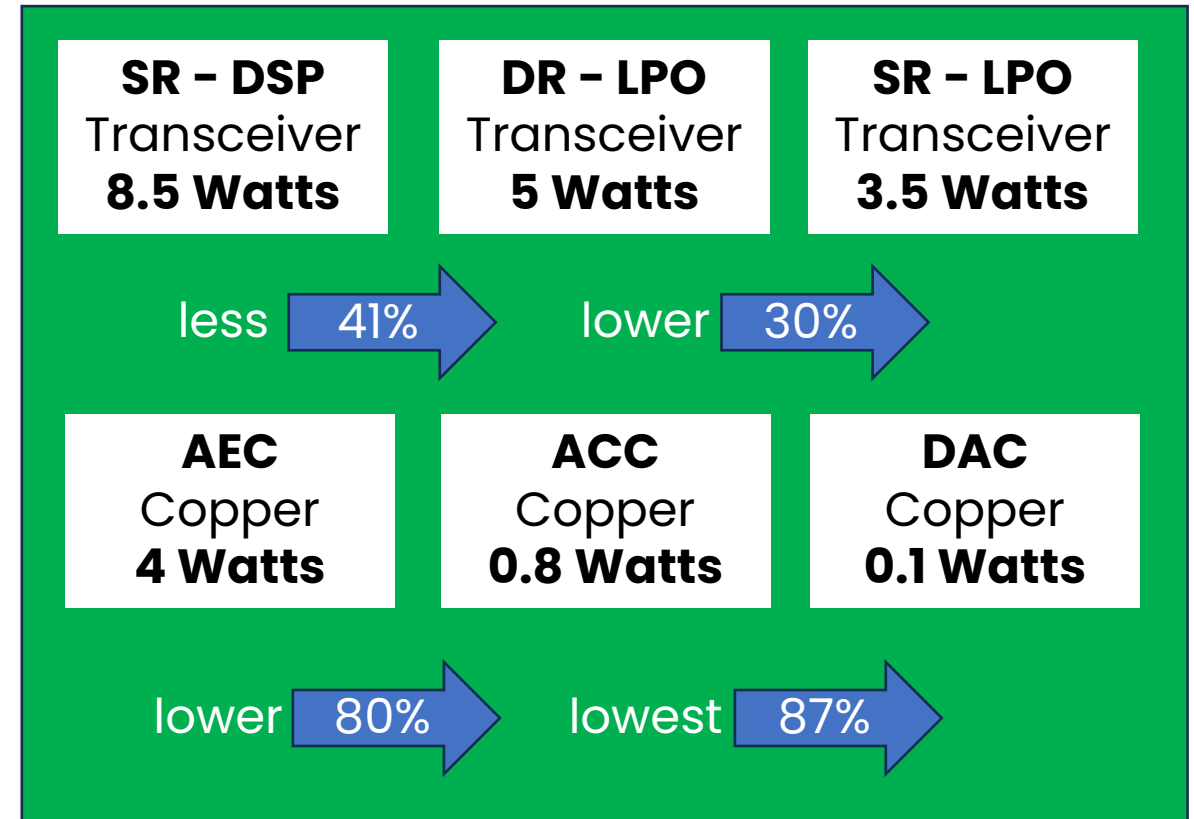
Lots of bending points for routing

Bending cables outside of specification can cause flapping and performance issues

# Cabling Considerations – Power and Heat per Switch Port - HSI



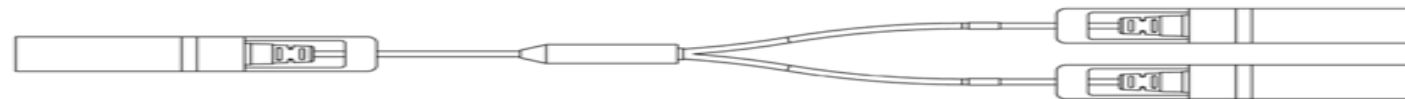
**800G Power Comparison**



**400G Power Comparison**



\*Reducing POWER also reduces HEAT in your system



# Efficient Cabling Comparison – 1024 GPU Scale-Out

## Typical – Reference Design

- DR DSP Optics – Leaf Switch to Aggregation/Spine
- DR DSP Optics – GPU to Leaf Switch

	Optic Type	Power Each	1024	
Leaf Switch to Aggregation/Spine	800G DR8 DSP	17 watts	34816	watts
GPU to Leaf Switch	800G DR8 DSP	17 watts	8704	watts
	2x400G DR4 DSP	9 watts	9216	watts
<b>Total</b>			<b>52736</b>	<b>watts</b>

## Efficient

- SR LPO Optics – Leaf Switch to Aggregation/Spine
- Copper – GPU to Leaf Switch

	Optic Type	Power Each	1024	
Leaf Switch to Aggregation/Spine	800G SR8 LPO	5 watts	10240	watts
GPU to Leaf Switch	800G AEC end	8 watts	4096	watts
	2x400G AEC end	4 watts	4096	watts
	800G DAC end	0.1 watts	51.2	watts
	2x400G DAC end	0.1 watts	102.4	watts
<b>Total</b>			<b>18586</b>	<b>watts</b>

Lower Power

64%

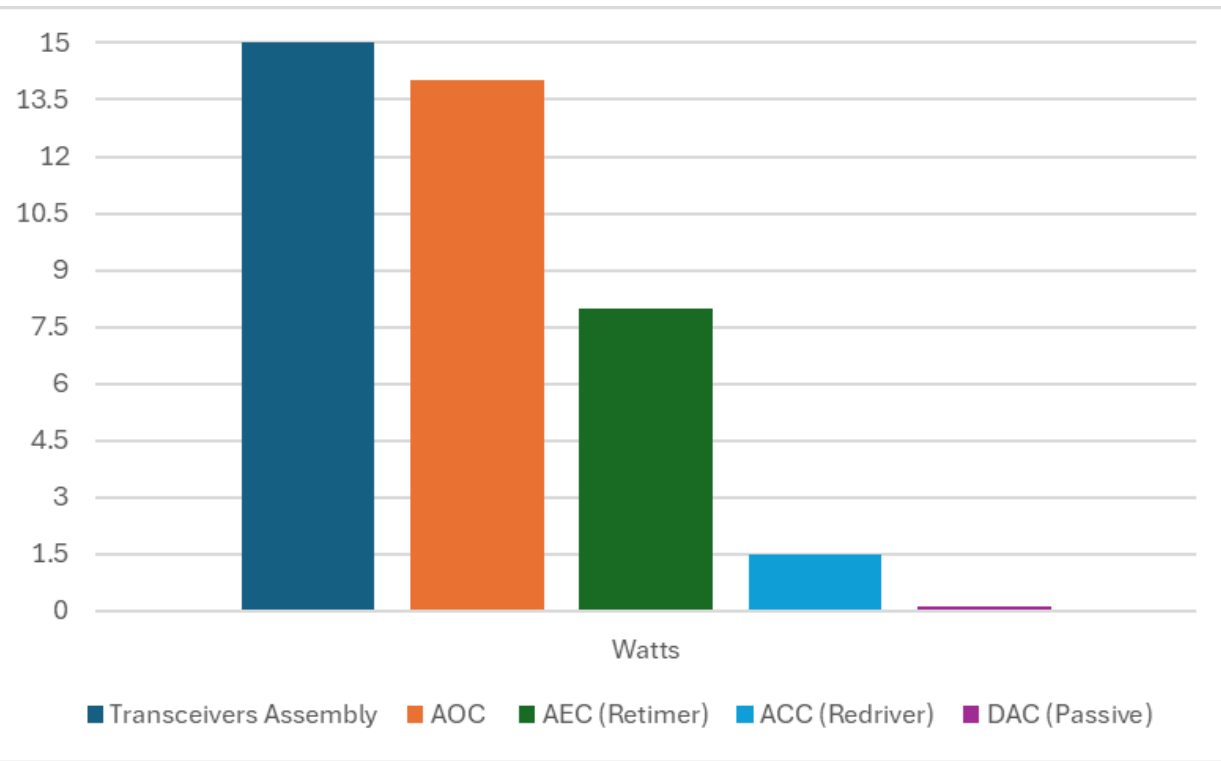
Adds Up Quickly

- ✓ Lowest Power = +60% Savings
- ✓ Lowest Price = +40% Savings
- ✓ Lowest Latency = Best Performance

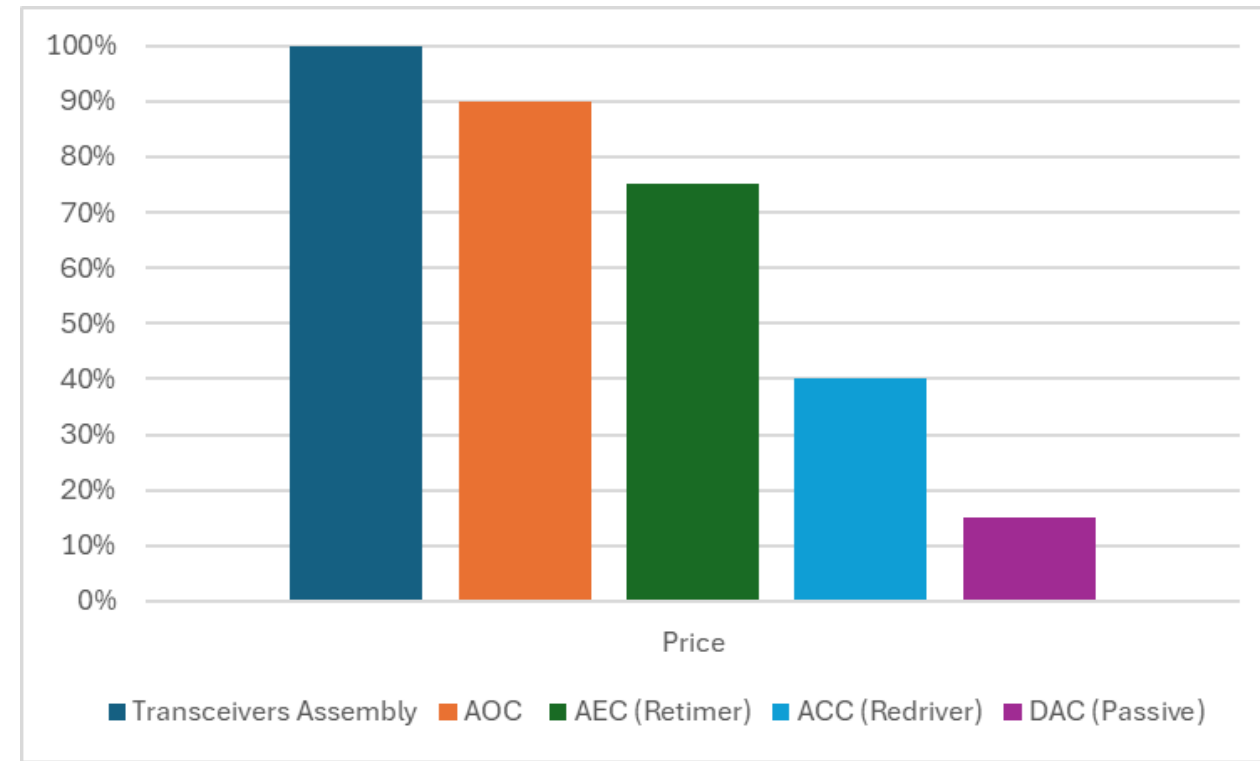
# of GPUs	2048	3072	4096
Typical	105kW	158kW	210kW
Efficient	37kW	55kW	74kW
Difference	<b>68kW</b>	<b>102kW</b>	<b>136kW</b>

# Cabling Considerations – Highest To Lowest

## POWER From Highest To Lowest




## PRICE From Highest To Lowest



\*Reducing POWER also reduces HEAT in your system

\*1.6T - +33.5 watts per port









# TIA 568.3

- **ANSI/TIA-568.3-E-1 [Optical Fiber Cabling and Components]:**
- Completed comment resolution
  - MM array connectors could be angled or non-angled
  - Text added to help explain polarity diagrams
  - Connector Type: Duplex, Array, Duplex Array
  - Replaced LC Connector with horizontally oriented duplex fiber connectors
  - Replaced MPO adapter with horizontally orientated array
  - Add section mirroring horizontally oriented array connectors for vertically oriented array connectors
  - Adding reference to alternative configurations in MPO16 section
    - MPO-16 (16f) to 4 x MPO-12 (4f)
    - MPO-12 (8f) to 2 x MPO-12 (4f)
  - Revising Method B schematic to show Type A or B adapter
- Approved new revision of 568.3

# What to Expect in the Future

- Anticipate diverse approaches to achieving requisite throughput.
- InfiniBand GDR
- NVIDIA is poised for an imminent release of 1.6T capabilities.
- Ethernet 200G per lane:\* A mid-2026 release is projected, 802.3dj This standard encompasses 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6Tbf physical layers.
- CFI Presentation for 400G per lane – Approved March 12, 2026
- 200G signaling with reaches much shorter than the shortest links defined in P802.3dj
- Twinax Integration: This includes a clear path toward Active Copper solutions.
- Expanded Single-Mode Optics: Expect a broader array of options in this domain. Prepare for continued evolution beyond 1.6T speeds,
- Adopt a Hybrid Approach: Leverage early AI training in cloud environments.
- Maintain continuous awareness of evolving industry trends..

Christian Ramírez Montilla  
RCDD, DCDC, CCRE  
Technical Manager LATAM  
christian\_ramirez@siemon.com  
Phone: + 57 3185359152

