




Co-Packaged Optics Informative Outlook and Recommended Practices

VERSION 1.0





Copyright

Copyright © 2026 Advanced Photonics Coalition. All Rights Reserved.

Limitations

THIS DOCUMENT IS PROVIDED “AS IS.” The publisher and contributors expressly disclaim any warranties (express, implied, or otherwise), including but not limited to implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title. The entire risk of using or implementing this document is assumed by the user. IN NO EVENT IS THE PUBLISHER OR ANY CONTRIBUTOR LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING LOST PROFITS) BASED ON ANY THEORY OF ACTION (INCLUDING NEGLIGENCE), EVEN IF THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. This document may be subject to third party rights.

CPO Working Group Members

2Pi Optics	Hakusan Inc.	Pilot Photonics Ltd.
Advantest Corporation	Hewlett Packard Enterprise	Qnity Electronics, Inc.
Alcyon Photonics	HYC Co., Ltd.	Quantifi Photonics (a Teradyne Company)
America Fujikura Ltd.	IBM Corporation	Resolight
Amphenol Corporation	Innolume GmbH	Resonac Corporation
Ansys Inc.	Intel Corporation	Santec Corporation
Ayar Labs	Keysight Technologies, Inc.	SENKO Advanced Components Inc.
Belden Inc.	Lightera	Siemens EDA
Ciena Corporation	Lightium AG	Shenzhen Adtek Technology Co., Ltd
ColorChip Ltd.	Lightmatter Inc.	SiLith Technology PTE Ltd.
CommsScope (an Amphenol Company)	LightSpeed Photonics PTE Ltd.	Silitronics Solutions, Inc.
Dexerials Corporation	Luceda Photonics	Sivers Semiconductor, Inc.
Doosan Corporation	MegaCool Technologies	SUNA Optoelectronics Co. Ltd.
EXFO Inc.	Mesh Semiconductor	Synopsys Inc.
ficonTEC Service GmbH	Microsoft Corporation	TE Connectivity
Flexcompute Inc.	MPI Corporation	Teradyne Incorporated
Focuslight Technologies Inc.	NEXUSTEST PTE Ltd.	Tokyo Ohka Kogyo Co., Ltd. (TOK)
Fujikura Ltd.	NTT DEVICES AMERICA	VIAVI Solutions
Furukawa Electric Co., Ltd.	OpenLight Photonics, Inc.	WaferForge LLC
Genuine Optics	Panduit Corporation	WUS Printed Circuit Co., Ltd.
Globalfoundries Inc.	PHIX Photonics Assembly	Xscape Photonics, Inc.
Go!Foton Corporation		



Table of Contents

1. Introduction	1
1.1. The Inflection Point In Data Center Connectivity	1
1.2. Where CPO Fits In The Data Center Architecture	4
1.3. The Critical Importance Of A Multi-Vendor Ecosystem	6
1.4. Market Momentum And Industry Validation	7
1.5. Extra Dense Pluggable Optics (XPO):	8
2. CPO Systems Overview	10
2.1. Optical Interface Requirements	12
2.2. Transmitter Requirements	13
2.3. Receiver Requirements	14
2.4. Thermal Management	18
2.5. Power Delivery Network (PDN) Outlook	20
2.6. Mechanical And Form Factor Specification	22
2.7. Technology Roadmap And Evolution	25
2.8. Multi-Vendor Interoperability Requirements	26
2.9. Test And Measurement	27
3. CPO System Enabled Overview	29
3.1. Connectivity Options Between Compute And Front Panel	30
3.2. Optical Engine Overview	32
3.3. Optical Engine Architecture	33
3.4. Optical Interface Informative Outlook	35
3.5. Fiber Attach And External Connector Interface	35
3.6. Electrical Interface Specifications	36
3.7. Mechanical And Form Factor Specifications	37
3.8. Power Delivery	40
3.9. Control, Management, And Telemetry	41
3.10. Reliability, Availability, And Serviceability	42
4.0 Mid – Board Connection	43
4.1. Mid-Board Connector as a Solution to Fiber Routing Concerns	44



5.0 Laser Sources	48
5.1. Informational Outlook and Recommended Practices Scope	48
5.2. Laser Source Technologies Overview	49
5.3. Laser Source Roadmap 2027-2023	53
5.4. Heterogeneous Integration Approaches	54
5.5. Redundancy and Fault Management	57
5.6. Laser Source Requirements at Scale-Out and Scale Across Levels	59
5.7. CPO Laser Integration Thermal Management	60
5.8. Laser Reliability, Qualification, and Lifetime	62
5.9. Manufacturing and Supply Chain	63
5.10. Regulatory and Environmental Compliance	64
5.11. Detailed Electrical Interface Specifications	65
6.0 Glass Interposer And Photonic Integrated Circuits	66
6.1. Fiber Attachment And FAU Interface	69
6.2. Photonic Integrated Circuit (PIC)	69
6.3. CPO Assembly And Co-Integration Standards	72
6.4. Thermal Management Specifications	73
6.5. Reliability And Environmental Qualification	74
6.6. Test And Measurement Methodology	74
6.7. Interoperability And Interface Requirements	75
7.0 Co-Packaged Optics Faceplate Design And Requirements	77
7.1. Mechanical Design Requirements	81
7.2. Thermal Requirements	83
7.3. Optical Interface Specifications	83
7.4. Scale-Up Architecture Requirements (Multi-Chassis)	86
7.5. Power Requirements	87
7.6. Management And OAM Interface Requirements	87
7.7. Technology Roadmap And Planning Horizon (2027–2035)	88
7.8. Compliance, Testing, And Qualification	89
7.9. Requirements Traceability Matrix	90



8.0 Test And Measurement	91
8.1. Multivendor Interoperability Testing Principles	92
8.2. Scale-Up Vs. Scale-Out Specific Requirements	92
9.0 CPO Systems And Power	94
9.1. Optical Interface Parameters	100
9.2. Multi-Rack Scale-Up System Implementations	101
9.3. OCS Control Plane	103
9.4. Scale-Out System Implementations	103
9.5. Management, Monitoring, And Telemetry	105
9.6. Reliability, Availability, And Serviceability Compliance	107
9.7. Technology Roadmap 2026–2034	109
9.8. Power Consumption Estimation Worksheet	110
10.0 Cooling Systems	111
10.1. Market Drivers	112
10.2. Thermal Load Growth Projections	112
10.3. Single-Chassis CPO Cooling Requirements	114
10.4. CPO Scale-Up Installations: Cooling Requirements	117
10.5. CPO Scale-Out Installations: Distributed Multi-Rack Requirements	120
10.6. CPO Cooling Compliance, Testing, And Future Roadmap	123
11.0 CPO Management Systems And Control	125
11.1. CMIS As The Chassis-Level Management Interface	126
11.2. Multi-Year CPO Deployment	127
11.3. Security Architecture And Security Design Principles	128
11.4. Sustaining CPO Management For Five To Ten Years	128
11.5. AI-Driven Management Extensions	129
12.0 Glossary Of Terms And Definitions	130
13.0 Acronyms	135
14.0 External Standards Alignment	139



List of Figures

Figure 1 - Characteristics of Traditional Modules and Co-Packaged Optics.....	2
Figure 2 - The Revolution of Co-Packaged Optics	3
Figure 3 - CPO's Primary Applications are Scale-Up and Scale-Out.....	5
Figure 4 - The Advanced Photonics Coalition Open Ecosystem.....	6
Figure 5 – XPO and CPO Key Characteristics	9
Figure 6 – Thermal Management in CPO	18
Figure 7 – Packaging both the Optical Engine and the Switching Chip.....	29
Figure 8 – 51.2T Switch CPO Layout Example.....	30
Figure 9 – The Host ASIC Connects to the Optical Engine over CEI.....	33
Figure 10 – Fiber Pigtail Lengths Required for Interior Connections.....	43
Figure 11 – Example of Mid-board Connector Solutions	44
Figure 12 – Connection Points between Transmit and Receive PMDs.....	44
Figure 13 – 204.8T / 1024 x 200G (-DR).....	45
Figure 14 – 204.8T / 1024 x 200G (-DWDM 8 lambda).....	46
Figure 15 – 204.8T / 1024 x 200G (-DR multicore 4CF)	46
Figure 16 – Mid-board Connector Option	47
Figure 17 – Splice / fan out / Shuffle Box Option.....	47
Figure 18 – Pluggable Optics Laser Failure vs CPO Package Fails	49
Figure 19 – Glass Interposer and Photonic Integrated Circuits	67
Figure 20 – The Faceplate of a 51.2T switch box.....	77
Figure 21 – ELS Optical Connectivity Design Options	79
Figure 22 - Chassis supporting direct liquid cooling (DLC) for ASIC cold.....	97



List of Tables

Table 1 - XPO vs. CPO: A Technical and Commercial Comparison	8
Table 2 - Four CPO module classes.....	12
Table 3 - Wavelength grids for CPO intra-chassis links.....	12
Table 4 - Modulation Formats	12
Table 5 - Transmitter Parameters	13
Table 6 - CPO Power Budget Requirements	14
Table 7 - Intra Chassis Link Budget.....	14
Table 8 - Scale-Up Intra-Rack Link Budget	15
Table 9 - Serial Interface Roadmap	15
Table 10 - Power Rails.....	16
Table 11 - Standardized Electrical Connectors.....	17
Table 12 - Operating Temperature Ranges.....	19
Table 13 - Voltage Regulator Requirement	20
Table 14 - Power Management State Machines.....	20
Table 15 - CPO Module Mechanical Dimensions.....	22
Table 16 - Environmental and Compliance	23
Table 17 - MTBF Target and Methodology	23
Table 18 - APC CPO Informative Outlook Roadmap	25
Table 19 - Silicon Photonics Technology Milestones.....	25
Table 20 - Coherent optics integrated in CPO modules.....	26
Table 21 - Backward and Forward Compatibility	27
Table 22 - Required Test Equipment	28
Table 23 - Electrical Capabilities of Pluggable, Optics, and CPO Designs.....	29
Table 24 - 51.2T switch CPO layout example	30
Table 25 - Optical fiber connectors and 1RU connector density.....	31
Table 26 - System Classes	32
Table 27 - Fiber Attach and External Connector Interface.....	35
Table 28 - Low-speed interface signal set required at the Engine Site.	37
Table 29 - Channel count and integration choices.....	37
Table 30 - Representative Floorplan Envelopes.....	38
Table 31 - Normative Thermal Operating Envelope	39
Table 32 - The Normative Power Rail Set.....	40
Table 33 - Power Budget.....	40
Table 34 - Laser Source Electrical Parameters	51
Table 35 - Laser Source Optical Parameters.....	51



List of Tables

Table 36 - Thermal and Environmental Parameters	52
Table 37 - Laser Reliability Requirements	52
Table 38 - Key Parameter Roadmap 2027-2030.....	53
Table 39 - Chassis-Level Optical Performance	56
Table 40 - Thermal Electric Cooling Specification and Control.....	60
Table 41 - Specific Stress Conditions in CPO Environments and Acceptance Criterion	62
Table 42 - Manufacturing Readiness Level (MRL) Milestones	63
Table 43 - Dimensional tolerances at the Photonic Engine Optical Interface (PEOI) plane	65
Table 44 - Material Parameters	66
Table 45 - Dimensional Parameters.....	67
Table 46 - Waveguide Parameters	68
Table 47 - Electrical Redistribution Layer Parameters.....	68
Table 48 - Fiber Attachment and FAU Interface Parameters	69
Table 49 - PIC Platform Classification	69
Table 50 - Electro-Optic Modulator Requirements.....	70
Table 51 - Photodetector Requirements	70
Table 52 - Optical Multiplexer / Demultiplexer Parameters.....	71
Table 53 - PIC Die Physical and Integration Parameters	71
Table 54 - Bonding and Attachment Technologies	72
Table 55 - Optical Path and Alignment Budget.....	72
Table 56 - Operating and Storage Temperature	73
Table 57 - Power Density Units	73
Table 58 - Host ASIC Interface.....	75
Table 59 - Fiber Channel Density for Various Optical Connectors	78
Table 60 - Dimensional Envelope	81
Table 61 - Port Density Outlook	81
Table 62 - Fiber Array Unit (FAU) Alignment.....	82
Table 63 - Environmental and Durability Requirements.....	82
Table 64 - Airflow and Thermal Envelope	83
Table 65 - Fiber Type and Reach Definitions.....	83
Table 66 - Fiber Type, Reach and Front Panel Connector Type vs. Applications.....	84
Table 67 - Optical Loss Budget (Link Level).....	84
Table 68 - Electrostatic Discharge Protection	84
Table 69 - Skew and Latency Requirements	85
Table 70 - Wavelength Management Requirements	85
Table 71 - Technology Roadmap Summary	86

List of Tables

Table 72 - Mechanical Qualification Tests.....88

Table 73 - Optical Qualification Tests 88

Table 74 - Thermal Qualification Tests.....89

Table 75 - Requirements Traceability Matrix.....89

Table 76 – Scale-up Reach.....91

Table 77 – Scale-out Reach.....92

Table 78 – Scale-Out Network Protocols by Tier.....102

Table 79 – Management and Monitoring Standards. 104

Table 80 – Reliability, Safety, and Environmental Compliance Requirements.....106

Table 81 – Thermal Load Growth 111

Table 82 – CPO Thermal Classification.....112

Table 83 – Configuration Profile.....114

Table 84 - Coolant Specifications for CPO Systems.....115

Table 85 - Zone-CDU Requirements for Scale-Up Installations.....116

Table 86 - Distribution Piping for Scale-Out CPO 119

Table 87 - Scale-Out Thermal Zone Layout.....121

Table 88 - Ongoing CPO Cooling Compliance Monitoring 123

1.0 Introduction

1.1. The Inflection Point in Data Center Connectivity

The global data center industry stands at a pivotal technological crossroads as the exponential growth of artificial intelligence workloads, large-scale machine learning training clusters, and high-performance cloud computing platforms is placing demands on interconnect bandwidth and energy budgets that future conventional pluggable transceiver architectures are structurally unable to satisfy. Co-Packaged Optics (CPO) is the integration of photonic components in close physical proximity to, or directly within, the switch or ASIC package and has emerged as the industry's most credible pathway to meet the demands of the today's AI enterprise infrastructure data centers.

Industry analysts project the global hyperscale data center market will surpass \$500 billion by 2030, with AI infrastructure representing the fastest-growing segment. Analyst firms Morgan Stanley, Dell'Oro Group, and LightCounting collectively forecast that CPO will capture a significant share of data center optical interconnect spend within this decade, with the CPO addressable market alone reaching eight to twelve billion annually by 2028.

The formation of Advanced Photonics Coalition dedicated CPO Standards Working Group reflects the industry's collective recognition to unlock CPO's full potential requiring coordinated, open, multi-vendor interoperability standards. This CPO Informative Outlook and Recommended Practices document provides the technical and commercial rationale underpinning Advanced Photonics Coalition's mission.

For more than two decades, the pluggable transceiver with its successive generations from SFP through QSFP-DD and OSFP, has served as the workhorse of data center optical interconnect. This architecture places the electrical-to-optical conversion at the face plate of the switch, connected to the switch ASIC via high-speed SerDes lanes routed across the printed circuit board. As data rates have scaled from 10G to 400G, and now toward 800G and 1.6T per port, this physical separation has become an increasingly severe liability for several interconnected reasons as the PCB trace from the ASIC to the faceplate should carry electrical signals at frequencies exceeding 100 GHz for next-generation rates.

At these speeds, the lossy nature of FR4 dielectric and connector interfaces demands aggressive equalization, digital signal processing, and retimer circuitry, adding latency, increasing bill-of-materials cost, and critically consuming significant power. Industry estimates place the per-port SerDes power penalty for pluggable 800G modules at five to ten watts relative to a Co-Packaged Optics equivalent.




Figure 1 – Characteristics of Traditional Modules and Co-Packaged Optics

A modern 51.2T switch chassis populated with $64 \times 800\text{G}$ OSFP pluggables dissipates excessive watts. Thermal engineering at this density forces air-flow compromises that degrade ASIC junction temperatures and limit overclocking headroom. CPO relocates this thermal load to a zone where heat management may be architecturally co-optimized with the package and system design.

The physical width of each pluggable cage imposes a hard limit on the number of ports that may be accommodated at the switch face plate. CPO eliminates this geometric constraint by taking the optical engine off the faceplate entirely, enabling the full I/O bandwidth to be routed through fiber arrays attached directly to the package substrate.




Figure 2 – The Revolution of Co-Packaged Optics

CPO addresses each of these limitations through a fundamental re-partitioning of where electro-optic conversion occurs. By integrating silicon photonics dies, micro-ring modulators, or indium phosphide laser arrays within or immediately adjacent to the switch package, the electrical signal path between the SerDes output and the optical modulator is reduced from centimeters to sub-millimeters. The resulting benefits are quantitative and substantial:

- **Insertion Loss:** Electrical insertion loss drops by 10–15 dB, eliminating the need for linear re-drive amplifiers.
- **Signal Quality:** SerDes eye opening improves materially, allowing lower-power equalization or removal of retraining circuitry entirely.
- **Latency:** Latency is reduced, benefiting time-sensitive distributed computing primitives such as collective operations in AI training.
- **Power Efficiency:** Optical engine power density improves because silicon photonics may be fabricated on advanced process nodes, and because the elimination of the PCIe/electrical interface to a remote module removes an entire power rail.

The net result is a system that is not merely incrementally better than pluggable at the same data rate, but one that fundamentally resets the power-per-bit and bandwidth-density equations. This is why CPO is not positioned as one option among many but is the architecture that the laws of physics favor as data rates continue to scale.

1.2. Where CPO Fits in the Data Center Architecture

Modern AI data center fabrics can be divided into three broad categories based on the scale of connectivity, that is, the number of XPU each fabric interconnects:

- **Scale-Up fabric (~100 kW):** Interconnects hundreds of XPUs in a low-latency, cache-coherent manner over an all-to-all network. This is the preferred fabric for inference workloads.
- **Scale-Out fabric (~1 GW):** Interconnects up to a million XPUs within a data center through a two-to three-tier Ethernet or InfiniBand network, and is typically deployed for training workloads.
- **Scale-Across fabric (2–5 GW):** Interconnects multiple data centers over long-haul networks, typically to run a single training job across multiple buildings or to aggregate capacity across a campus.

CPO is not a single-use technology confined to one segment of the data center and its value proposition applies across multiple interconnect domains. The most immediate and commercially urgent application for CPO is within the scale-up fabric of AI accelerator clusters. Modern large language model training requires synchronization across hundreds or thousands of GPU and AI accelerator dies, exchanging gradient data through collective operations at extremely high bandwidth and extremely low latency. CPO's low-latency, high-bandwidth, and power-efficient profile makes it ideally suited for the switch tier within scale-up domains, where port counts in the range of 512–1,024 per switch are being targeted and where the power consumed by optical I/O directly impacts total cluster operations.

In traditional cloud and enterprise data center architectures, CPO offers compelling advantages in the scale-out spine-leaf fabric that interconnects server racks across the facility. At the spine layer, where 51.2T and 102.4T switches are deployed in dense configurations, the power and space savings of CPO translate directly into lower power usage effectiveness metrics and reduced cooling infrastructure capital expenditure.




Figure 3 – CPO's Primary Applications are Scale-Up and Scale-Out

Beyond the primary scale-up and scale-out use cases, CPO presents opportunities in several adjacent domains:

- **Disaggregated Storage:** Storage networking and disaggregated non-volatile memory fabrics, where bandwidth density per rack unit is constrained by pluggable form factors.
- **ASIC Integration:** Co-packaged optical engines integrated with ASIC, enabling line-rate offload with reduced host CPU involvement.
- **Coherent Edge Applications:** Front-end optical transceivers for coherent metro applications, where the integration of DSP and photonics in a common package reduces module cost and power.
- **HPC Interconnects:** High-performance computing interconnects in national laboratory and scientific computing environments, where bisection bandwidth and Message Passing Interface (MPI) latency are dominant design metrics.

1.3. The Critical Importance of a Multi-Vendor Ecosystem

The history of optical networking contains important lessons about the relationship between standardization and market growth. The explosive adoption of pluggable transceivers beginning with SFP and continuing through QSFP-DD and OSFP was not driven by any single vendor's proprietary solution. It was driven by the confidence that purchasers could source compatible modules from multiple qualified suppliers, that system vendors could design to a stable interface specification, and that the resulting competition drove rapid cost reduction.

CPO faces a more complex integration challenge than pluggable modules as it requires co-design of the optical engine, the switch ASIC package, the PCB, the system thermal architecture, and the fiber management infrastructure. Without open ecosystem driven standards, the natural gravitational pull is toward vertically integrated proprietary solutions, an outcome that would fragment the market, increase customer lock-in risk, slow adoption, and ultimately diminish the opportunities of innovation for CPO's technical benefits.




Figure 4 – The Advanced Photonics Coalition Open Ecosystem as of Sept. 2026

Advanced Photonics Coalition Multi-Vendor Standardization Innovation Hub and Open Ecosystem offers:

- **Supply Chain Resilience:** Operators may qualify multiple CPO optical engine suppliers against a common interface specification, preventing sole-source dependency for a critical infrastructure component.
- **Platform Design Efficiency:** System vendors may design merchant-silicon switch platforms that accept CPO engines from any compliant supplier, lowering NRE costs and accelerating time to market.
- **Cost Reduction:** Competitive sourcing across a multi-vendor ecosystem drives down cost per bit over time, as has been demonstrated in every previous optical interconnect generation.
- **Operational Simplicity:** Open diagnostic and management interfaces allow operators to use common tooling, monitoring platforms, and open application model frameworks regardless of CPO vendor.
- **Ecosystem Innovation:** Startups and emerging photonics innovators may enter the CPO market without requiring deep integration agreements with ASIC vendors, accelerating the innovation cycle.
- **Industry Signals on Standardization:** The industry's growing commitment to CPO standardization is evident from several concurrent initiatives. The Advanced Photonics Coalition integrates implementation efforts into a coherent, end-to-end specification framework covering the optical, electrical, mechanical, thermal, and management interfaces simultaneously, the gap that this organization is formed to fill.

1.4. Market Momentum and Industry Validation

The commercial momentum behind CPO has accelerated substantially in 2025. Several industry leaders have demonstrated proprietary solutions of CPO-enabled switch and compute platforms. The common thread across these programs is validation of the core thesis: CPO delivers a step-change improvement in power-per-bit and bandwidth-density that is not achievable through incremental evolution of the pluggable paradigm.

From a regulatory and sustainability standpoint, CPO's energy efficiency credentials are increasingly relevant. CPO's 30–50% reduction in optical interconnect power represents a material contribution to these sustainability goals, a factor that will accelerate procurement decisions and regulatory favor. Every watt saved at the optical interface translates directly into reduced cooling load, lower carbon emissions, and greater computational throughput per dollar of infrastructure investment.

Advanced Photonics Coalition's (APC) open ecosystem membership is open to all organizations with a material interest in all aspects of the CPO supply chain. The Advanced Photonics Coalition operates under a consensus-based governance model with intellectual property policies designed to encourage broad participation and timely specification publication.

1.5. Extra Dense Pluggable Optics (XPO):

As the data center industry accelerates toward 800G per lane and beyond, a parallel technological trajectory has emerged alongside co-packaged optics called Extra Dense Pluggable Optics (XPO). XPO represents a fundamental re-engineering of the pluggable transceiver paradigm — rather than relocating the optical engine to the package, XPO seeks to radically compress the form factor, power envelope, and electrical interface of pluggable modules to close the gap with CPO on key performance dimensions.

The table below summarizes the principal technical and commercial dimensions across which XPO and CPO diverge as the industry scales toward 800G per lane and 1.6T per port.

Dimension	Extra Dense Pluggable Optics (XPO)	Co-Packaged Optics (CPO)
Architecture	Optical engine remains external to ASIC package; mounted on PCB in near-package or faceplate topology. Module is pluggable and field-replaceable.	Optical engine integrated within or immediately adjacent to ASIC package substrate. Electrical path to modulator reduced to sub-millimeter scale.
Electrical Path Length	30–250 mm. Retimers may still be required at 800G and above in faceplate configurations.	Sub-1 mm to ~5 mm. Eliminates retimers entirely; equalization requirements are fundamentally reduced at any lane rate.
Power Efficiency at 800G	Estimated 12–18 W per 800G port in optimized near-package XPO configurations. Conventional faceplate OSFP 800G: 15–22 W per port. 15–25% improvement over conventional pluggable possible.	Estimated 6–10 W per 800G port. 30–50% total system power reduction versus conventional pluggable. Physics-driven advantage grows as lane rates increase beyond 800G.
Serviceability	Field-replaceable without switch downtime. Operators may swap individual modules independently.	Optical engine replacement requires switch removal or complex in-situ procedures. Detachable FAU architectures partially address this, but full field serviceability remains an active standardization challenge.
Ecosystem Readiness	Mature multi-vendor ecosystem. Operators, system integrators, and supply chains are fully optimized for pluggable form factors. Near-term deployment confidence is high for 800G XPO.	Ecosystem in active formation. Multi-vendor interoperability standards (Advanced Photonics Coalition mandate) are the critical missing piece.
Long-Reach Optical Interface	Native long-reach capability through standard LC/MPO connectors. Supports DR4, FR4, and LR4 optical interfaces without external components.	Optimized for intra-rack and intra-pod distances. Longer reach requires external optical amplification or a hybrid CPO/pluggable border architecture.
System Redesign Requirement	Near-package XPO requires PCB layout changes but does not require new ASIC packaging. Incremental change from existing switch designs. Lower NRE cost for system vendors.	Requires co-design of ASIC package, optical engine, PCB, and thermal architecture. Significant NRE investment but delivers step-change performance that justifies the investment at 800G+ scale.

TABLE 1 – XPO vs. CPO: A Technical and Commercial Comparison

The technology trajectory toward 1.6T per lane is where the fundamental divergence between XPO and CPO becomes most consequential. The industry's current generation of deployments at 100G and 200G per lane has sustained the pluggable paradigm through a combination of PAM4 modulation, advanced DSP, and process node scaling. As lane rates double again to 400G and 800G, the physics of electrical signal transport over PCB traces become increasingly unforgiving in ways that incremental XPO improvements may not fully mitigate.

CPO's long-term roadmap at 800G per lane and above is anchored in a fundamentally more favorable physical regime. Because the electrical interface between SerDes and modulator is measured in microns rather than millimeters, the signal integrity constraints that scale adversely with lane rate in pluggable topologies do not apply in the same way. CPO optical engines designed for 800G per lane operation may leverage the same silicon photonics process node scaling that has driven improvements in ring modulator efficiency, photodetector responsivity, and wavelength division multiplexing channel count, all without the compounding penalty of long electrical traces.




Figure 5 – XPO and CPO Key Characteristics

At the system level, the transition to 800G per lane also enables a reduction in lane count for a given port bandwidth: a 3.2T port may be served by 4 lanes at 800G rather than 16 lanes at 200G, dramatically simplifying the fiber array unit, FAU receptacle, and on-package optical routing architecture. This lane count reduction is a secondary benefit of CPO at high lane rates that further improves system density and reduces assembly complexity in ways that do not apply to XPO. CPO is the dominant architecture in hyperscale AI training clusters, scale-up fabrics operating at 800G per lane and above, and any deployment context where power-per-bit is the dominant design constraint. As the CPO ecosystem matures through the standards work of the Advanced Photonics Coalition, the costs of CPO adoption including NRE, supply chain development, and serviceability complexity will decrease, progressively expanding the share of the market in which CPO is the rational choice over XPO.

2. CPO Systems Overview

CPO Systems Overview covers:

- Single-chassis systems accommodating 4, 8, 16, and 32 or more ASIC or switch configurations.
- Scale-up architectures extending across two or more chassis within the same rack.
- Scale-out architectures spanning multiple racks connected via an optical fabric.
- The full-stack optical interface from the photonic integrated circuit (PIC) die boundary through the optical fiber management system to the rack-level patch panel.
- Electrical interfaces between ASIC packages and CPO modules on the system board.
- Power delivery networks (PDN) optimized for co-packaged optic modules.
- Thermal design constraints for single-chassis and multi-chassis CPO deployments.
- Management, telemetry, and in-band/out-of-band monitoring.
- Electromagnetic compatibility (EMC) and signal integrity (SI) specifications.
- Reliability, qualification, and lifecycle specifications.

CPO Architectural Overview

There are three hierarchical system tiers that share a common optical interface standard but differ in topology, scale, and latency profile:

- **Single Chassis:** 4 ASICs within a single mechanical enclosure sharing a mid-plane or backplane. CPO modules attach to ASIC packages via ultra-short electrical traces or direct bump attachment.
- **Scale-Up Multi-Chassis:** 2 to 8 chassis within a single rack cabinet connected by an intra-rack optical fabric. CPO modules interface between chassis via short-reach or medium-reach single-mode fiber.
- **Scale-Out Multi-Rack:** Multiple racks (up to a full pod) interconnected by a high-radix optical switching fabric. CPO modules interface via medium-reach single-mode fiber, with optional coherent pluggable uplinks for inter-pod connectivity.

The 4-ASIC configuration is the minimum compliant baseline. It is intended for inference pods, edge inference, and cost-sensitive training workloads. Physical form factor should accommodate 2U chassis height minimum. All four ASICs should be connected in a full-mesh optical topology with point-to-point CPO links.

Scale-Up Architecture

Scale-up extends a single coherent logical domain across multiple chassis. The defining characteristic of scale-up is shared memory or tightly coupled collective operations at sub-microsecond latency across chassis boundaries. For 2 to 4 chassis in a rack, a dedicated Intra-Rack Optical Switch (IROS) module may be installed. The IROS module is a 1OU unit supporting up to 128 x 128 single-mode ports. Latency through the IROS switch should be ≤ 100 ns end-to-end (including fiber propagation for ≤ 3 m patch cables). The IROS may use photonic circuit switching (microsecond reconfiguration), silicon photonic switch matrices, or passive star coupler configurations depending on traffic pattern.

Scale-Out Architecture

Scale-out connects independent compute domains via a high-radix, non-blocking optical fabric. Latency requirements are relaxed compared to scale-up, but bandwidth and oversubscription ratios are tightly controlled.

- **Leaf tier:** CPO-equipped chassis-level switch or top-of-rack (TOR) optical module; ≥ 204.8 T switching capacity; connects ≥ 16 chassis downlinks and ≥ 16 spine uplinks.
- **Spine tier:** High-radix optical switch; ≥ 204.8 T per spine; radix ≥ 512 optical ports; wavelength-selective switching (WSS) based, targeting 2027+.
- **Super-spine / DCI tier (optional):** Coherent WDM based; 400ZR or 800ZR modules; inter-pod or inter-datacenter.

2.1 Optical Interface Requirements

APC defines CPO module classes based on ASIC attachment method and photonic integration level:

Class	Integration Level	Attachment Method	Fiber Type	Target Generation
On-Package	PIC co-packaged with ASIC in MCM	Electrical bumps / μ Bumps ≤ 2 mm	ITU-T Standard G.652.D or G.657	2027–2029
Flip-Chip PIC	PIC flip-chipped directly on ASIC interposer	Direct electrical bond	Lensed fiber / FAU	2028–2031
Monolithic Photonic	Photonic die and compute die co-integrated	Optical waveguide on-chip	Nano-coupler to SMF	2031–2035

Table 2 – A Portion of CPO Module Classes

Note: For transparent-switch deployments the OCS-facing outputs may require PMF FAU.

Short-Reach and Medium-Reach WDM

Band	Wavelength Range (nm)	Channel Spacing	Max Channels	Application	Year
O-Band	1295–1310	4.5 nm (CWDM4)	4	Short-reach intra-chassis, 100m–2km	2027+
C-Band (8ch)	1530–1565	400 GHz ITU-T	8	Medium-reach scale-up, 2–10km	2027+
C-Band (16ch)	1530–1565	200 GHz ITU-T	16	Scale-up, scale-out uplinks	2027+
C-Band (32ch)	1530–1565	100 GHz ITU-T	32	High-density scale-out fabric	2028+
C+L-Band (80ch)	1530–1625	50 GHz ITU-T	80	Long-haul DCI / multi-pod	2030+

Table 3 – Wavelength Grids for CPO intra-Chassis Links

Modulation Formats

Format	Symbol Rate	Bit Rate/Lane	Lanes/Module	Aggregate/Module	Application	FEC Required
NRZ	56GBd	56G	16	896G	Baseline 2027	KP4
PAM4	56GBd	112G	16	1.792T	Primary 2027–2028	KP4
PAM4	112GBd	224G	16	3.584T	Next-gen 2028–2030	KP4+
16-QAM coherent	100GBd	400G	8 (dual-pol)	3.2T	Scale-out 2028+	SD-FEC
64-QAM coherent	128GBd	~768G	4 (dual-pol)	~3.0T	Long-reach 2031+	SD-FEC

Table 4 – Modulation Formats

2.2 Transmitter Requirements

CPO modules should use one of the following compliant laser source technologies:

- **External Cavity Laser (ECL) with on-chip modulator:** preferred for lowest RIN and narrowest line-width; suitable for coherent and direct-detect WDM.
- **Distributed Feedback (DFB) laser array:** acceptable for CWDM and short-reach applications; RIN < -140 dB/Hz at operating bias.
- **Vertical Cavity Surface Emitting Laser (VCSEL):** permitted only for CPO modules from intra-chassis to inter-rack applications up to 100m with OM4/OM5 fiber . up to 100m applications with OM4/ITU-T Standard G.652.D or G.657 fiber.
- **Integrated Optical Frequency Comb (OFC):** emerging technology approved for multi-wavelength modules after 2029; comb spacing should be ITU-T compliant.
- **Quantum dot (QD) lasers:** Semiconductor lasers that use nanoscale semiconductor crystals (2-10 nm) as the active gain medium. Because charge carriers are confined in all three spatial dimensions, these devices offer superior temperature stability, lower threshold currents, and a broader gain spec-

Transmitter Optical Specifications

Parameter	2027	2028	2029–2031	Units
Average Launch Power	-1 to +4	+1 to +5	+2 to +6	dBm per lane
Extinction Ratio (PAM4)	≥ 3.5	≥ 4.0	≥ 4.5	dB
TDEC (PAM4, 56 GBd)	≤ 5.5	≤ 5.0	≤ 4.5	dB
TDEC (PAM4, 112 GBd)	—	≤ 6.5	≤ 5.5	dB
RIN (DFB source)	< -140	< -142	< -145	dB/Hz
SMSR	≥ 35	≥ 38	≥ 40	dB
Wavelength Accuracy	±0.1	±0.05	±0.025	nm
Laser Startup Time (cold)	≤ 2000	≤ 1500	≤ 1000	μs

Table 5 – Transmitter Parameters

2.3 Receiver Requirements

Compliant CPO receiver sub-assemblies should use germanium-on-silicon photodetectors for silicon photonic implementations or InGaAs photodetectors for InP-based designs. Avalanche photodiodes (APDs) are permitted for extended reach links with ≥ 10 dB power budget advantage required to justify increased bias complexity.

Parameter	2027	2028	2029–2031	Units
Responsivity (1310 nm)	≥ 0.8	≥ 0.9	≥ 0.95	A/W
Responsivity (1550 nm)	≥ 0.9	≥ 1.0	≥ 1.05	A/W
Overload Power (PAM4)	$\geq +2$	$\geq +3$	$\geq +4$	dBm
Sensitivity (BER 1e-4, PAM4, 56GBd)	≤ -7.5	≤ -8.5	≤ -9.5	dBm
Sensitivity (BER 1e-4, PAM4, 112GBd)	—	≤ -6.0	≤ -7.5	dBm
TIA Bandwidth (–3dB)	$\geq 0.75 \times \text{BitRate}$	$\geq 0.75 \times \text{BitRate}$	$\geq 0.75 \times \text{BitRate}$	—
TIA Input Noise Current	≤ 12	≤ 10	≤ 8	pA/ $\sqrt{\text{Hz}}$
CDR Lock Time	≤ 1	≤ 0.5	≤ 0.2	ms

TABLE 6 – CPO Power Budget Requirements

Intra-Chassis Link Budget

Parameter	Minimum	Typical	Maximum	Units
Fiber Span Loss (0.5m–3m)	0.1	0.2	0.5	dB
FAU Coupling Loss (each end)	0.3	0.5	1.0	dB
MPO/MTP Connector (each)	0.1	0.2	0.5	dB
Splicing Loss (if any)	0.0	0.05	0.1	dB
Receiver Sensitivity (PAM4, BER 1e-4)	—	–7.5	—	dBm
Optical Power Budget	—	8.5	—	dB
Margin (after all penalties)	≥ 3.0	4.5	—	dB

TABLE 7 – Intra Chassis Link Budget

Scale-Up Intra-Rack Link Budget (1–10 m SMF)

Parameter	Value	Units
Fiber Loss (10m @ 0.35 dB/km)	0.004	dB
Connector Pairs (4x MPO)	0.8–2.0	dB
Dispersion Penalty (10m SMF @ 112GBd)	< 0.1	dB
Receiver Sensitivity	≤ -8.5	dBm
Total Link Power Budget	≥ 10	dB
Minimum System Margin	≥ 3	dB

TABLE 8 – Scale-Up Intra-Rack Link Budget

All CPO modules should be compatible with defined fiber management assemblies:

- FAU-12: 12-fiber single-mode array unit, 127 μ m pitch.
- FAU-24: 24-fiber array for high-density CPO modules.
- MPO-24 Type B: industry-standard ribbon fiber connector for chassis-level patching.
- MPO-16: angled-physical-contact finish preferred for return-loss-sensitive coherent links (≤ -60 dB RL).

Bend radius: fiber management trays should maintain ≥ 10 mm bend radius for ITU-T Standard G.652.D or G.657 and ≥ 30 mm for ribbon fiber bundles. CPO modules should tolerate insertion forces ≤ 10 N at the FAU interface without exceeding 0.2 dB IL change.

The electrical interface between the ASIC and CPO module is the most critical electrical design challenge in CPO systems. As bit rates increase, the manageable electrical trace length drops sharply, making CPO physically adjacent to the ASIC mandatory.

Year	Per-Lane Rate	Modulation	Standard Reference	Max Trace Length (PCB)	Max Bump Reach (MCM)
2027	112G	PAM4	OIF CEI-112G-MR/VSR	≤ 50 mm	≤ 20 mm
2028+	224G	PAM4	OIF CEI-224G-MR (draft)	≤ 25 mm	≤ 12 mm
2031+	896G 1.12T	PAM4 / QAM	TBD	Direct die integration	N/A

TABLE 9 – Serial Interface Roadmap

All electrical host interfaces should comply with the following SI budget:

- Insertion Loss (IL) at Nyquist frequency: ≤ -15 dB for MR class, ≤ -10 dB for SR class, ≤ -5 dB for VSR class (measured from package ball to CPO module ball).
- Return Loss: ≤ -10 dB from 100 MHz to Nyquist.
- Differential to common mode conversion (Scd21): ≤ -30 dB at Nyquist.
- Intra-pair skew: ≤ 0.5 ps for 112G lanes; ≤ 0.25 ps for 224G lanes.
- Inter-lane skew (within a group of 4): ≤ 5 ps for 112G; ≤ 3 ps for 224G.
- Crosstalk (NEXT + FEXT composite): ≤ -35 dB at Nyquist.
- Impedance: $100\ \Omega \pm 10\%$ differential.

CPO Module Power Interface – Power Rails

Rail	Voltage (Nominal)	Tolerance	Max Current (per module)	Usage
VDDIO-Optical	1.8 V	$\pm 2\%$	8 A	TIA, CDR, bias circuits
VDD-Laser	3.3 V	$\pm 2\%$	4 A	Laser drive and TEC bias
VDD-Core	0.9 V	$\pm 1.5\%$	20 A	DSP and serializer cores
VDD-PD	1.5 V (reverse bias ref)	$\pm 1\%$	0.2 A	Photodetector reverse bias
VPZT (coherent only)	± 15 V	$\pm 3\%$	0.5 A	Piezo phase control
VMOD	0.9–1.5 V (adj)	$\pm 1\%$	12 A	Modulator driver
VSUP (management)	3.3 V	$\pm 5\%$	0.5 A	CMIS controller, I2C/GPIO

TABLE 10 – Power Rails

Power rails should be applied in the following order during module initialization: VSUP → VDD-Core → VDDIO-Optical → VDD-Laser → VMOD → VDD-PD. Reverse order on power-down. Sequencing violation should trigger an irreversible hardware latch until explicit soft-reset via CMIS register 0x0026 (Module Reset). CPO modules should tolerate the following dynamic power variation patterns characteristic of AI training workloads:

- Load step: 0 to 100% in $< 1\ \mu\text{s}$ with PSRR ≥ 30 dB at 10 MHz on all optical power rails.
- All Reduce burst: 60–100% load cycling at 1–10 kHz for sustained periods ≥ 60 seconds.
- Idle-to-active ramp: $< 100\ \mu\text{s}$ optical channel activation time from low-power idle state.
- Laser aging compensation: automatic bias current increase ≤ 10 mA over 50,000 ETOP hours with < 0.5 dB launch power change.

Module Connector Interface

Connector Type	Pin Count	Mating Cycles	Application	Availability
Direct MCM	480 signal / LGA	≥ 100 (ZIF)	On-package	2028
Flip-chip	N/A (bumped)	Single permanent	Flip-chip PIC	2029
UCle extended	Compatible with UCle 2.0 bump map	Single permanent	Monolithic	2032

TABLE 11 – Standardized Electrical Connectors

System boards hosting CPO modules and ASIC packages should comply with the following substrate requirements:

- Dielectric loss tangent: ≤ 0.002 at 10 GHz for signal layers in the CPO/ASIC routing region.
- Conductor roughness: $\leq 0.3 \mu\text{m Ra}$ for outer layers in high-speed zones; low-profile or reverse-treated copper foil required above 112 GBd.
- Differential impedance control: $\pm 5\%$ or $\pm 5 \Omega$, whichever is tighter.
- Via stub length: $\leq 100 \mu\text{m}$ back-drilled via stubs at CPO connector vias.
- Ground plane continuity: No slot in reference plane within $3\times$ trace width of any 112G+ signal routing.
- Thermal vias: ≥ 200 vias/cm² under CPO module thermal pad; via thermal resistance $\leq 0.3 \text{ K}\cdot\text{mm}^2/\text{W}$.

2.4 Thermal Management

CPO modules introduce unique thermal challenges: the PIC and laser components are thermally sensitive, yet co-location with high-TDP ASICs (typically 700 W–2 kW per die by 2027) creates extreme local thermal environments. We separate the ASIC thermal domain from the CPO optical domain with mandatory thermal isolation features.




Figure 6 – Thermal Management in CPO

We separate the ASIC thermal domain from the CPO optical domain with mandatory thermal isolation features.

Component	T _{case} Min	T _{case} Max	T _{junction} Max	Laser Max Junction Max	Notes
Module (flip-chip)	0 °C	85 °C	95 °C	40 °C	Microfluidic cooling option
ASIC Package (reference)	5 °C	95 °C	105 °C	N/A	Thermal isolation $\geq 5 \text{ K}\cdot\text{cm}^2/\text{W}$ T _{between} CPO and ASIC
Switch Tile (intra-chassis)	0 °C	75 °C	85 °C	N/A	Shared heatsink with CPO permitted if $\Delta T \leq 15^\circ\text{C}$

TABLE 12 – Operating Temperature Ranges

To prevent ASIC-induced thermal runaway in the laser die, the following isolation measures should be mandatory in all compliant designs:

- Minimum horizontal separation between ASIC lid and CPO PIC substrate: 3 mm for 1st generation; 1 mm for 2nd generation with TEC; 0.5 mm for 3rd generation with microfluidic isolation wall.
- Thermal isolation barrier: aerogel or vacuum gap $\geq 2 \text{ mm}$ equivalent R-value between ASIC and laser die; thermal resistance $\geq 5 \text{ K}\cdot\text{cm}^2/\text{W}$.
- TEC integration: 1st and 2nd generation CPO modules in environments where ambient $> 45^\circ\text{C}$ should incorporate a thermoelectric cooler maintaining laser junction $\leq 35^\circ\text{C}$; TEC efficiency $\geq 40\%$ ($\text{COP} \geq 0.4$) at target cooling delta.
- Thermal runaway protection: hardware-enforced laser shutdown within 10 ms when laser junction temperature exceeds $T_{j\text{max}} + 5^\circ\text{C}$, as measured by on-chip thermal diode with 0.5°C resolution.

Air-forced cooling is most likely be not sufficient and liquid cooling will be required. For air-forced cooling, the front-to-rear airflow at $\geq 300 \text{ LFM}$ over CPO module heatsinks. Inlet air temperature $\leq 35^\circ\text{C}$ (ASHRAE A2) preferred; ASHRAE A3 ($\leq 40^\circ\text{C}$) is required compliance minimum. CPO module fin-array heatsink should achieve $\leq 0.15 \text{ }^\circ\text{C}/\text{W}$ from case to air at 300 LFM.

2.5 Power Delivery Network (PDN) Outlook

CPO chassis should implement a multi-rail hierarchical PDN. Point-of-load (POL) conversion for CPO modules should occur on the system board using high-frequency GaN or GaN-SiC hybrid regulators.

Voltage Regulator Requirements for CPO Rails

VR Parameter	Requirement	Rationale
VDD-Core (0.9V) Transient Response	< 5 mV deviation at 20 A/ μ s step load	DSP cores in CPO modules have aggressive di/dt during equalization adaptation
Output Noise (100 kHz–100 MHz)	< 50 μ Vrms	Required for < –140 dB/Hz RIN floor maintenance
Efficiency	$\geq$ 92% at 50% load, $\geq$ 88% at 10% load	Optical rail efficiency directly impacts system PUE
Switching Frequency	$\geq$ 2 MHz (GaN preferred)	Pushes switching harmonics above 2 MHz optical noise floor
Current Sharing (parallel VRs)	$\pm$ 5% of per-unit rated current	Required for thermal balance in high-current optical rails
Hold-Up Time	Not required on CPO optical rails	Fast shutdown preferred over hold-up to protect lasers
PMBus Telemetry	Vin, Vout, Iout, Temp at $\geq$ 1 Hz	Required for predictive failure detection per Section 10

TABLE 13 – Voltage Regulator Requirement

Power Management State Machine

State	Name	Laser	Tx	Rx CDR	DSP	Power (typical)	Entry Trigger	Exit Latency
P0	Full Active	On	On	On	Full	100%	Default operational	N/A
P1	Low-Power Rx	On	Off	On	Rx only	45%	Software or idle Tx	< 50 μ s
P2	Standby	On (bias only)	Off	Off	Off	15%	Traffic idle > 100 ms	< 1 ms
P3	Sleep	Off	Off	Off	Off	2%	Traffic idle > 10 s	< 100 ms
P4	Deep Sleep	Off	Off	Off	Off	< 0.5%	Directed by host	< 2 s (laser warmup)

TABLE 14 – Power Management State Machines

CPO systems should implement a hierarchical management stack:

- Module Level: CMIS 5.x over I2C/I3C, accessible via host ASIC driver or BMC.
- Chassis Level: IPMI 2.0 or OpenBMC.
- Rack Level: Redfish 2.x REST API aggregating chassis-level data; SNMP v3 for legacy integration.
- Fabric Level: NETCONF/YANG for switch tiles and optical amplifiers; OpenConfig optical transport model.
- Cluster Level: gRPC-based streaming telemetry to infrastructure monitoring (Prometheus, InfluxDB, Grafana).

Telemetry Streaming

High-rate telemetry for AI workload monitoring should be supported via a dedicated telemetry interface (separate from CMIS I2C) operating at ≥ 1 MHz sample rate for the following signals:

- Per-lane optical power (Tx and Rx) at ≥ 1 kHz sample rate.
- Per-lane pre-FEC BER counter delta at ≥ 100 Hz.
- Laser junction temperature at ≥ 10 Hz.
- Module aggregate power consumption at ≥ 100 Hz.
- CDR lock status (1-bit per lane) at ≥ 10 kHz (interrupt-driven event for unlocks).

Telemetry data should be streamed via a sideband PCIe-based DMA interface (preferred) or via a dedicated SMBus-to-Ethernet bridge for out-of-band paths. Telemetry format should comply with OpenConfig Telemetry (RFC 7951 YANG-encoded JSON or Protobuf).

2.6 Mechanical and Form Factor Specification

CPO Module Mechanical Dimensions

Module Class	Length (mm)	Width (mm)	Height (mm)	Weight (g max)	Optical Connector Face	Electrical Connector
Standard	120 ±0.5	40 ±0.3	8 ±0.2	75	FAU (rear)	E1 (front)
High Density	80 ±0.5	30 ±0.3	6 ±0.2	50	FAU (rear)	E1 (front)
On-Package	60 ±0.3	35 ±0.3	5 ±0.15	45	FAU (top)	E2 (bottom LGA)
Flip-Chip	45 ±0.2	20 ±0.2	3 ±0.1	20	Nano-coupler (edge)	Bumped (direct)
Switch Tile (1-port WDM)	200 ±1.0	80 ±0.5	15 ±0.5	300	MPO-24 Angle polished Connector (×4)	E1 (×2)
IROS Module (1U)	482 ×1U rack	300 (depth)	44.45	5,000	LC-Angle polished Connector (×256)	10GbE Mgmt

TABLE 15 – CPO Module Mechanical Dimensions

CPO Module Keying and Anti-Mis-Insertion

Physical keying should prevent incorrect module insertion in all compliant chassis:

- Generation keying: 2026+ modules should have different guide pin locations (offset by 5 mm per generation) preventing insertion into an incompatible socket.
- Speed keying: within current CPO generation, 112G and 224G variants should use different guide keys at the electrical connector.
- Polarization keying: optical fiber connector is asymmetric with a 0.3 mm protrusion on one side preventing 180° misorientation.
- Hot-plug capability: 2026+ modules should support hot-plug insertion and removal with the chassis powered, without ASIC fabric disruption longer than 10 ms (graceful reroute).

Environmental and Compliance

Parameter	Specification	Standard
Operating Temperature	0 to +45°C inlet air (ASHRAE A3)	ASHRAE TC9.9 2021
Storage Temperature	–40 to +70°C	Ericsson GR-63-CORE
Operating Humidity (RH)	8–90% non-condensing	IEC 60068-2-78
Altitude	Up to 3,048 m (10,000 ft) operational	IEC 60068-2-13
Vibration (operating)	0.5 Grms, 5–500 Hz, IEC 68-2-64	NEBS Level 3
Shock (non-operating)	30 G, 11 ms half-sine	IEC 60068-2-27
Seismic	Zone 4 per GR-63-CORE NEBS	Ericsson GR-63-CORE
ESD Protection	HBM 2 kV, CDM 500 V at all external pins	ANSI/ESDA/JEDEC JS-001
EMI Radiated	FCC Part 15 Class A; EN 55032 Class A	CISPR 22 Class A
EMI Conducted	EN 55032 Class A limits	CISPR 22 Class A
Safety	UL 62368-1; IEC 62368-1; CB Scheme	IEC 62368-1:2023
Laser Safety	Class 1M per IEC 60825-1; interlocked ports	IEC 60825-1:2014+A1:2021
RoHS Compliance	RoHS 3 (EU 2015/863)	EU Directive
REACH Compliance	SVHC concentration < 0.1% by weight	ECHA REACH Reg.

TABLE 16 – Environmental and Compliance

Reliability, MTBF, and Qualification

Component	MTBF Target	FIT Rate (target)	AFR	Methodology
2025 CPO Module	≥ 500,000 hours	≤ 2,000 FIT	≤ 1.75% @ 25°C	Telcordia GR-468 + APC-REL-001
2026+ Module	≥ 750,000 hours	≤ 1,333 FIT	≤ 1.17%	HALT/HASS
Laser sub-assembly (DFB)	≥ 1,000,000 hours	≤ 1,000 FIT	≤ 0.88%	Telcordia GR-468 Laser HALT
Switch Tile (silicon photonic)	≥ 2,000,000 hours	≤ 500 FIT	≤ 0.44%	IPC-9592B + AEC-Q100
Scale-Out IROS Switch	≥ 200,000 hours	≤ 5,000 FIT	≤ 4.38%	Ericsson SR-332 Issue 4

TABLE 17 – MTBF Target and Methodology

NOTE: FIT = Failures In Time (failures per 10⁹ device-hours). AFR = Annualized Failure Rate.

Qualification Test Flow

All CPO modules may complete the following qualification sequence before production shipment:

- Design Validation Testing (DVT): Full compliance test suite. Minimum 3 lots × 30 samples per lot.
- HALT (Highly Accelerated Life Test): Temperature cycling –40 to +95°C at 20°C/min; vibration 5 Grms broadband; combined stress 8+ hours. Purpose: identify latent design weaknesses.
- HASS (Highly Accelerated Stress Screen): Production screening per lot. Temperature cycling 3× from Tmin to Tmax at 10°C/min; vibration 2 Grms.
- Temperature Life Test: 2,000 hours at Tcase = Tmax; 77 units per stress cell (per Telcordia GR-468 sample plan). Accept criterion: 0 fails (C=0).
- Damp Heat: 85°C / 85% RH, 1,000 hours. Accept: < 1 dB IL change; no condensation on PIC.
- Thermal Shock: –40 to +85°C, 100 cycles. Accept: < 0.5 dB IL change.
- Mechanical Vibration: per Section 9.4 profile, 3 axes × 2 hours. Accept: no functional failure; < 0.3 dB IL change.
- ESD: HBM 2 kV on all external pins. CDM 500 V. Accept: no permanent failure.
- EMC: Radiated emissions test in accredited chamber; FCC Class A, EN 55032. Accept: ≥ 6 dB margin at all measured frequencies.
- Fiber Pull Test: MPO connector pull-out force ≥ 50 N without IL > 0.5 dB. 200 mating cycle endurance with < 0.3 dB IL change.
- Long-Term Wavelength Stability: 1,000 hours at 70°C; wavelength drift < 0.05 nm.
- Power Cycling: 10,000 on/off cycles; accept < 1 dB IL change; no CDR lock failures.

For laser lifetime prediction, the Arrhenius acceleration model should be used per Telcordia GR-468:

Acceleration Factor (AF) = $\exp[(E_a/k) \times (1/T_{\text{use}} - 1/T_{\text{stress}})]$ Where $E_a = 0.7$ eV for InP DFB lasers; $E_a = 0.5$ eV for Si photonic passive components; $k = \text{Boltzmann constant } (8.617 \times 10^{-5} \text{ eV/K})$. Minimum acceleration factor ≥ 100× to reduce test duration below 12 months.

2.7 Technology Roadmap and Evolution

The below APC informative outlook roadmap is structured around three phases, each defined by a dominant CPO class and per-lane electrical rate.

Phase	Years	Per-Lane Rate	Key Technology Milestone
Foundation	2028	112G PAM4	Near-package PIC; EML laser array; CWDM4 O-band
Scale	2030	224G PAM4	On-package MCM PIC; 16-ch C-band WDM; optical switch tiles
Integration	2029–2032	448 Gbps PAM4	Flip-chip SiPh; integrated laser on Si; microfluidic cooling
Convergence	2032–2035	896G+	Monolithic photonic-electronic integration; on-die waveguides

TABLE 18 – APC CPO Informative Outlook Roadmap

Silicon Photonics Technology Milestones

Milestone	Target Year	Description	Impact
GeSi EAM 112 GBd	2027	Germanium-Silicon electro-absorption modulator at 112GBd	Eliminates InP modulator need for O-band CPO-1
On-chip EDFA (rare-earth doped Si)	2028	Erbium-doped waveguide amplifier on SiPh platform	Removes discrete EDFA from chassis design
Hybrid III-V/Si laser at wafer scale	2028–2030	Bonded InP gain chip on SiPh foundry wafer	Integrated laser source for CPO-2/3; eliminates external coupling losses
Optical Phase Array (OPA) switch	2029–2031	Silicon OPA-based non-mechanical optical switching	Enables sub- μ s photonic circuit switching for scale-up fabric
Optical Memory (photonic RAM)	2031–2033	Photonic RAM for in-fabric buffering (research)	Eliminates O-E-O conversion in switching fabric
Monolithic SiGe BiCMOS PIC	2032–2035	Photonic and electronic on single BiCMOS wafer	CPO-4 class; eliminates all bump connections

TABLE 19 – Silicon Photonics Technology Milestones

CPO Informative Outlook Roadmap (Scale-Across)

Coherent optics integrated in CPO modules for scale-out applications:

Year	Modulation	Symbol Rate	Net Bitrate/ λ	Reach	DSP Node
2027	DP-16QAM	96GBd	400G	≤ 80 km	7 nm FinFET
2028–2029	DP-32QAM	130GBd	800G	≤ 80 km	5 nm FinFET
2029–2031	DP-64QAM	130GBd	1.2T	≤ 40 km	3 nm GAA
2031–2033	DP-64QAM	200GBd	1.6T	≤ 40 km	2 nm GAA
2033–2035	DP-256QAM	200GBd	2.4T+	≤ 20 km	1.4 nm (target)

TABLE 20 – Coherent Optics Integrated in CPO Modules

2.8 Multi-Vendor Interoperability Requirements

CPO modules from different vendors should interoperate within the same chassis if:

- Both modules are APC Compliant with the same generation.
- Wavelength plans are identical or the CWDM4 O-band baseline is used.
- CMIS register map is fully implemented.
- FEC mode negotiation completes successfully via CMIS FEC negotiation registers.

Backward and Forward Compatibility

Scenario	Support Required	Notes
Current Generation CPO module in CPO Generation 2 module socket	Mandatory mechanical adapter available	Reduced bandwidth; optical coupling loss ≤ 2 dB adapter penalty
112G module in 224G socket	Required (downgrade mode)	CMIS speed negotiation to Gen 1 parameters; laser wavelength unchanged
224G module in 112G host	Required (downgrade mode)	Module operates at 112G on each lane; excess lanes held in P3 state

TABLE 21 – Backward and Forward Compatibility

2.9 Nanosecond-Level OCS

Scale-up extends a single coherent logical domain across multiple chassis. The defining characteristic of scale-up is shared memory or tightly coupled collective operations at sub-microsecond latency across chassis boundaries. For 2 to 4 chassis in a rack, a dedicated Intra-Rack Optical Switch (IROS) module may be installed. The IROS module is a 1OU unit supporting up to 128 x 128 single-mode ports. Latency through the IROS switch should be ≤ 100 ns end-to-end (including fiber propagation for ≤ 3 m patch cables). The IROS may use photonic circuit switching (microsecond reconfiguration), silicon photonic switch matrices, or passive star coupler configurations depending on traffic pattern. For transparent-switch deployments the OCS-facing outputs may require PMF FAU.

Supporting this architecture requires extending the control model from configuration-time OCS control to a real-time fabric control plane. Endpoints and optical switching operation must operate from a common time base and coordinate transmission opportunities at the nanosecond scale. The standards package therefore needs to address both fixed and demand-driven scheduling; demand, grant and credit exchange; time and cycle synchronization; switch-state programming and confirmation; guard intervals and transition behavior; and recovery from missed slots, link faults or errors.

This real-time fabric control plane should remain distinct from conventional management, telemetry and SDN interfaces, which continue to provide configuration, monitoring and orchestration at much slower timescales. A dedicated out-of-band optical control channel or wavelength may also be used so that control traffic is independent of the instantaneous state of the data fabric.

A corresponding optical and system-interface framework is required so that fast OCS implementations remain interoperable without prescribing a particular switching technology. This should establish end-to-end optical-loss and power budgets for transparent multi-stage paths, wavelength and polarization requirements, control-channel separation, switching-time and settling-time definitions, synchronization accuracy, and common test points for measuring fabric performance. At the same time, the standard should preserve architectural choices - such as lane rate, number of wavelengths, fiber count, switch radix, number of switching stages and endpoint implementation - where different combinations can provide equivalent fabric capacity and behavior.

The objective is not to standardize a particular fast OCS implementation, but to define the control-plane, optical-interface and behavioral framework that allows fast, transparent OCS to operate as a first-class component of an open AI scale-up fabric.

2.9 Test and Measurement

Required Test Equipment

Test Category	Equipment	Key Spec
Optical BER Testing	Optical BER tester with PAM4 capability	$\geq 224\text{GBd}$; SJ/RJ injection; error insertion
Optical Power Measurement	Calibrated optical power meter	Wavelength range 1260–1625 nm; ± 0.1 dB accuracy; NIST traceable
Optical Spectrum Analysis	OSA or OFT	RBW ≤ 0.1 nm (optical domain); ≥ 70 dB dynamic range
Electrical Signal Integrity	Oscilloscope with coherent sampling	≥ 256 GHz analog BW; TDR/TDT; PAM4 eye analysis software
Vector Network Analyzer	4-port VNA	100 MHz – 200 GHz; S-parameter measurement to 0.02 dB accuracy
Wavelength Meter	Fizeau-based wavemeter	0.0001 nm absolute accuracy; 1260–1630 nm range
Noise Figure Analyzer	RIN measurement system	–160 to –130 dB/Hz range; 10 MHz – 20 GHz RF spectrum
Time-Interval Analyzer	Phase noise / jitter analyzer	< 50 fs rms floor; 1 Hz – 100 MHz offset range
Thermal Imaging	MWIR/LWIR thermal camera	$\pm 1^\circ\text{C}$ accuracy; 25 μm spatial resolution for PIC die imaging

TABLE 22 – Required Test Equipment

All CPO modules may pass the following production acceptance tests before shipping:

- Optical power test: All Tx channels within ± 1 dB of target launch power.
- Wavelength test: All Tx channels within ± 0.05 nm of nominal wavelength.
- BER test: 10 seconds $\times$ all Rx lanes at rated optical input; zero uncorrected errors post-KP4-FEC.
- CMIS register readback: Verify all APC-defined registers read/write correctly.
- Power consumption test: Module total power within $\pm 5\%$ of specification at full load.
- Thermal sense calibration: Laser temperature sense calibrated $\pm 1^\circ\text{C}$ vs. external reference.
- Mating cycle test (sample, 5% of production): 5 insertions/removals per module; IL change < 0.1 dB.

At chassis assembly completion, the following system-level tests should be required:

- End-to-end fabric BER: All ASIC pairs; inject 10^{14} bits; zero uncorrected FEC errors.
- Latency characterization: Measure fabric latency at 10%, 50%, 90%, and 100% load; compare to Section 3 targets.
- Thermal steady-state: Run chassis at full optical load for 1 hour; verify no thermal alarms.
- Power cycling: 10 chassis power cycles; verify all CPO modules achieve P0 state within 2 seconds.
- Hot-plug test: Remove and re-insert each CPO module while remaining ASICs carry live traffic; verify < 10 ms service interruption.
- EMC pre-smay: In-house smay at 1 m; verify ≥ 6 dB margin to FCC Class A limits before submission to accredited lab.

3. CPO System Enabled Overview

The CPO system still performs the same structure as a pluggable-based system, routing data between individual servers and toward the overall data center. A CPO based system integrates or co-packages the ASIC and the optical engine onto a single substrate. Solutions needs to co-package the ASIC with the correct number of optical engines to reach the target bandwidth. The shift of optical transceivers away from the front plate creates new opportunities for revised layouts and modifications across the entire CPO based system. Capabilities for electrical, optical, and thermal cooling, and packaging designs will depend on the design choices for various solutions.




Figure 7 – By Packaging both the Optical Engine and the Switching Chip, CPO Enables Optical Signals to be Converted Directly within the Chip, Eliminating the Need for High Speed PCB Signal Routing

The electrical interconnect between the host ASIC and the (optical) engine/module may be selected to minimize power and cost at the same time achieving the required performance metrics and mechanical requirements. The electrical interface performance is not only a function of the insertion loss of the host electrical interconnect, it is also dependent on the reflections and crosstalk of the interconnect. A variety of options have been proposed as bandwidth requirements continue to increase. A high level summary is provided in the below table.

Characteristics	Pluggable	Optics	Co-Packaged Optics
Electrical Interface Reach	LR, VSR	VSR, XSR, XSR+	XSR, XSR+, Linear
Electrical Interface Width	Serial	Serial	Serial or Wide
Bandwidth Density (Electrical)	Serial PCB phy & bump pitch dependent	OBO module size dependent	High. Package level bump pitch dependent
Target Link Bandwidths	All link bandwidths	Link BW: 1, 2, 4, 8+ Lanes: 25 to 100G	Link BW: 1, 2, 4, 8+ Lanes: 50, 100+G

TABLE 23 – Electrical Capabilities of Pluggable, Optics, and CPO Designs




Figure 8 – 51.2T Switch CPO Layout Example

3.1. Connectivity Options between Compute and Front panel

For a 51.2T switch, most of the ports on the front panel will be optical connections. A brief discussion of connections and definitions is needed to define the optical path from the compute engine to the front panel. The term PMD (Physical Medium Dependent) from the IEEE is used to incorporate the electro-optical package mounted to the system circuit board without the fiber necessary to take the signal to and from the card edge. In a co-packaged optical design, the PMD is the integrated optical engine (OE). Some form of optical cabling is required to reach the Medium Dependent Interface (MDI), which is defined in the same manner as IEEE ethernet standards. The end-user of the system simply sees an MDI interface at the faceplate regardless of connectivity type.

Characteristics	Pluggable	Co-Packaged Optics
Internal or External Laser	Internal	Internal or External
Signaling Technology	56G and 112G PAM4	56G and 112G PAM4 (112G focus)
Signal Integrity	Reflected in power and BER rows.	Reflected in power and BER rows.
Retimers	Potentially (Dependent on rate, distance)	No
Bandwidth Density (Optical)	Determined by transceiver module width	High. PIC size dependent
Optical Interface Reach	Wide range: cost dependent options	Intra-data center
Optical Lane Formats	SR4/SR8, DR4, FR4	SR4/SR8, DR4, FR4

TABLE 24 – 51.2T Switch CPO Layout Example

Both the connection method and connection type are design considerations that impact overall system requirements for signal, heat, and cabling density. Pigtailed cabling is directly attached to the PMD, while connectorized cabling uses a patch cord with connectors at both ends. In either case, connections at the panel are needed for data transmit/receive along with optional external laser sources. A design showing the basic elements of a switch ASIC with co-packaged optics is shown above.

Connector types vary dependent on implementation and the selection of PMD. For a 51.2T switch, multiple multi-fiber connectors are expected to be required due to single channel bandwidth constraints. For external laser sources which would plug into the face plate, current form factors are not standardized. This informative outlook document follows the proposal where either OSFP form factor dimensions or QSFP-DD form factor dimensions would be used. These designs have different fiber densities and area requirements, which define the overall concentration and layout of connectors on the faceplate layout.

Fiber Connector	Fibers per Connector	Connectors per 1RU	Fibers per 1RU	Reference Standard
Duplex LC	2	144	144	IEC 61754-20:2012 [7]
CS	2	320	320	TIA-604-19 [8]
MDC	2	432	432	IEC 61754-37 [9]
SN	2	432	432	IEC 61754-36 [10]
MPO-12	12	80	960	TIA-604-5 [11] IEC 61754-7-1 [12]
MPO-16	16	80	1280	TIA-604-18 [13] IEC 61754-7-4 [14]
MPO-24	24	80	1920	TIA-604-5 [11] IEC 61754-7-2 [16]
MPO-32	32	80	2560	TIA-604-18 [13] IEC 61754-7-3 [17]
MXC-32	32	104	3328	
MMC-16	16	216	3456	
SN-MT16	16	216	3456	

Table 25 – Optical Fiber Connectors and 1RU Connector Density

The size of the overall connection area on the rack faceplate will vary depending on the overall design and spacing requirements of each receptacle. Blind mate connectors, used to protect the end-user from eye-sight damage, may also require designs with receptacles significantly larger than the minimum receptacle size requirements.

3.2. Optical Engine Overview

This section describes the optical engines and electrical interfaces within Co-Packaged Optics (CPO) systems containing four (4) to thirty-two (32) Optical Engines per co-packaged module or system. A CPO system consists of one to four host ASIC co-packaged on a shared substrate or interposer with a population of Optical Engines arrayed around its periphery or, in higher-density variants, on multiple sides and tiers of a 2.5D/3D assembly. Each Optical Engine connects to the host ASIC over a short, dense electrical interface and to the external optical fiber plant through a fiber attach unit and external connector.

Optical Engine Definition and Boundaries

- Electrical boundary (host side): electrical interface at the Optical Engine's substrate-facing pads or bumps.
- Optical boundary (line side): the external fiber interface at the Optical Engine's connector or pigtail termination point.
- Mechanical boundary: the Optical Engine footprint and attachment interface to the host substrate.
- Thermal boundary: the top (or bottom) thermal interface surface of the Optical Engine lid or case.
- Control boundary: the sideband management interface exposed at the substrate connection.

Scalability Requirements: 4 to 32 Optical Engines

The below table summarizes a 4-32 Optical Engine scaling range. Actual channel counts and per-lane rates depend on the generation in deployment.

Density	Optical Engines	Typical Use Case	Illustrative Aggregate I/O (Gen-2 baseline)
Low	4–6	Compact switch / accelerator module	12.8 – 19.2T
Mid	8–16	Top-of-rack / leaf switch, multi-die accelerator package	25.6 – 51.2T
High	20–32	Spine / core switch	– 102.4T

Table 26 – System Classes

3.3. Optical Engine Architecture

The Optical Engine performs bidirectional conversion between the electrical domain (high-speed SerDes lanes to/from the host ASIC) and the optical domain (modulated optical channels to/from the external fiber plant). At minimum, a conformant Optical Engine comprises the following functional blocks:

- One or more Photonic Integrated Circuits (PICs) implementing optical modulation, photodetection, and passive routing (multiplexing/demultiplexing, splitting, polarization handling).
- One or more Electronic Integrated Circuits (EICs) implementing driver amplifiers for the modulators and transimpedance/limiting amplifiers for the photodetectors.
- A laser source, which may be integrated on the PIC or supplied externally via external laser source (ELS) module.
- A Fiber Attach Unit (FAU) providing mechanical and optical coupling between the PIC and the external fiber interface.
- Local power conditioning (regulation/decoupling) sufficient to meet supply quality requirements.
- A management controller or passive register interface exposing the control and telemetry model.

Optical Engine Functional Block Diagram




Figure 9 - The Host ASIC Connects to the Optical Engine over a Common Electrical Interface CEI

The host ASIC connects to the Optical Engine over a common electrical interface (CEI) carrying transmit and receive SerDes lanes. On the transmit path, the Electrical Integrated Circuit (EIC) driver stage modulates the PIC's optical modulators, which encode data onto the laser source's optical carrier. On the receive path, the PIC's photodetectors convert incoming optical signals to current, which the EIC's TIA/limiting amplifier stage converts to a voltage-domain SerDes signal returned to the host ASIC. The FAU couples the PIC's optical ports to the external fiber interface. A sideband management path connects the EIC or an associated microcontroller to the host substrate management bus.

The E/O path should support, per channel, modulation formats including NRZ and PAM4 and provide per-channel bias and modulation amplitude control accessible through the management interface to enable factory calibration and field re-calibration over device lifetime and temperature range. The O/E path should meet the receiver sensitivity and overload requirements across the full specified case temperature range of the Optical Engine and provide loss-of-signal (LOS) and received optical power monitoring per channel.

There are two architectural options for the laser source: Integrated Laser: the laser source is integrated within the Optical Engine boundary, and External Laser Source (ELS), an unmodulated continuous-wave light is supplied to the Optical Engine from a separate, replaceable ELS module via fiber interface, decoupling laser field-replaceability from the co-packaged module. An Optical Engine design should declare its laser source option and host substrates should provide the supply-fiber routing and ELS bulkhead interface. Systems implementing ELS should implement per-engine optical safety interlocks consistent with the applicable laser safety classification, disabling supply light upon detection of a fiber fault or open connector.

This section is process and technology neutral with respect to PIC platform (silicon photonics, silicon nitride, indium phosphide, thin-film lithium niobate) and EIC process node. This neutrality is a deliberate design goal to preserve multivendor sourcing options during which the dominant PIC platform may evolve.

3.4. Optical Interface Informative Outlook

Each Optical Engine supports a number of optical channels per direction (transmit and receive) as shown below. All channels of a given Optical Engine instance should operate at the same nominal line rate.

Two wavelength-plan options are defined to accommodate differing PIC integration strategies:

- WDM, O-band: a four- or eight-wavelength CWDM-style grid in the O-band.
- SDM, single wavelength: a single nominal wavelength per fiber, with channel parallelism achieved through spatial division rather than wavelength multiplexing.

3.5. Fiber Attach and External Connector Interface

The Fiber Attach Unit should terminate at a detachable external optical connector at the host substrate edge or designated fiber breakout region, using one of the connector families listed the below table.

Connector Family	Fiber Count	Typical Application
MPO-16	16	Aggregated multi-engine fiber shuffle to front panel or external patch field
SN / MDC (single or duplex)	1–2	Per-channel breakout for SDM configurations
Custom high-density array	≥32	Maximum-density 32-engine system fiber shuffle

Table 27 – Fiber Attach and External Connector Interface

The external connector interface should be field-accessible without disassembly of the co-packaged module, to permit fiber plant servicing without requiring Optical Engine replacement. Fiber routing between the Optical Engine and the external connector (the “fiber shuffle”) should meet a minimum bend radius and routing channel requirements to avoid excess macrobend loss across the full system thermal and mechanical operating range. Optical Engines using polarization-sensitive coupling should specify polarization-dependent loss (PDL) not exceeding 0.5 dB per facet and should document any required polarization-maintaining fiber handling in the fiber attach interface document.

3.6 Electrical Interface Specifications

The Optical Engine's host-facing electrical interface should comply with an OIF CEI extra-short-reach (XSR) interface class appropriate to the channel rate. The XSR class is selected specifically because the co-packaged trace length between the host ASIC and Optical Engine is short relative to front-panel pluggable reach classes, enabling lower per-bit energy and relaxed equalization requirements relative to long-reach electrical interfaces.

A conformant Optical Engine should implement CEI-224G-XSR compliant transmit and receive electrical interfaces for all lanes, meeting the transmitter and receiver eye mask, jitter, and return loss requirements of [OIF-CEI] at the package-level test points. The number of electrical lanes per Optical Engine should equal the number of optical channels per direction, a 1:1 electrical-lane-to-optical-channel mapping should be supported as the baseline mapping; gearbox or lane-aggregation mappings may be supported as an optional declared capability.

Lane numbering should be consistent across all Optical Engine vendors to ensure substrate routing compatibility regardless of which conformant Optical Engine is populated at a given Engine Site. Lane polarity (true/complement) and lane order should be fixed; any required lane reordering or polarity inversion to accommodate substrate routing should be resolved in the host ASIC's configuration, not by deviation from the Optical Engine pad map.

Where an Optical Engine requires a reference clock for internal timing, it should accept a reference clock distributed from the host substrate at a frequency and jitter mask declared in its identification record, selected from the common set (156.25 MHz, 312.5 MHz) to maximize host clock-tree reuse across vendors. Optical Engines implementing fully self-clocked operation may declare reference-clock-free operation in their identification record; host substrates should not be required to route a reference clock to Engine Sites populated exclusively by such engines. Each Optical Engine should expose a low-speed sideband management interface using I2C, electrically compatible with the voltage and timing parameters. Each Engine Site may provide a unique hardware address or strap mechanism such that up to a number of Optical Engines on a common management bus segment can be individually addressed without address collision. The below table summarizes the low-speed interface signal set required at the Engine Site.

Signal	Direction (host→OE / OE→host)	Function
SCL / SDA	Bidirectional	Register access, identification, DDM telemetry, firmware update
RESET_N	Host → OE	Hardware reset
INT_N / ALERT_N	OE → Host	Asynchronous interrupt for fault/alarm conditions
MODSEL_N / ADDR	Host → OE	Engine Site address strap, supporting up to 32 unique addresses
LASER_DISABLE	Host → OE	Optical safety interlock control

Table 28 – Low-speed Interface Signal Set Required at the Engine Site.

3.7 Mechanical and Form Factor Specifications

To enable multivendor sourcing of Optical Engines into a common substrate design, this specification defines a small set of standard footprint classes rather than a single fixed footprint, recognizing that channel count and integration choices affect the practical minimum footprint.

Footprint Class	Nominal Envelope (L × W × H)	Typical Channel Count	Notes
Small	12.5 × 8.5 × 3.0 mm	8	Single-row PIC, integrated laser
Medium	17.5 × 12.5 × 3.5 mm	8–16	Dual-row or 2-PIC tile
Large	25.0 × 17.5 × 4.0 mm	16	Additional area for supply-fiber coupling

Table 29 – Channel Count and Integration Choices

Host substrates intended for general multivendor sourcing should implement Engine Sites compatible with footprint class Medium as the baseline; support for small and large is optional and should be declared in the substrate identification record. As system designs scale from 4+ Optical Engines, the substrate must accommodate proportionally denser or more numerous Engine Sites while preserving uniform electrical reach, thermal access, and fiber shuffle routing. This CPO forward looking outlook document defines a unit-cell placement model rather than mandating a single fixed array geometry, since 4+ engine will reasonably differ in floorplan.

The below table illustrates representative floorplan envelopes consistent with these placement rules across the supported scaling range.

Engine Count (N)	Representative Array	Approx. Substrate Optical Edge Length
4	2 × 2 or single-edge linear	≥ 76 mm (linear) or compact quad
8	Single-edge linear or 2-edge split (4+4)	≥ 152 mm (single edge) or 2 × 76 mm
16	2-edge split (8+8) or 4-edge (4 per edge)	≥ 152 mm per populated edge
32	4-edge (8 per edge) or dual-tier 2.5D assembly	≥ 152 mm per populated edge, or vertical tiering

Table 30 – Representative Floorplan Envelopes

Each Engine Site should provide:

- A defined electrical attachment interface (e.g., flip-chip bump array, land grid, or socketed interposer connector).
- A defined mechanical keep-out volume, sufficient to clear neighboring engines, fiber shuffle channels, and the thermal interface structure.
- A defined fiber egress channel routing from the Engine Site's FAU location to the external connector field.

Because Optical Engines are co-packaged rather than front-panel pluggable, field serviceability differs fundamentally from conventional transceiver modules. This CPO forward looking outlook document defines three serviceability tiers to allow systems to make an explicit, declared trade-off between field repairability and integration density/performance.

The below table defines the normative thermal operating envelope per Optical Engine. These values are defined at the Optical Engine's thermal interface surface and are independent of the cooling technology used to remove heat, provided the case temperature limit is met.

Parameter	Value	Notes
Case temperature operating range	0°C to 75°C	Measured at thermal interface surface per
Typical power dissipation per engine	4.5 – 6.5 W	EIC + PIC + integrated laser, all channels active
Typical power dissipation per engine	3.0 – 4.5 W	EIC + PIC only; laser dissipation external to module
Maximum allowable thermal interface resistance	≤ 0.4 °C/W	Aggregate across PIC and EIC thermal paths
Laser-region maximum local temperature	70°C	Where TEC-assisted local control is not implemented

Table 31 – Normative Thermal Operating Envelope

Because system designs in scope range from 4+ Optical Engines, the aggregate optical-subsystem thermal load scales accordingly and must be accounted for in host module cooling design, independent of host ASIC thermal load. Host module thermal designs should provision cooling capacity for the aggregate optical thermal load at the maximum populated Engine Site count supported by that module design, in addition to host ASIC and other co-packaged thermal loads. Where a shared cold plate spans multiple Optical Engines and the host ASIC, the cold plate design should accommodate the differing nominal thermal interface heights of populated footprint classes through compliant TIM or local shim features, without constraining the host ASIC to a single Optical Engine footprint class. Each Optical Engine should implement an autonomous thermal protection response if its internal temperature exceeds a vendor-declared protection threshold, independent of host intervention.

3.8 Power Delivery

The below table defines the normative power rail set at the Optical Engine's substrate-facing interface. Vendors may implement additional internally regulated rails derived from these inputs but should not require additional rails to be routed from the host substrate beyond those listed, to preserve substrate design commonality across vendors.

Rail	Nominal Voltage	Tolerance	Typical Use
VDD_HS	0.80 V	±3%	EIC high-speed driver/TIA core
VDD_IO	1.80 V	±5%	Low-speed I/O, management controller
VDD_LSR	3.3 V or vendor-declared	±5%	Integrated laser bias supply
VSS	0 V (reference)	—	Common return

Table 32 – The Normative Power Rail Set

Each declared rail should meet the voltage tolerance measured at the Optical Engine's power input pads, inclusive of substrate IR drop, across the full load range from idle to all-channels-active operation. Optical Engines should specify maximum allowable power supply ripple and noise on each rail, which should not be tighter than 1% peak-to-peak on VDD_HS without documented justification. Optical Engines should tolerate power rail application in any order and should not require a host-enforced sequencing order, unless a specific sequencing requirement is declared in the identification record; where sequencing is required, the maximum allowable skew between rail assertions should be declared. Optical Engines should limit inrush current at power-on to a vendor-declared maximum, and host power delivery network designs should account for the worst-case simultaneous inrush of all populated Engine Sites at module power-on.

The below table illustrates representative aggregate power delivery provisioning across the supported scaling range. These figures inform host power delivery network (PDN) design margin since electrical power and dissipated heat are not always numerically identical once optical output power is accounted for.

Engine Count	Typical Aggregate Optical-Subsystem Power	Typical Aggregate Optical-Subsystem Power
4	20 – 28 W	13 – 19 W
8	40 – 56 W	26 – 38 W
16	80 – 112 W	52 – 76 W
32	160 – 224 W	104 – 152 W

Table 33 – Power Budget

Host PDN designs supporting field upgrade from a lower to a higher Engine Site population within a single product generation should provision for the higher population's power and thermal budget at initial design, even if shipped initially in a lower-population configuration.

3.9 Control, Management, and Telemetry

Each Optical Engine should implement a register-based management interface, structured as a paged register map adapted from the [CMIS] model to the co-packaged context. The register map provides identification, configuration, control, and diagnostic monitoring functions through a common addressing scheme regardless of vendor, enabling host management software to interoperate with any conformant Optical Engine without vendor-specific drivers for baseline functions. The register map should include a lower memory page containing, at minimum: vendor identification, part number, serial number, declared generation, declared footprint class, declared laser source option, declared wavelength plan, and declared channel count and line rate.

Each Optical Engine should report the following diagnostic parameters through the register map, at a minimum, updated at an interval not to exceed 100 ms:

- Per-channel transmit optical power or supply optical power.
- Per-channel received optical power.
- Per-channel laser bias current or supply-light presence indicator.
- Module-level temperature.
- Supply voltage readback for each rail.
- Per-channel loss-of-signal (LOS) and loss-of-lock (LOL) status flags.
- Aggregated alarm and warning flags with configurable high/low thresholds for each monitored parameter.

Where an Optical Engine implements an onboard microcontroller with field-updatable firmware, it should support an in-field update mechanism meeting the following requirements: Firmware update should be performed over the existing low-speed management interface without requiring removal of the Optical Engine from the host substrate; The Optical Engine should support a dual-bank (active/inactive) firmware image model, such that an update failure or interruption does not leave the device unable to boot a previously functional image; Firmware images should be authenticated prior to activation, to prevent unauthorized or corrupted firmware from being placed into the active bank.

3.10 Reliability, Availability, and Serviceability

Optical Engine vendors may publish a reliability qualification report, including accelerated life test results and Failure in Time (FIT) rate derivation methodology, as part of their compliance declaration. Because a CPO system may contain several Optical Engines, the failure of a single engine represents a smaller fraction of total system bandwidth than in lower-radix designs, but also implies a larger number of potential failure points. This CPO Informative Outlook and Recommended Practices document does not mandate system-level design practices to exploit this scaling characteristic. Systems should support graceful bandwidth degradation upon single Optical Engine failure, disabling only the affected engine's channels rather than the full module, where the host ASIC and switching/routing layer support per-engine channel isolation.

Optical compliance is evaluated at three defined test points to separate PIC-level, fiber-attach-level, and system-level contributions to link budget. Electrical compliance is evaluated at the Optical Engine's substrate-facing electrical interface, using the test fixtures and de-embedding methodology defined in [OIF-CEI] for the declared XSR class. Electrical compliance testing should include, at minimum: transmitter eye mask compliance, transmitter jitter (random and deterministic components), receiver sensitivity/tolerance (stressed eye), and return loss, for both transmit and receive directions.

This CPO Informative Outlook and Recommended Practices document anticipates that the dominant per-lane electrical and optical line rate will advance roughly twofold per generation across the planning window, consistent with historical Ethernet and electrical SerDes scaling trends, while the Optical Engine's external mechanical and management interfaces are designed to remain stable across at least two generations to protect substrate and host ASIC investment.

Future years are documented here as informative roadmap guidance to support long-lead substrate and system planning; they will be formalized as normative requirements in future major revisions of this forward looking document as the underlying electrical and optical technologies mature and are validated. A host substrate designed to the FC-M footprint and pad map of this revision should be capable of accepting a current Optical Engine of compatible footprint, with the understanding that aggregate bandwidth will be limited to the current rate, provided the host ASIC SerDes supports the corresponding CEI-112G-XSR fallback rate. Mixed-generation population within a single substrate may be supported at vendor discretion but is not required for conformance. Systems supporting mixed population should document per-site rate negotiation behavior in their system-level documentation.

4. Mid – Board Connection

For a CPO switch ASIC, dealing with fiber routing is inevitable from fibers from optical engines to the front panel inside the switch chassis. The optical engines are positioned closely around the ASIC to minimize the distance of the electrical path, thereby maximizing the electrical performance and lowering the overall ASIC power consumption. However, design selection complicates the fiber routing inside the system. Even though electrical paths may be minimized between OE and ASIC, the distance from each OE to the faceplate would vary. In addition, the fibers would exit the optical engine in four different directions in the case where the optical engines are positioned as illustrated below.

For an electrically separable socketed optical engine, fibers in the figure below design would need to be long enough to extend to the faceplate, requiring various lengths of fiber to route each optical engine. For a pigtailed optical engine design as shown below, there may be 4 to 8 variants of pigtail length required. Multiple lengths increase the total components in the BOM of the switch system, and the OE manufacturers would be required to prepare the pigtails in various lengths as illustrated below.




Figure 10 – Fiber Pigtail Lengths Required for Interior Connections

Pigtail length variation brings concerns not only about design but also the risk of fiber damage during installation. The fiber pigtails should be routed in tight spaces surrounded by other components. This may impact the overall production yield, especially in the case that the OEs are soldered in the switch package and are not replaceable, where damage to a single fiber in the pigtail will compromise the entire CPO switch package. This yield risk may contribute negatively to the overall costs of the switch production. In the case where pigtails may be connectorized, connector positions and size could constrain options for heating or fiber density. Pigtail lengths, if varied, may document detailing fiber layouts to prevent unwanted breakage.

4.1 Mid-Board Connector as a Solution to Fiber Routing Concerns

Implementing a mid-board/on-board optical interconnect solution may solve these concerns. By adding a mid-board connector between the OE and the faceplate, and providing jumper cords of various lengths, the OE pigtail length may be reduced to just one design. This simplifies the manufacturing for the OE vendors, while also reducing the risk of damage to the OE and the attached pigtail. Furthermore, the shorter pigtails simplify the installation of the ASIC OE to the faceplate. Although the jumper cords require routing around the system, fiber breakage do not compromise the CPO switch ASIC subassembly. Jumper fibers may be easily replaced with another jumper. Yield-affecting risks are shifted to jumper cables from the more expensive OE and ASIC components. The change to connectorized cables helps to improve the cost impact of the CPO switch ASIC. There are multiple solutions for mid-board/on-board connectivity. Options for mid-board connectors are shown below.




Figure 11 – Example of Mid-board Connector Solutions




Figure 12 – Connection Points between Transmit and Receive PMDs

The use of IEC 61753-1 Grade B rated connectors may be considered to achieve a lower loss consistently. It is one of the best performance grades for the fiber optic connector insertion loss. This performance grade is providing a manufacturer independent interoperability, and the connector loss is guaranteed by a test with randomly mated pairs among different manufacturers and manufacturing batches. Based on this random mating specification for MPO connector (IEC 61755-3-3 or IEC 61755-3-31), Grade B MPO 12F single mode fiber connectors should be 0.25dB or better at 97.3% of each fiber channel.





Figure 13 – Future 204.8T / 1024 x 200G (-DR)



Mid-board connector option




Figure 16 – Mid-board Connector Option

Splice / fan out / shuffle box option




Figure 17 – Splice / Fan out / Shuffle Box Option

5. Laser Sources

This section describes the performance specifications of external and integrated laser sources for CPO applications. This informational outlook and recommended practices document defines requirements, interfaces, test methods, and a phased performance roadmap for laser sources integrated within CPO assemblies. The document outlines approaches for applications from single-chassis implementations to hyperscale datacenter fabrics aggregating more than one million accelerators.

Laser sources occupy the highest-criticality subsystem position in CPO design. Their wavelength stability, relative intensity noise (RIN), phase noise, thermal budget, reliability, and polarization management determine the potential system-level optical signal-to-noise ratio (OSNR), bit-error rate (BER) floor, and the viability of advanced modulation formats.

5.1 Informational Outlook and Recommended Practices Scope

- Laser source types considered for CPO integration, including semiconductor Fabry-Pérot (FP), distributed feedback (DFB), vertical-cavity surface-emitting (VCSEL), optical frequency comb (OFC), and silicon photonics micro-ring resonator (MRR) coupled external-cavity laser (ECL) sources.
- Electro-optic and thermo-optic performance requirements for each modulation and multiplexing scheme.
- Electrical interface overview.
- Optical interface overview.
- Thermal design constraints and thermal interface overview.
- Reliability and qualification requirements compliant with Telcordia GR-468, IEC 62343.
- Deployment architectures for scale-up (NVLink/Fabric domain) and scale-out (rack, pod, datacenter, and multi-datacenter) configurations.
- Manufacturing readiness levels (MRL) and supply chain considerations.
- Power, energy-efficiency, and thermal management requirements.
- Security, provenance, and export-control considerations.

5.2 Laser Source Technologies Overview

In a CPO implementation, the laser source provides the continuous-wave (CW) optical carrier that is subsequently modulated by silicon photonic Mach-Zehnder or micro-ring resonator modulators within the photonic engine. The laser is the single highest-value and highest-risk component in the optical link: it determines OSNR budget, modulation format viability, wavelength plan, energy-per-bit, and — critically — field reliability over the targeted 10-to-15-year operating lifetime of the host platform.

Unlike pluggable optics where laser failure causes a replaceable module to be swapped, a laser integrated within a CPO package that fails in the field requires board-level replacement of the entire switching or compute subsystem. This requirement elevates laser reliability specifications by one to two orders of magnitude relative to conventional transceiver-grade sources, driving unique qualification regimes and redundancy architecture requirements.




Figure 18 – Pluggable Optics Laser Failure (Module needs to be swapped) and a CPO Package that Fails (needs total replacement).

Indium phosphide (InP)-based distributed feedback lasers are the current industry baseline for single-mode, wavelength-stabilized CW sources. DFB lasers operate on a single longitudinal mode defined by a periodic grating etched near the active region. They deliver excellent side-mode suppression ratios (SMSR >45 dB), moderate linewidths (typically 100 kHz–10 MHz), and output powers of 10–40 mW from single-element emitters.

For CPO, DFB arrays — typically 4, 8, or 16 wavelengths on a single InP bar — reduce packaging complexity and cost. Butt-coupling to silicon photonic waveguides using micro-lensed fiber or direct edge-coupling with spot-size converters achieves coupling efficiency of 75–90% (–0.5 to –1.2 dB) in state-of-the-art implementations. Polarization maintaining output (PER >25 dB) is achievable with epitaxial strain engineering.

External cavity lasers extend the effective cavity length by coupling the InP gain chip to a low-loss silicon nitride (Si₃N₄) or silica waveguide resonator. This reduces linewidth dramatically — to below 100 Hz in laboratory demonstrations and below 1 kHz in product-grade implementations — making ECL sources essential for coherent modulation formats beyond QPSK.

An optical frequency comb generates tens to hundreds of equally spaced, phase-coherent spectral lines from a single source die. For WDM CPO, a single OFC source may replace an array of discrete DFBs, substantially reducing component count, wavelength management complexity, and the number of individual reliability-critical elements. OFC sources impose demanding requirements on comb line power uniformity (<3 dB variation across ≥ 64 lines), repetition rate stability (phase-locked to a frequency reference), and the ability to thermally tune the comb offset frequency across a full free spectral range.

Vertical-cavity surface-emitting lasers are inherently low-cost, surface-emitting devices compatible with 2D array integration and wafer-level testing. GaAs-based VCSELs at 850 nm deliver multimode output suitable for multimode fiber (MMF) short-reach links and are the dominant technology in current 400G QSFP-DD SR4 and 800G QSFP-DD SR8 modules. For CPO/NPO, solutions based on VCSEL arrays over MMF remain an option for intra-chassis and intra-racks to 100 m. Single-mode 1310 nm InP VCSELs, still at early TRL for high-volume production, may enable low-cost single-mode CPO for reaches up to 2 km.

The aspirational goal of CMOS-compatible laser integration — avoiding III-V hybrid bonding entirely — drives research into rare-earth-doped waveguide amplifiers, germanium-tin (GeSn) band-engineered lasers, and color-center emitters in silicon. While none of these approaches offers a viable datacenter-grade product within the 2025–2030 horizon, GeSn lasers operating at room temperature with continuous-wave lasing above -15 °C were demonstrated in 2023, suggesting potential for CPO-compatible integration in the 2032–2035 timeframe.

Laser Source Electrical Specifications

Parameter	Min	Typical	Max	Unit	Notes
Threshold Current (I_{th})	8	12	18	mA	At 25 °C case temperature
Operating Bias Current	40	60	100	mA	Above threshold; APC-controlled
Slope Efficiency	0.20	0.28	—	W/A	At operating bias
Wall-Plug Efficiency (WPE)	28	33	—	%	At $T_{case} = 45$ °C
Modulation Bandwidth (–3 dB)	—	25	—	GHz	Small-signal, at operating bias
Loop Bandwidth	—	1	—	MHz	Monitor PD feedback loop
Drive Voltage (single-ended)	—	1.8 / 3.3	—	V	Compatible with 7/5 nm CMOS
Monitor PD Responsivity	0.8	1.0	—	A/W	Integrated back-facet PD
Laser Supply Voltage	3.0	3.3	3.6	V	Regulated on-package LDO
Transient Current Overshoot	—	—	±5%	% of I_{bias}	During transients

Table 34 – Laser Source Electrical Parameters

Laser Source Optical Parameters

Parameter	Min	Typical	Max	Unit	Notes
Wavelength Plan (CWDM4)	1271	—	1331	nm	Channels: 1271, 1291, 1311, 1331 nm
Wavelength Plan (LR4 DWDM)	1295.56	—	1309.14	nm	Channels per IEEE 802.3bs Annex 88D
Wavelength Accuracy (AWC locked)	–0.05	0	+0.05	nm	Over full operating temperature range
Output Power at Waveguide Coupler	3.0	5.0	8.0	dBm	After coupling optics, into SiPh waveguide
SMSR (Single-Mode Suppression Ratio)	45	50	—	dB	At all operating conditions
RIN (10 MHz – 20 GHz)	—	–155	–150	dBc/Hz	Integrated over band
Linewidth (–20 dB)	—	1	5	MHz	Lorentzian fit
Polarization Extinction Ratio	25	30	—	dB	TE output, into SiPh waveguide
Back-Reflection Tolerance	—	—	–30	dB	No degradation; isolator required if higher
Optical Return Loss (ORL) at Input	40	—	—	dB	Isolator minimum requirement
Fiber-to-Waveguide Coupling Loss	—	0.8	1.5	dB	Lensed fiber / spot-size converter

Table 35 – Laser Source Optical Parameters

Thermal and Environmental Parameters

Parameter	Min	Typ	Max	Unit
Operating Case Temperature (Liquid Cooled)	45	55	65	°C
Storage Temperature	−40	—	+85	°C
Junction Temperature (T _{jmax})	—	70	85	°C
TEC Setpoint Temperature	25	30	35	°C
TEC COP (Coefficient of Performance)	0.5	0.7	—	—
TEC Maximum Delta-T	—	—	30	K
Thermal Resistance Junction-to-Coolant	—	4.0	5.0	K/W
Maximum Laser Package Power	—	1.5	2.5	W
Maximum TEC Power	—	0.8	1.5	W
Humidity (non-condensing)	5	—	85	% RH

Table 36 – Thermal and Environmental Parameters

Laser Reliability Requirements

Reliability Metric	Requirement	Test Method
MTBF (Median, 55 °C case)	$\geq 1.0 \times 10^7$ hours	ALT at 85 °C, 150% Ibias; Arrhenius extrapolation
FIT Rate (Failures In Time)	≤ 100 FIT at 55 °C	Statistical sample ≥ 77 units, 1000 hr at 85 °C
Catastrophic Optical Degradation (COD)	Zero occurrences	Tested at 200% rated power for 500 hr
Wavelength Drift over Life	$< \pm 0.1$ nm	Measured post-1000 hr ALT, extrapolated to 10^7 hr
Output Power Degradation at EoL	< 1 dB	Measured post-ALT; APC compensation assumed
Humidity / Condensation (85/85)	No degradation	1000 hr at 85 °C, 85% RH per JESD22-A101
Thermal Cycling (−40 to +85 °C)	No degradation	500 cycles per IEC 60068-2-14
Electrostatic Discharge (ESD)	± 500 V HBM	Per JESD22-A114

Table 37 – Laser Reliability Requirements

5.3 Laser Source Roadmap 2027-2033

CPO laser sources transition from discrete externally packaged laser modules to chiplet-compatible co-integrated assemblies. The introduction of higher-order modulation formats, initially DP-QPSK and DP-16QAM, enable a fourfold increase in spectral efficiency relative to Phase 1 PAM4 IM-DD, supporting per-lane data rates of 400G and aggregate package bandwidths approaching 100T for 128-port switch ASICs manufactured on the 2 nm process node.

The introduction of DWDM channel plans with 50 GHz (0.4 nm) channel spacing across the C-band (1530–1565 nm) and L-band (1565–1610 nm), require laser wavelength accuracy of ± 25 pm and integrated isolators with >35 dB return loss. The ECL architecture becomes the specification baseline for coherent lanes, while InP DFB arrays continue for intensity-modulated short-reach applications.

Parameter	Year 2027	Year 2030	Change Driver
Per-lane Data Rate	100G PAM4	400G DP-16QAM	Coherent DSP chiplet integration
Aggregate Bandwidth / Package	25.6T	102.4T	128-port @ 800G
Laser Linewidth (coherent)	1–5 MHz	<10 kHz	ECL + Si ₃ N ₄ resonator
Wavelength Channels per Source	4 (CWDM)	80 (C-band DWDM)	OFC or DFB comb array
Wavelength Accuracy	± 50 pm	± 25 pm	Improved AWC, micro-heater trim
Wall-Plug Efficiency	30–33%	38–42%	Improved InP epitaxy, reduced threshold
Coupling Architecture	Lensed fiber array	Photonic wire bond / flip-chip	Heterogeneous integration maturity
Isolator Integration	Discrete module	Monolithically bonded	Magneto-optic thin-film maturity
RIN Requirement	-150 dBc/Hz	-160 dBc/Hz	Higher OSNR requirement for DP-16QAM
TEC Thermal Resistance	4–5 K/W	2.5–3.5 K/W	Micro-channel liquid cooling
Laser Source MRL	8–9	9–10	Volume production established

Table 38 – Key Parameter Roadmap 2027-2030

Coherent External Cavity Laser Specifications

For coherent lanes, this forward-looking outlook document mandates ECL sources meeting the following additional requirements:

- **Linewidth:** <10 kHz (Lorentzian, measured by delayed self-heterodyne method per IEC 61282-9).
- **Frequency Noise PSD:** <1000 Hz²/Hz at 1 MHz offset from carrier.
- **Relative Intensity Noise:** <-160 dBc/Hz from 10 MHz to 50 GHz.
- **Wavelength Tuning Range:** $\geq$ Full C-band (35 nm) with <1 mW tuning power via micro-heater.
- **Tuning Speed:** Full C-band in <100 ms; 1 GHz in <100 μ s for fast channel provisioning.
- **Polarization Maintaining Output:** PER ≥ 30 dB into PM-waveguide coupler.
- **Coherent Optical Output Power:** $\geq +10$ dBm into Si waveguide (accounting for coupling loss).

5.4 Heterogeneous Integration Approaches

MTP deposits pre-fabricated InP QD gain coupons (typical size 100×500 μ m) onto designated regions of a processed SiPh wafer using a PDMS stamp with <1 μ m placement accuracy. Evanescent coupling between the bonded InP layer and underlying Si or Si₃N₄ waveguide is achieved through wafer-level adiabatic taper design. MTP throughput of $>10,000$ coupons/hour at the die-attach step is targeted for high-volume manufacturing.

Plasma-activated direct bonding of InP epitaxial structures (grown on GaAs or InP handle wafers) to SiO₂-terminated SiPh wafers achieves sub-nanometer bond interface uniformity at temperatures compatible with processed back-end-of-line metal layers (<300 °C when using surface-activated bonding — SAB). Post-bond handle-wafer removal via selective wet etch exposes the InP gain layer for lithographic patterning.

Photonic wire bonds achieve insertion loss <0.5 dB per bond, enabling multi-die photonic systems with $>95\%$ coupling efficiency. Automated PWB tools operating at >500 bonds/hour support production volumes. Hermetic encapsulation of PWBs with index-matched epoxy withstands 1000 thermal cycles (-40 to $+85$ °C).

Laser Source Outlook

Photonic interconnect requirements of exascale AI datacenter fabrics projected to aggregate more than ten million accelerators within a single logical training domain by 2035. At this scale, per-package optical bandwidth should exceed 800T, driving wavelength plans spanning C, L, S, and O bands simultaneously with agile routing, and laser sources delivering sub-100-Hz linewidth for DP-256QAM modulation at 1.6T per lane.

Fully integrated laser-on-chip (LoC) sources fabricated in commercial foundry silicon photonics (SIP) processes, leveraging rare-earth (Er^{3+} , Yb^{3+}) doped waveguide gain sections pumped by on-chip 976 nm pump lasers, GeSn lasers operating at room temperature, or color-center (G-center) emitters in silicon-on-insulator. While these technologies remain at technology readiness level (TRL) 3–4 in 2025, the foundry and heterogeneous integration infrastructure is expected to mature them to production readiness within this time frame.

Areas of Development

GeSn Room-Temperature CW Lasing: $\text{Ge}_{1-x}\text{Sn}_x$ alloys with $x > 0.10$ have demonstrated lasing at room temperature in pulsed mode. Achieving threshold current densities below 1 kA/cm² in CW mode at 85 °C junction temperature requires strain engineering and resonator quality-factor improvements.

Silicon G-Center Emitters: Carbon-interstitial (G-center) emitters in implanted silicon provide photoluminescence at 1278 nm (O-band). Integration with Si photonic cavities for laser amplification is an active MIT/imec research program.

Rare-Earth Doped Si₃N₄ Waveguide Lasers: Er^{3+} -doped Si_3N_4 waveguides pumped at 1480 nm have achieved net optical gain. On-chip integration of InP pump lasers with Er^{3+} -doped resonators may provide narrow-linewidth C-band lasing without epitaxial III-V transfer.

Photonic Neuromorphic Control: Machine-learning-based laser control loops using on-chip neuromorphic processors may adaptively compensate for aging-induced wavelength drift and RIN degradation without explicit model updates, extending effective laser lifetime beyond mechanical reliability limits.

Centralized Laser Source (CLS) Architecture

In the CLS architecture, a single high-power laser source bank provides CW light to all ASIC photonic engines within the chassis via a passive optical distribution network. The CLS bank is mounted in a dedicated laser tray within the chassis, thermally isolated from the compute cards. Light is distributed via PM fiber ribbons routed through the chassis backplane, with optical power splitters (1:N, N up to 16) delivering equal-power inputs to each ASIC card.

Advantages of CLS: simplified redundancy (N+1 or N+2 laser sparing at the tray level), centralized wavelength management, reduced per-card laser component count, and easier replacement serviceability. Disadvantages: optical splitting loss (up to 12 dB for 1:16 passive split), requiring higher initial laser output power (+18 to +22 dBm), and long optical fiber routing paths with associated insertion and polarization management challenges.

Distributed Laser Source (DLS) Architecture

In the DLS architecture, each ASIC card or group of 2–4 ASIC cards carries its own LSU, typically mounted at the card edge or in a dedicated thermal zone adjacent to the photonic engine. Each LSU provides CW light only to the photonic engines on its local card or sub-group. DLS eliminates the splitting loss penalty and reduces optical distribution complexity but multiplies the number of laser reliability-critical elements within the chassis.

Users may use DLS for implementations, transitioning to CLS when OFC sources with sufficient per-line output power (>+10 dBm per line), enabling a single comb source to serve the full chassis with manageable splitting loss.

Chassis-Level Optical Budget

Link Element	Year 2027 Loss (dB)	Year 2028 – 2030 Loss (dB)	Year 2031- 2033 Loss (dB)
Laser coupling to PM fiber (per engine)	0.8–1.5	0.5–0.8	<0.3
PM fiber routing in chassis (2 m)	0.1	0.1	0.1
Fiber-to-waveguide coupler at PIC	0.8–1.5	0.5–0.8	<0.3
On-chip waveguide routing (5 cm)	0.2	0.15	0.1
Power splitter (1:4)	6.5	6.5	6.2
Modulator insertion loss	3.0	2.5	2.0
Total insertion loss estimate	11.4–12.8	10.25–11.0	8.7–9.0
Required laser output power (into fiber)	+14 to +15 dBm	+13 to +14 dBm	+12 to +13 dBm

Table 39 – Chassis-Level Optical Performance

5.5 Redundancy and Fault Management

Given the non-replaceable nature of laser sources in fully integrated CPO designs, single-chassis systems should implement the following redundancy features:

- **Wavelength-Diverse Redundancy:** For DFB array sources, a spare wavelength channel is held in reserve. On detection of single-channel failure by the monitor PD array, the APC controller activates the spare channel and updates the AWC setpoint within <100 ms, maintaining link availability.
- **Power-Boosted Degraded Mode:** On detection of >2 dB output power degradation on any laser channel, the APC loop increases bias current to the remaining healthy lasers by up to 20%, maintaining aggregate optical power delivery to the fabric at the cost of reduced per-laser lifetime margin.
- **Chassis-Level Laser Health Monitoring:** Each LSU reports laser junction temperature, bias current, monitor PD current, wavelength locker error signal, and TEC power via a chassis management bus (CMB) I²C/SPI interface. The chassis BMC logs these telemetry streams at 1-second intervals and generates predictive failure alerts when trends exceed trained thresholds.
- **Hot-Swap Laser Source Cartridge (Phase 1–2):** In CLS and hybrid DLS architectures, the LSC mechanical interface permits tool-free removal and replacement while the chassis remains powered, with optical blind-mate connector break occurring <50 ms before electrical disconnect

Scale-Up Datacenter Deployments

Scale-up AI clusters extend accelerator integration beyond a single chassis to encompass groups of 2–64 chassis interconnected by an ultra-high-bandwidth switched fabric operating within a single administrative domain. Within a scale-up domain, the optical fabric should deliver sustained all-to-all bandwidth approaching the aggregate ASIC memory bandwidth. For 64-chassis × 16-GPU systems at 2T GPU memory bandwidth, this implies aggregate fabric bandwidth of 2,048T across the domain, a requirement only optical interconnects may satisfy at acceptable energy cost.

Each chassis connects its CPO-enabled switching ASIC to ToR optical switches via short-reach intra-rack optical links (≤2 m, Phase 1 DFB CWDM, 400G per fiber pair). The ToR switch aggregates 32–64 chassis-to-switch links and provides any-to-any connectivity within the rack. Multiple ToR switches are interconnected via spine switches at the pod level.

Future scale-up domains benefit from optical circuit switching (OCS) integrated at the ToR level, enabling bandwidth-on-demand provisioning across the fabric with sub-millisecond reconfiguration times. CPO laser sources should support OCS wavelength-assignment requests, including rapid wavelength tuning (C-band sweep in <100 ms for ECL sources) and stable output during tuning transients (optical power variation <0.5 dB during tuning).

At scale-up dimensions, the total number of CPO laser source units in a 64-chassis × 16-GPU system is 1,024 LSUs (one per ASIC card in DLS architecture). At a FIT rate of 100 per LSU, the system-level FIT is 102,400 — implying a system MTBF of ~9,800 hours, or approximately 13 months. This is unacceptable for a production AI training cluster.

Training frontier AI models through 2030 and beyond will require compute clusters exceeding one million accelerators, operating as a single logical training job over periods of weeks to months. At this scale, the datacenter fabric transitions from a switched Ethernet or InfiniBand topology to a purpose-built optical fabric organized as a multi-stage Clos or dragonfly topology, with multiple levels of hierarchical aggregation spanning individual racks, rows, pods, halls, and datacenters.

For 1,000,000 GPUs with 3.2T bidirectional bandwidth per GPU, the aggregate fabric capacity requirement is 1.6 exabytes per second (EB/s), a number that makes optical interconnect not merely advantageous but the only viable physics solution. At this scale, CPO laser sources represent both a massive technical challenge (>62.5 million individual laser elements at 1 LSU per GPU in DLS architecture) and an equally massive supply chain and reliability engineering challenge.

5.6 Laser Source Requirements at Scale-Out and Scale Across Levels

At the pod-spine and hall-core levels, CPO-based short-reach optics give way to coherent DWDM transceivers (initially pluggable 400ZR/800ZR, transitioning to CPO-integrated coherent engines). The laser sources for these links should meet:

- **Linewidth:** <100 Hz (distributed Bragg reflector ECL with Si₃N₄ resonator).
- **Wavelength Plan:** ITU-T G.694.1 flex-grid, 6.25 GHz granularity across C+L bands (160+ channels).
- **Output Power:** +13 to +17 dBm before EDFA or SOA boost, to achieve required OSNR after fiber spans.
- **Frequency Stability:** ±1 GHz over operating lifetime (maintained by GPS-disciplined optical frequency reference).
- **Modulation Format Support:** DP-64QAM (Phase 2), DP-256QAM with probabilistic constellation shaping.

APC defines the Datacenter-scale Laser Health Monitoring (DLHM) framework, comprising:

- **Edge Collection Layer:** Each chassis BMC collects laser telemetry at 1-second granularity and forwards summarized 1-minute statistics to the rack-level aggregator via CMB/Redfish API.
- **Rack/Row Aggregation Layer:** Row-level servers aggregate telemetry from 32 racks, applying real-time anomaly scoring using pre-trained LSTM or Transformer models to detect incipient laser failures.
- **Cluster Intelligence Layer:** A dedicated AI inference cluster (irony intentional) runs fleet-wide laser health models, correlating failure patterns across vendors, chassis generations, and environmental variables (ambient temperature, humidity, vibration). It generates work orders for predictive replacements 30–60 days in advance.
- **Open Telemetry Standard:** DLHM telemetry should be exported in OpenTelemetry format via OTLP protocol to enable vendor-agnostic monitoring integration. Minimum telemetry payload per LSU: bias current (mA), monitor PD current (μA), junction temperature (°C), TEC current (mA), wavelength locker error (pm), output power (dBm), and LSU uptime (seconds).

5.7 CPO Laser Integration Thermal Management

Thermal management is the most critical design challenge for CPO laser source integration. Semiconductor lasers exhibit strong output-wavelength and threshold-current dependence on junction temperature: wavelength shifts approximately $+0.1 \text{ nm}/^{\circ}\text{C}$ for InP DFB lasers, while threshold current roughly doubles for every 50°C increase. In a co-packaged environment, the laser should maintain junction temperatures below 85°C while neighboring compute and memory dice dissipate up to $1,000 \text{ W}/\text{cm}^2$ in the same package.

Thermal Ground Plane with Differential Cooling - The laser source unit is mounted on a dedicated thermal sub-island within the package, separated from the compute die by a thermal break (low-conductivity adhesive or air gap). A localized TEC maintains the laser sub-island at a setpoint temperature (typically $25\text{--}35^{\circ}\text{C}$) independent of the compute die temperature.

Micro-Channel Liquid Cooling – Integration of micro-channel liquid cooling directly into the laser interposer, with coolant channels machined at $50\text{--}100 \mu\text{m}$ pitch in silicon or copper. Liquid coolant (deionized water or fluorinated fluid) flows at $0.1\text{--}1 \text{ L}/\text{min}$ per laser cluster, removing heat with a thermal resistance of $<1.5 \text{ K}/\text{W}$ — substantially lower than TEC-based cooling.

Future Immersion Cooling Compatibility - CPO implementations targeting hyperscale deployments with two-phase immersion cooling should ensure all laser source materials, adhesives, encapsulants, and optical coatings are compatible with the selected dielectric fluid.

TEC Parameter	Year 2027	Year 2030	Year 2033	Unit
Maximum Delta-T ($T_h - T_c$)	30	25	20	K
COP at operating point	0.5–0.7	0.6–0.8	0.7–0.9	—
TEC supply voltage	3.3–5	2.5–3.3	1.8–2.5	V
TEC maximum current	3	2	1.5	A
Temperature control loop bandwidth	1	5	10	Hz
Setpoint accuracy	± 0.5	± 0.2	± 0.1	$^{\circ}\text{C}$
Thermal resistance (junction-to-TEC cold side)	<3	<2	<1	K/W
TEC MTBF (operational)	$>5 \times 10^6$	$>1 \times 10^7$	$>1 \times 10^7$	hr

Table 40 – Thermal Electric Cooling Specification and Control

All CPO laser source integrators should provide thermal simulation models conforming to the JEDEC JEP30 compact thermal model (CTM) format, enabling system integrators to validate chassis-level thermal budgets. Models should be validated against measured data at three operating points (minimum, typical, maximum power) with <5% error in junction-temperature prediction.

At datacenter scale, laser source energy efficiency is measured in femtojoules per bit (fJ/bit) transmitted. Advanced Photonics Coalition defines the Laser Energy Efficiency Ratio (LEER) as the ratio of total optical energy delivered to the photonic engine modulator to the total electrical energy consumed by the laser source subsystem (including laser drive, TEC, Automatic Power Control electronics, and wavelength locker).

Efficiency Improvement Pathways

- **Epitaxial Engineering:** Reducing threshold current density in DFB lasers from 2 kA/cm² (current commercial) to <0.5 kA/cm² (quantum-dot, 2030 target) directly reduces wasted power at operating point.
- **TEC Elimination:** Replacing thermoelectric cooling with a thermal laser designs (quantum-dot lasers with intrinsically flat gain spectrum over 0–85 °C) or micro-channel liquid cooling achieving sufficiently low thermal resistance eliminates the 20–40% overhead TEC power penalty.
- **Optical Amplification Offset:** On-chip SOA amplification of higher splitting-loss architectures enables using lower-power lasers (less waste heat, better WPE at lower output power operating points) at the cost of SOA power, a net energy benefit when SOA efficiency >40%.
- **Integrated Photonic Power Management:** Dynamic per-channel laser power control (reducing optical power during low-traffic periods or under reduced modulation constellation) via the DLHM system may reduce fleet-average laser power by 15–25%.
- **PUE Impact of Laser Sources:** A 1-million-GPU datacenter operating at PUE 1.15 with 400 MW total power (typical 2030 projection) allocates approximately 5 MW to laser sources (1.25% of IT load). Improving LEER from 0.30 to 0.55 across the 10-year horizon reduces laser source power consumption proportionally, saving approximately 2.1 MW per facility, equivalent to 18,400 MWh/year and approximately \$920,000/year at \$0.05/kWh industrial electricity pricing.

5.8 Laser Reliability, Qualification, and Lifetime

CPO laser sources may be qualified to a higher standard than traditional pluggable transceiver lasers because: (1) they are not field-replaceable in fully integrated implementations, (2) they should survive the full chassis lifetime of 7–10 years without scheduled replacement, and (3) a single laser failure in a densely integrated package may cause the entire compute card to fail, impacting training jobs running across thousands of interdependent nodes.

The Telcordia GR-468 SONET protocol with additional stresses specific to CPO environments:

Stress Condition	Test Temperature	Bias	Duration	Sample Size	Acceptance Criterion
High-Temperature Operating Life (HTOL)	85 °C case	1.5× I _{bias}	2,000 hr	≥77 units	FIT ≤ 100 at 55 °C
Damp Heat (85/85)	85 °C, 85% RH	0% (unpowered)	1,000 hr	≥22 units	No catastrophic failure
Temperature Cycling	–40 to +85 °C	0% (unpowered)	500 cycles	≥22 units	Pout change < 1 dB
Power Cycling	Operating T _{case}	0% ↔ 100% I _{bias}	10,000 cycles	≥22 units	Pout change < 1 dB
Vibration (random)	25 °C	100% I _{bias}	30 min/axis	≥10 units	No structural failure
Mechanical Shock	25 °C	0%	1500 g, 0.5 ms	≥10 units	No structural failure
Electrostatic Discharge	25 °C	0%	±500 V HBM	≥10 units	No threshold shift
CPO In-Package High-Temp	105 °C ambient	100% I _{bias}	500 hr	≥22 units	FIT ≤ 50 extrapolated

Table 41 – Specific Stress Conditions in CPO Environments and Acceptance Criterion

All laser source suppliers may provide a completed FMEA that should cover:

- Catastrophic optical degradation (COD) of front facet coating.
- Diffusion of metallization contacts into active region over life.
- TEC solder joint fatigue from thermal cycling.
- Wavelength locker etalon degradation (coating delamination).
- Monitor photodiode dark current increase over temperature and time.
- Optical fiber coupling degradation (adhesive aging, micro-crack in lensed fiber).
- Gain chip misalignment from thermal expansion mismatch in ECL configurations.
- ESD-induced active region damage during handling or in-system transients.

5.9 Manufacturing and Supply Chain

Total global InP wafer production capacity in 2027 is estimated at 50,000–80,000 4-inch equivalent wafers per year, supporting current global laser demand of approximately 500 million DFB die per year across telecommunications, datacom, and sensing applications. CPO adoption at the scale projected in this document will require a 10–100× expansion in InP manufacturing capacity by 2035. The transition path involves:

- **Wafer Size Transition:** Migration from 3" and 4" InP wafers to 6" wafers (currently in development at Coherent and UCSB) will increase die yield per wafer by 2.25–4×, partially offsetting capacity constraints.
- **Fab Automation and AI Quality Control:** Applying automated optical inspection (AOI), machine-vision die selection, and ML-based process control to InP fabrication lines reduces yield loss from 30–40% to <15% in state-of-the-art facilities.
- **New Fab Investment:** 2–4 new greenfield InP fabrication facilities will be required globally by 2030, with capital investment of \$500M–\$1B per fab. Government incentive programs may partially fund this expansion.
- **GaAs Substrate Migration for O-Band:** O-band lasers (for 1310 nm applications) may leverage the larger, more mature GaAs substrate supply chain, which is already at 6" volume production and offers lower material cost than InP.
- **Supplier Qualification and Diversity Requirements:** To mitigate supply chain concentration risk, CPO platform builders deploying more than 10,000 chassis should qualify at least two independent laser source suppliers per wavelength band. Users may use a primary/secondary supplier model with a minimum 60%/40% volume split, reviewed annually.
- **Supplier qualification may include:**
 - Full Level 2 qualification per Section 14.
 - On-site process audit by APC WG-Manufacturing auditors.
 - Demonstration of 90-day buffer inventory at geographic distribution points.
 - Export control certification and ITAR/EAR compliance documentation.
 - Business continuity plan (BCP) covering single-site disaster scenarios.

Manufacturing Readiness Level (MRL) Milestones

MRL	Definition	Phase Target
MRL 1	Capability to produce in laboratory, technology demonstrated in production-relevant environment	OFC / ECL chiplets (2025)
MRL 2	Capability to produce prototype components in production-relevant environment	ECL flip-chip (2026)
MRL 3	Capability to produce prototype system in production-relevant environment	OFC + SiPh (2027)
MRL 4	Capability to produce in pilot/early production environment	MTP InP-on-SiPh (2029)
MRL 5	Pilot line capability demonstrated, ready for low rate initial production (LRIP)	All Phase 2 sources (2030)
MRL 6	LRIP demonstrated, major issues resolved	Wafer-bonded laser (2032)
MRL 7	Full production demonstrated, lean manufacturing in place	All Phase 3 sources (2033)

Table 42 – Manufacturing Readiness Level (MRL) Milestones

5.10 Regulatory and Environmental Compliance

CPO laser sources should be classified per IEC 60825-1 Edition 3.0. In normal CPO operation with the laser fully enclosed within the package and chassis, the accessible emission limit (AEL) at any service point should not exceed Class 1M (intrinsically safe with the naked eye). During maintenance operations where the LSC is removed from a powered chassis, exposed fiber ends may present Class 3R or Class 3B hazards; appropriate interlocks and warning labels are required per IEC 60825-1 Section 5.

RoHS and REACH Compliance

All laser source materials and assembly processes should comply with EU Directive 2011/65/EU (RoHS III) and REACH Regulation (EC) No 1907/2006 as amended. Specifically:

- Lead-free solder for all non-hermetic package assembly.
- Absence of restricted substances (Cd, Hg, Cr⁶⁺, PBB, PBDE) above threshold concentrations.
- SVHC (Substances of Very High Concern) disclosure for all materials present above 0.1% w/w.
- 18.3 Environmental Impact and Circularity.

Advanced Photonics Coalition is committed to reducing the environmental footprint of CPO laser sources across their full lifecycle:

- **Energy Efficiency (Manufacturing):** Suppliers should report Scope 1 and Scope 2 GHG emissions per 1,000 LSUs manufactured and set annual reduction targets aligned with the Science Based Targets initiative (SBTi).
- **Material Efficiency:** Epitaxial material utilization should be tracked; suppliers are encouraged to implement InP substrate reclaim programs (substrate repolishing after die removal) to reduce rare material consumption.
- **End-of-Life (EoL) Management:** LSUs containing InP (a compound of indium, a critical material per EU Critical Raw Materials Act) should be collected at EoL and returned to InP recycling programs.
- **Conflict Mineral Compliance:** Suppliers should document absence of conflict minerals (tin, tantalum, tungsten, gold) sourced from conflict-affected areas per Dodd-Frank Section 1502 / EU Conflict Minerals Regulation.

5.11 Detailed Electrical Interface Specifications

The laser bias current should be supplied from an on-package low-dropout (LDO) voltage regulator with output voltage accuracy $\pm 1\%$ and noise spectral density $< 10 \text{ nV}/\sqrt{\text{Hz}}$ from 10 Hz to 1 GHz. The LDO output should be externally adjustable via an 8-bit DAC with resolution of 0.5 mA/LSB and monotonicity guaranteed. Maximum slew rate of the bias current during step changes: 1 mA/ μs . Faster slew rates cause transient power overshoots that may trigger optical power alarms.

The TEC driver should be a bidirectional H-bridge PWM driver with:

- **Supply voltage:** 5.0 V $\pm$ 5%.
- **Maximum output current:** 3 A .
- **PWM switching frequency:** 200 kHz (to avoid acoustic noise in TEC).
- **Current sense accuracy:** $\pm 2\%$ over full operating range.
- **Thermal shutdown:** at 125 °C driver die temperature.

At the Photonic Engine Optical Interface (PEOI) plane, the following dimensional tolerances apply:

Parameter	Si Waveguide	Si3N4 Waveguide	Unit
Waveguide width	450 $\pm$ 20	800 $\pm$ 30	nm
Waveguide height (etch depth)	220 $\pm$ 10	400 $\pm$ 15	nm
SSC output MFD	3.0 $\pm$ 0.3	4.0 $\pm$ 0.4	μm
SSC length	≥ 200	≥ 300	μm
Lateral alignment tolerance (3σ)	± 0.5	± 0.5	μm
Vertical alignment tolerance (3σ)	± 0.3	± 0.3	μm
Return loss at PEOI	≥ 40	≥ 45	dB

Table 43 – Dimensional Tolerances at the Photonic Engine Optical Interface (PEOI) Plane

6. Glass Interposer and Photonic Integrated Circuits

The Glass Interposer (GI) serves as the central substrate integrating optical waveguide routing, electrical redistribution layers (RDL), and micro-bump or hybrid bonding interfaces between the ASIC, PIC, and fiber array. This section defines the material, dimensional, optical, and electrical requirements for conformant GI substrates.

Material, Dimensional, Waveguide and Electrical Redistribution Layer Parameters

Parameter	Min	Typical	Max	Unit	Test Method
Glass Type	—	Borosilicate / AF32 / D263T	—	—	Manufacturer CoA
Refractive Index (core, 1310 nm)	1.444	1.4480	1.455	—	ISO 489
Refractive Index (clad, 1310 nm)	1.438	1.4440	1.450	—	ISO 489
CTE (0–300°C)	3.0	3.25	3.5	ppm/°C	ASTM E228
CTE mismatch to silicon	—	< 0.5	0.8	ppm/°C	APC-TM-GI-01
Young's Modulus	60	72	—	GPa	ASTM C623
Poisson's Ratio	0.18	0.22	0.26	—	ASTM C623
Birefringence	—	< 5	10	nm/cm	ISO 11455
Optical Transmission (1260–1625 nm)	99.0%	99.5%	—	%/cm	IEC 60584
Surface Roughness Ra (waveguide region)	—	< 0.5	1.0	nm	ISO 4287
Surface Roughness Ra (bonding region)	—	< 0.3	0.5	nm	ISO 4287
Dielectric Constant (10 GHz)	—	4.1	4.6	—	IEC 60250
Dissipation Factor (10 GHz)	—	0.002	0.005	—	IEC 60250
Breakdown Voltage	300–400	—	—	kV/μm	IEC 60243

Table 44 – Material Parameters

All dimensions are specified at 23°C ± 2°C, 45–55% RH, unless otherwise noted.

Parameter	Min	Typical	Max	Unit
Substrate Size (square)	—	50 × 50	100 × 100	mm
Substrate Size (rectangular, X)	30	—	120	mm
Substrate Size (rectangular, Y)	25	—	80	mm
Substrate Thickness	0.30	0.50	0.80	mm
Thickness Uniformity (TTV)	—	< 2	5	μm
Bow / Warp	—	< 10	20	μm
Edge Chamfer	0.1	0.2	0.3	mm
Dicing Street Width	50	100	—	μm
Positional Accuracy of Features	—	±0.5	±1.0	μm
Chip-to-chip alignment tolerance	—	±0.2	±0.5	μm

Table 45 – Dimensional Parameters

Embedded optical waveguides within the glass interposer should be fabricated by laser direct writing, ion exchange, or etching-and-deposition processes. The following parameters apply to single-mode waveguides operating in the O-band (1260–1360 nm) and C-band (1530–1565 nm).




Figure 19 – Glass Interposer and Photonic Integrated Circuits

Parameter	O-Band (1310 nm)	C-Band (1550 nm)	Unit
Propagation Loss	≤ 0.5	≤ 0.3	dB/cm
Waveguide Width (target)	3.0 ± 0.3	2.5 ± 0.3	μm
Waveguide Height (target)	2.5 ± 0.3	2.0 ± 0.3	μm
Mode Field Diameter (MFD)	8.5 ± 1.0	9.5 ± 1.0	μm
Numerical Aperture (NA)	0.13 ± 0.02	0.12 ± 0.02	—
Minimum Bend Radius (0.1 dB/turn)	3.0	5.0	mm
Bend Loss (at min. bend radius)	≤ 0.1	≤ 0.1	dB/90°
Polarization Dependent Loss (PDL)	≤ 0.2	≤ 0.15	dB
Polarization Mode Dispersion (PMD)	≤ 0.05	≤ 0.05	ps/ $\sqrt{\text{km}}$
Group Velocity Dispersion (GVD)	—	< 3	ps/nm-cm
Fiber-to-Waveguide Coupling Loss	≤ 0.8	≤ 0.6	dB
Waveguide-to-PIC Coupling Loss	≤ 0.5	≤ 0.4	dB
Crosstalk (adjacent waveguide, 10 μm gap)	≤ -40	≤ -40	dB
Return Loss at all interfaces	≥ 40	≥ 45	dB

Table 46 – Waveguide Parameters

The GI may include one or more metallic redistribution layers enabling electrical signal routing between co-integrated ICs. These layers must meet the following requirements for 112G PAM4 and 224G PAM4 signaling:

Parameter	Min	Typical	Max	Unit
Number of RDL Metal Layers	1	2	4	layers
Line Width / Space (L/S)	5/5	8/8	—	μm
Via Diameter (through-glass)	30	50	100	μm
Via Pitch	100	150	—	μm
Metal Conductivity (Cu)	—	5.8×10^7	—	S/m
Transmission Line Loss (50 Ω , 112G)	—	< 3.0	4.0	dB/cm
Transmission Line Loss (50 Ω , 224G)	—	< 5.5	7.0	dB/cm
Impedance Control ($\pm$ tolerance)	—	50 ± 5	—	Ω
Micro-bump pitch (ASIC-to-GI)	40	55	100	μm
Micro-bump pitch (PIC-to-GI)	30	40	80	μm
Micro-bump height	15	25	40	μm
Micro-bump shear strength	≥ 10	—	—	mN/bump

Table 47 – Electrical Redistribution Layer Parameters

6.1 Fiber Attachment and FAU Interface

The Glass Interposer should provide a defined mechanical reference surface for fiber array unit (FAU) attachment. Fiber arrays may be V-groove bonded, edge-coupled, or surface-coupled via grating couplers. The following parameters apply:

Parameter	Specification	Unit
FAU adhesive cure temperature	$\leq 85^{\circ}\text{C}$ (to protect co-integrated die)	$^{\circ}\text{C}$
Edge-facet surface roughness (Ra)	≤ 2.0 nm after polish	nm
Edge-facet angular tolerance	$\pm 0.3^{\circ}$ from design angle (typically 8°)	$^{\circ}$
Fiber array pitch (standard)	127 ± 0.5	μm
Fiber array pitch (dense)	80 ± 0.5	μm
Vertical alignment tolerance (V-groove)	± 0.5	μm
Horizontal alignment tolerance	± 0.5	μm
FAU insertion loss (per channel)	≤ 1.0 dB (incl. coupling + reflection)	dB
FAU return loss	≥ 40 dB (PC), ≥ 55 dB (APC-polished)	dB
Fiber type compatibility	ITU-T Standard G.652.D or G.657, HI-1060, PM-1550	—

Table 48 – Fiber Attachment and FAU Interface Parameters

6.2 Photonic Integrated Circuit (PIC)

This section defines the performance and interface requirements for PICs. PIC platforms addressed include Silicon Photonics (SiPh), Thin-Film Lithium Niobate (TFLN), Indium Phosphide (InP), and Gallium Arsenide (GaAs) where applicable.

PIC Platform Classification

Platform	Wavelength Range	Primary Application	Integration Density
Silicon Photonics (SiPh)	1260–1625 nm	Modulation, routing, detection	High
Thin-Film LiNbO_3 (TFLN)	800–2000 nm	High-speed modulation	Medium
Indium Phosphide (InP)	1260–1625 nm	Gain, detection, modulation	Medium
Silicon Nitride (SiN)	400–2350 nm	Low-loss routing, O-band PICs	Medium-High
GaAs (for O-band VCSELs)	800–1100 nm	Short-reach, low-power	Low-Medium

Table 49 – PIC Platform Classification

Modulators integrated within compliant PICs should meet the following performance requirements:

Parameter	112G PAM4 Min	224G PAM4 Min	Unit	Notes
Small-signal EO Bandwidth (S21)	≥ 70	≥ 110	GHz	-3 dB reference
Half-wave voltage V_{π} (DC)	—	≤ 3.5	V	TFLN: ≤ 2.0 V
$V_{\pi} \cdot L$ product (phase shifter)	—	≤ 2.0	V·cm	—
Extinction Ratio (ER)	≥ 7.0	≥ 6.5	dB	PAM4 levels 0→3
Insertion Loss (modulator)	≤ 5.0	≤ 4.5	dB	Includes coupling
Return Loss (all ports)	≥ 25	≥ 25	dB	—
EO chirp parameter α	—	< 0.2	—	$ \alpha $ for SiPh
Drive voltage (single-ended, pp)	—	≤ 2.0	V _{pp}	OIF CEI-224G compatible
Thermal drift (IL)	—	< 0.05	dB/°C	0 to 85°C
Thermal drift (phase, with TEC-less design)	—	$< \pi/10$	rad/°C	—

Table 50 – Electro-Optic Modulator Requirements

Integrated photodetectors should comply with the following performance requirements:

Parameter	Min	Typical	Max	Unit
Responsivity (1310 nm)	0.8	0.9	—	A/W
Responsivity (1550 nm)	0.9	1.0	—	A/W
Electrical Bandwidth (-3 dB)	70	—	—	GHz
Dark Current ($V_{\text{bias}} = -2\text{V}$)	—	< 1	5	nA
Junction Capacitance ($V_{\text{bias}} = -2\text{V}$)	—	15	25	fF
Saturation Power	10	—	—	dBm
Polarization Dependent Responsivity	—	< 0.3	0.5	dB
Return Loss (optical input)	≥ 30	≥ 35	—	dB
Operating Reverse Bias	1.5	2.0	4.0	V

Table 51 – Photodetector Requirements

Wavelength multiplexing components co-integrated within the PIC should meet:

Parameter	CWDM4 (25 nm ch. spacing)	LAN-WDM (4.5 nm spacing)	DWDM (100 GHz)	Unit
Channel Insertion Loss	≤ 1.8	≤ 2.0	≤ 2.5	dB
Adjacent Channel Crosstalk	≥ 25	≥ 30	≥ 35	dB
Non-adjacent Crosstalk	≥ 40	≥ 45	≥ 50	dB
Passband Width (0.5 dB)	≥ 10	≥ 3.0	≥ 0.5	nm / GHz
Center Wavelength Accuracy	± 2.0	± 0.5	± 0.08	nm / GHz
Thermal Wavelength Drift	< 0.1	< 0.015	< 0.001	nm/°C
Return Loss (all ports)	≥ 35	≥ 35	≥ 40	dB

Table 52 – Optical Multiplexer / Demultiplexer Parameters

PIC Die Physical and Integration Parameters

Parameter	Specification	Unit
Die Thickness (before thinning)	725 ± 25	μm
Die Thickness (after thinning)	50–200	μm
Thinned Die Bow	< 20	μm
Edge Surface Quality (post-dicing)	Chipping $< 3 \mu\text{m}$	—
Hybrid Bonding Pitch (Cu-Cu)	≥ 1.0 (target: 0.5)	μm
Hybrid Bonding Planarity	$< 1 \text{ nm RMS over } 1 \text{ cm}^2$	nm
Flip-Chip Bump Pitch	40–100	μm
PIC-to-GI Alignment Accuracy	± 0.3	μm
Die Shear Strength (per MIL-STD-883-2019)	$\geq 5 \text{ N per mm}^2 \text{ die area}$	N/mm ²
Maximum Die Temperature (operation)	125	°C
Maximum Die Temperature (solder reflow)	260	°C

Table 53 – PIC Die Physical and Integration Parameters

6.3 CPO Assembly and Co-Integration Standards

The following bonding technologies are recognized within the CPO assembly framework. Each is associated with specific performance tiers and process constraints:

Bonding Method	Pitch Range	Bandwidth/Channel	Process Temp.	APC Compliance Tier
Micro-bump (SnAg solder)	40–100 μm	Up to 112G	< 260°C reflow	Tier 1 (Baseline)
Copper Pillar + Underfill	30–60 μm	Up to 224G	< 250°C	Tier 2 (Advanced)
Hybrid Bonding (Cu-Cu/dielectric)	1–10 μm	Beyond 224G	200–400°C anneal	Tier 3 (Leading Edge)
Direct Glass Bonding (optical only)	—	Optical coupling	< 200°C	Tier 4 (Optical Only)

Table 54 – Bonding and Attachment Technologies

Underfill and Encapsulation

- Underfill should be no-flow or capillary type with CTE matched to substrate within 5 ppm/°C.
- Underfill Young's Modulus should be between 5 and 12 GPa after cure.
- Underfill should be RoHS-compliant and compatible with IEC 60068-2-14 thermal shock testing.
- void fraction within underfill should not exceed 2% by cross-sectional area.
- Encapsulant moisture absorption should not exceed 0.5% weight gain after 168h / 85°C / 85% RH (JEDEC JESD22-A110).

The total optical channel insertion loss budget for a CPO GI/PIC assembly from fiber array input to PIC receiver is specified in Table 68. All values are worst-case at end-of-life (EOL) and 70°C operating temperature.

Loss Element	O-Band (1310 nm)	C-Band (1550 nm)	Unit
FAU fiber attachment	0.5	0.4	dB
Waveguide propagation (5 cm)	0.25	0.15	dB
Waveguide bends (total)	0.10	0.10	dB
Waveguide splitter/coupler	0.30	0.25	dB
Waveguide-to-PIC coupling	0.5	0.4	dB
PIC on-chip routing	0.5	0.4	dB
Demultiplexer / filter	1.8	2.0	dB
Connector aging + end-of-life margin	0.5	0.5	dB

Table 55 – Optical Path and Alignment Budget

6.4 Thermal Management Specifications

Parameter	Min	Max	Unit
Ambient operating temperature	0	55	°C
Junction temperature (ASIC)	—	105	°C
GI substrate temperature	–5	95	°C
Storage temperature	–40	85	°C
Humidity (operating)	5	85	% RH non-condensing

Table 56 – Operating and Storage Temperature

Thermal Interface Requirements

- Thermal Interface Material (TIM) between lid and PIC should achieve thermal resistance $\leq 0.05 \text{ mm}^2\cdot\text{K/W}$.
- Thermal resistance junction-to-case (θ_{jc}) for PIC should be $\leq 0.5 \text{ }^\circ\text{C/W}$ per mm^2 of die area.
- Thermoelectric cooler (TEC) integration, if present, should maintain PIC junction temperature within $\pm 2^\circ\text{C}$ of set point across $0\text{--}55^\circ\text{C}$ ambient range.
- TEC power consumption should not exceed 15% of total CPO module power budget.
- GI substrate should be designed to tolerate $\Delta T \leq 50^\circ\text{C}$ between top (ASIC) and bottom (PCB) surfaces without exceeding allowable strain limits.

Power Density Units

Component	Typical Power Density	Maximum Allowable	Unit
ASIC (co-packaged)	0.5–1.5	2.0	W/mm ²
PIC (active region)	0.05–0.2	0.5	W/mm ²
TIA array (per channel)	50–100	150	mW/channel
Laser sub-assembly (if integrated)	100–500	800	mW/laser
Entire CPO module	—	350	W (absolute)

Table 57 – Power Density Units

6.5 Reliability and Environmental Qualification

MTBF and Reliability Targets

- Target MTBF for CPO module (Telcordia SR-332): $\geq 2,000,000$ hours at 40°C.
- Optical path reliability target: $< 0.1\%$ IL degradation per year under rated operating conditions.
- Glass substrate fracture probability: < 1 per million substrates per operational year (Weibull analysis required).
- PIC electro-optic degradation: < 0.2 dB IL increase over 7 years at 70°C junction temperature.

6.6 Test and Measurement Methodology

Optical Characterization

- **Insertion Loss and Return Loss:** measured per IEC 61300-3-4 using swept-wavelength optical vector analyzer (OVA) at 0.1 pm resolution across full operating wavelength range.
- **Modulator EO bandwidth:** measured via lightwave component analyzer (LCA) with calibrated optical reference receiver, -3 dB point determined from S21 magnitude.
- **Photodetector bandwidth:** measured via optical frequency-domain reflectometry (OFDR) or LCA, with correction for reference detector response.
- **Waveguide MFD:** measured by far-field scanning method per ISO/IEC 61280-1-4.
- **Crosstalk:** measured with DUT channel under test terminated and all adjacent channels driven at -3 dBm CW.
- **PDL:** measured by Stokes parameter rotation method using polarization synthesizer.

Electrical Characterization

- **S-parameter characterization:** Vector network analyzer (VNA) with 2.4 mm or 1.0 mm ports to 110 GHz (224G) or 67 GHz (112G) minimum.
- **Eye diagram / BER testing:** BERT with calibrated channel emulation per OIF CEI-224G normative annex.
- **Jitter characterization:** Real-time oscilloscope (> 160 GSa/s) with TDECQ and DJ/RJ decomposition.
- **Impedance profiling:** Time-domain reflectometry (TDR) with 10 ps rise time reference.

Physical and Structural Characterization

- **Waveguide geometry:** scanning electron microscopy (SEM) or atomic force microscopy (AFM) on cross-section samples — minimum 3 sites per substrate.
- **Glass surface roughness:** AFM with 1 μm × 1 μm scan area, 5 sites per substrate per lot.
- **Bonding void detection:** Scanning acoustic microscopy (SAM) at ≥ 200 MHz transducer frequency.
- **CTE measurement:** Thermomechanical analyzer (TMA) per ASTM E228 over –50°C to +200°C range.
- **Bow/warp:** Non-contact optical profilometer, 9-point measurement pattern per SEMI M49.

6.7 Interoperability and Interface Requirements

The CPO GI/PIC assembly should present the following electrical interface to the host ASIC through the interposer:

Interface Parameter	112G PAM4	224G PAM4	Unit
Signaling standard	OIF CEI-112G-XSR	OIF CEI-224G	—
Lane rate	112	224	G
Lanes per direction (per port)	4 or 8	4 or 8	—
Aggregate bandwidth (8 lanes × 2)	1.6 Tb/s	3.2 Tb/s	—
Differential impedance	85 ± 10	85 ± 10	Ω
Common-mode impedance	42.5 ± 5	42.5 ± 5	Ω
Trace length (GI, ASIC-to-TIA bump)	≤ 3	≤ 1.5	mm
Package insertion loss (TX)	≤ 3.0	≤ 4.5	dB at Nyquist

Table 58 – Host ASIC Interface

External Fiber Interface

- The GI/PIC assembly should support standard MPO-12 or MPO-16 connectors per IEC 61754-7.
- Multi-fiber interface insertion loss should not exceed 1.5 dB per connector pair (mated).
- Optical return loss at the external interface should be ≥ 40 dB (PC) or ≥ 60 dB.
- Fiber type should be compatible with SMF-28 (ITU-T G.652.D) or OM5 multimode (ITU-T G.651.1) as specified by the wavelength plan.
- The assembly should be optically compatible with coherent 400G-ZR (OIF) and 800G-ZR+ systems when operating in C-band DWDM configuration.

Each production lot should be accompanied by a full documentation package including:

- Certificate of Conformance (CoC) signed by Quality Director.
- Optical characterization data sheet (per-lot statistical summary: mean, sigma, Cpk).
- Reliability test coupon results for the active qualification lot.
- Material Safety Data Sheet (MSDS / SDS) for all adhesives, encapsulants, and metallization.
- Process control chart data demonstrating process stability ($Cpk \geq 1.33$ for all critical parameters).

Key technology transitions anticipated in future revisions include:

- Transition from micro-bump to hybrid bonding as the default Tier 1 baseline in Rev 2.0.
- Inclusion of Integrated Voltage Regulator Module (IVRM) specifications within the CPO package.
- Specification of on-chip wavelength locking and laser control for uncooled operation.
- Standardized 3D integration stacking of TIA + DSP + PIC within a single interposer package.
- Expansion of GI specifications to cover glass-ceramic and ultra-low-expansion (ULE) materials for extreme thermal environments.

7.0 Co-packaged Optics Faceplate Design and Requirements

This section establishes design requirements, mechanical recommendations, thermal budgets, optical interface definitions, management interface recommendations, and interoperability requirements for co-packaged optics (CPO) faceplate assemblies. It covers:

- Single-chassis deployments: 1U–12U spine and leaf switches employing CPO ASICs.
- Scale-up multi-chassis systems: electrically interconnected chassis sharing a backplane or mid-plane.
- Scale-out fabric interconnects: optically interconnected pods, rows, and halls in hyperscale and enterprise datacenter environments.

Compliance with the below recommendations should be required for all faceplate components, cable management accessories, and management software deployed in conjunction with CPO-based network equipment.

The faceplate may be configured in multiple ways, with different combinations of the optical port granularity and connector types. It is important that these are carefully chosen to serve the operational needs of the data center as well as to assist in the thermal management of the switch box, particularly if this involves the air-cooling of >1kW systems. For example, the faceplate of a 51.2T switch box may be configured as 32 ports with 1.6T throughput per port, or 128 ports x 400G per port, with impact on the required panel height (rack real estate) and faceplate surface available for airflow. Basic guidance is provided in this document on the possible configurations and their influence on thermal management.




Figure 20 – The Faceplate of a 51.2T Switch Box may be Configured Multiple Ways however the ELS Ports Require PM fiber

Optical Connectors for Data Ports

The table below lists the fiber channel density for various optical connectors, and the corresponding panel height required DR-based and FR-based switch systems, assuming 100G lane communication. For ease of discussion, DR-based 51.2T bandwidth switches were assumed to require 1024 fibers, while FR-based 51.2T bandwidth switches were assumed to require 256 fibers. Due the reduced fiber requirements for FR-based front panels, additional connector options could be available for a 1U design since a smaller number of ports would be necessary to fit 256 fibers on the faceplate. For DR-based front panels, only connectors with high number of fibers per port would fit within 1U switch design.

Fiber Connector	Fibers per Connector	Connectors per 1RU	Fibers per 1RU	DR (1024f)	FR (256f)
LC	2	72	144	2U+	2U
CS	2	160	320	2U+	1U
MDC	2	216	432	2U+	1U
SN	2	216	432	2U+	1U
MPO-12	12	80	960	2U	1U
MPO-16	16	80	1280	1U	1U
AirMT-12	12	128	1536	1U	1U
MPO-24	24	80	1920	1U	1U
MPO-32	32	80	2560	1U	1U
AirMT-24	24	128	3072	1U	1U
MXC-32	32	104	3328	1U	1U
MMC-16	16	216	3456	1U	1U
SN-MT16	16	216	3456	1U	1U

Table 59 – Fiber Channel Density for Various Optical Connectors

Connector choices also differ in size and functionality: the air-gap (AirMT) and expanded beam connectors are more dust resistant. Connectors such as AirMT, MXC, MMC, or SN-MT16, have a smaller connector size and high fiber/port density to minimize the number of ports and the overall area required on the faceplate.

External Light Source (ELS)

The implementation of optical source light sources is also a matter of consideration for the faceplate design. In the case that the optical source is implemented on the front panel in the form of a pluggable external light source (ELS) modules, its form factor and method of optical connection will influence the face-plate design. A blind-mate connection is preferred, whereby the optical connection between the ELS and the CPO is made through the host-connector inside the module cage, as opposed to having a dedicated connector on the faceplate. This not only improves the eye-safety of the system, but also reduces the footprint, thereby enabling a smaller panel height and/or increased space for airflow.




Figure 21 – ELS Optical Connectivity Design Options

The number of ELS modules required for a CPO system depends on the optical power required by the optical engine, and the output power of the laser channels. For example, if the output of each laser channel is sufficient to power the transmission of 4 x 100G lanes, this would require 128 laser channels in a 51.2T system. For ELS module carrying 8 laser channels, 16 of such modules would be required.

Faceplate Comparison

Once connectors have been selected, any available remaining space could be used for cooling or other components at the switch faceplate. The figure below provides a visual representation of potential connector footprint area. Higher density connector configurations all allow significant reduction in overall space, providing room for other components or reducing the switch from a 2U to a 1U design.

While 64 MPO ports in a 1U panel is possible, there is only a small amount of faceplate area left for airflow or ELS modules. Smaller, denser, or ganged connectors would be needed to free front panel area for other components. Configurations could be envisioned where ELS modules and a connector could be incorporated within a 1U based design. The addition of both the ELS module and optionally the ELS optical connector would require additional faceplate area to incorporate. The configuration and layout of such a design is likely to be specific to that module and require layout of both ELS and fiber connectors in a manner that matches the internal fiber layout.

With a CPO having electrical to optical connections within the switch itself, connector choice availability and placement is no longer constrained to only the pluggable module design. Smaller footprint connectors may be implemented and still maintain the number of fibers or ports necessary for a 51.2T bandwidth switch. The connector choice may be made in combination with other design considerations.

Operational factors such as connector cleanability and breakout capability will need to be analyzed to ensure acceptable ease of use while in operation. The layout of the connectors and ELS modules is also important, as it affects the fiber routing and management inside the switch box. However, specific port layout recommendations are beyond the scope of this specification as stated in a previous paragraph. The faceplate configuration requires careful consideration of thermal management, particularly since the CPO could potentially need to dissipate several kW.

7.1 Mechanical Design Requirements

The faceplate is the primary structural and optical interface element of a CPO-based line card or switch ASIC module. It should simultaneously satisfy high-density port requirements, thermal airflow management, mechanical durability, and deterministic optical alignment under all operating conditions.

Dimensional Envelope:

Parameter	Specification / Requirement
Chassis height (single chassis)	1U (44.45 mm), 2U (88.9 mm), 4U (177.8 mm); EIA-310-D compliant rack units
Faceplate width (full-width card)	482.6 mm $\pm$ 0.3 mm (standard 19-inch rack) or 530 mm $\pm$ 0.3 mm (21-inch)
Faceplate depth (protrusion from rack rail)	$\leq$ 25 mm unmated; $\leq$ 40 mm with cable management installed
Front-to-back airflow clearance	$\geq$ 1U of clear space above and below per ASHRAE A2
Panel thickness (structural region)	3.0 mm $\pm$ 0.1 mm aluminum alloy 6061-T6 or equivalent
Panel thickness (optical window region)	1.5 mm $\pm$ 0.05 mm; minimized to reduce insertion depth
Port cutout tolerance (MPO window)	$\pm$ 0.05 mm on all axes for Gen 1/2; $\pm$ 0.02 mm for future photonic FP
Bending stiffness	Deflection $\leq$ 0.1 mm at mid-span under 50 N transverse load

Table 60 – Dimensional Envelope

Port density should scale with each CPO generation as follows:

Generation	Switch Capacity	Port Count (400G eq.)	Face-Area per Port
Year 2025	51.2T	128× 400G	$\leq$ 5.5 cm ²
Year 2026	102.4T	64× 800G or 128× 400G	$\leq$ 4.5 cm ²
Year 2031	204.8+T	128× 1.6T	$\leq$ 3.0 cm ²

Table 61 – Port Density Outlook

Parameter	Specification / Requirement
FAU fiber pitch (current)	250 $\mu\text{m} \pm 1 \mu\text{m}$
FAU fiber pitch (future)	127 $\mu\text{m} \pm 0.5 \mu\text{m}$ (high-density PIC interface)
Transverse alignment accuracy	$\leq \pm 1 \mu\text{m}$ at faceplate-to-FAU mating surface
Angular tilt tolerance	$\leq 0.1^\circ$ in all axes
FAU latch retention force	$\geq 20 \text{ N}$ (pull-out); $\leq 8 \text{ N}$ (latch release)
FAU temperature stability	Alignment maintained from -40°C to $+85^\circ\text{C}$ (storage); 0°C to $+70^\circ\text{C}$ (operation)

Table 62 – Fiber Array Unit (FAU) Alignment

Environmental and Durability Requirements

Parameter	Specification / Requirement
Operating temperature (inlet air)	0°C to $+45^\circ\text{C}$ (ASHRAE A2 envelope)
Operating humidity (non-condensing)	5% to 85% RH
Storage temperature	-40°C to $+85^\circ\text{C}$
Operating altitude	Sea level to 3000 m (10,000 ft); derated above 1500 m per ASHRAE
Vibration (operating)	IEC 60068-2-6: 5–200 Hz, 0.5 g, 10 sweep cycles per axis
Vibration (shipping)	IEC 60068-2-64: random, 5–500 Hz, 1.5 g RMS
Shock (operating)	IEC 60068-2-27: 11 ms, 15 g half-sine, 3 shocks per axis
Shock (non-operating)	IEC 60068-2-27: 11 ms, 40 g half-sine
Dust/particle ingress	IEC 60529 IP4X (faceplate apertures equipped with dust shutters)
Corrosion resistance	ASTM B117 salt spray, 48 h, no base metal exposure
Flammability	UL 94 V-0 for all polymeric components
Faceplate mechanical lifetime	≥ 15 years in rated operating environment

Table 63 – Environmental and Durability Requirements

7.2 Thermal Requirements

Co-packaged optics place significant thermal management demands on the faceplate region. The silicon photonics die and the switch ASIC share a thermal envelope; the faceplate should not impede airflow to either the ASIC heat sink or the optical die thermal management structures.

Parameter	Specification / Requirement
Airflow direction	Front-to-back (F2B) or back-to-front (B2F) — should be specified per SKU; mixed-direction not permitted within a chassis
Faceplate pressure drop (F2B)	≤ 10 Pa at 200 LFM (1 m/s) inlet velocity
Minimum open area ratio (faceplate)	≥ 65% open cross-section in the optical port zone; ≥ 55% in structural zones
Blockage factor (cable management)	≤ 15% incremental airflow restriction when fully cabled at rated port density
ASIC junction temperature budget	≤ T _{junction_max} – 10°C margin at maximum port utilization and ASHRAE A2 inlet
Thermal gradient across faceplate	≤ 10°C ΔT between hottest and coolest optical port at steady state

Table 64 – Airflow and Thermal Envelope

7.3 Optical Interface Specifications

This section defines the optical transmission parameters at the faceplate interface (point of demarcation between the CPO module and the external fiber plant). All parameters are at the physical medium dependent (PMD) interface.

Fiber Type and Reach Definitions

Parameter	Gen 1 (400G, 800G)	Gen 2 (800G, 1.6T)	Gen 3 (1.6T, 3.2T)
Modulation format	PAM4	PAM4	PAM4 / coherent (TBD)
Lane rate	200 Gbaud IM-DD technology	200 Gbaud IM-DD technology	200 Gbaud IM-DD and Coherent (TBD)
Lanes per port	4 or 8 (DR4)	4 or 8	4 or 8 (1.6T-DR4)
TX output power (per lane)	–3.0 to +3.0 dBm	–2.0 to +4.0 dBm	0.0 to +5.0 dBm
RX sensitivity (BER 1e-6 pre-FEC)	≤ –10.0 dBm	≤ –9.0 dBm	≤ –8.0 dBm
RX overload (damage threshold)	≥ +3.0 dBm	≥ +4.0 dBm	≥ +5.0 dBm
Extinction ratio (PAM4)	≥ 3.5 dB	≥ 3.5 dB	≥ 4.0 dB
Wavelength (primary)	1310nm O-band or CWDM4 (1271, 1291, 1311, 1331nm)	1310nm O-band or CWDM4 (1271, 1291, 1311, 1331 nm) or DWDM C-band TBD	1310nm O-band or CWDM4 (1271, 1291, 1311, 1331 nm) or DWDM C-band TBD

Table 65 – Fiber Type and Reach Definitions

Application	Fiber Type	Reach	Connector Type
DR4 within-row	OS2 G.652.D	500 m	MPO-12
DR4 cross-row	OS2 G.652.D	2 km	MPO-12
FR4 inter-pod	OS2 G.652.D	2 km	MPO-12
LR4 campus	OS2 G.652.D	10 km	MPO-12
ITU-T Standard G.652.D or G.657 within-rack (legacy)	OM5 G.651.1	50 m (100G SR)	MPO-12 UPC
Gen 3 DWDM inter-hall	SSMF G.654.E	80 km (EDFA)	LC/APC duplex
Hollow-core (2031+)	HCPCF (TBD)	≤ 1 km (anticipated)	FAU custom

Table 66 – Fiber Type, Reach and Front Panel Connector Type vs. Applications

Optical Loss Budget (Link Level): The following end-to-end loss budgets apply to each architecture tier.

Architecture Tier	Max Allowed Loss (source to receiver)	Remarks
Top-of-Rack (ToR) to ToR (same row)	≤ 3.5 dB	≤ 15 m cable, 2 mated pairs
ToR to End-of-Row (EoR) switch	≤ 4.5 dB	≤ 100 m, trunk + breakout
EoR to Core (same hall)	≤ 6.0 dB	≤ 500 m, pre-terminated trunk
Cross-hall (scale-out pod mesh)	≤ 7.5 dB	≤ 2 km, MDA/HDA patching
Campus multi-building	≤ 12.0 dB	≤ 10 km, passive plant
Future DWDM inter-campus	≤ 20 dB + EDFA gain	G.654 fiber, Raman+EDFA

Table 67 – Optical Loss Budget (Link Level)

Electrostatic Discharge Protection

Parameter	Specification / Requirement
Human Body Model (HBM)	≥ ±2 kV on all optical ports and connectors (IEC 61000-4-2)
Machine Model (MM)	≥ ±200 V (ESDA/JEDEC JS-001)
Charged Device Model (CDM)	≥ ±500 V for all optical interface pins
System-level ESD (IEC 61000-4-2)	±8 kV contact, ±15 kV air discharge — no permanent damage
ESD wrist strap jack	1 per faceplate (ANSI/ESD S6.1 compliant, 1 MΩ resistor)

Table 68 – Electrostatic Discharge Protection

7.4 Scale-Up Architecture Requirements (Multi-Chassis)

Scale-up architectures interconnect multiple chassis electrically or optically to form a single logical switch or routing domain. CPO faceplates in scale-up configurations should satisfy additional requirements for inter-chassis link uniformity, skew compensation, and fault isolation.

Skew and Latency

Parameter	Specification / Requirement
Inter-chassis fiber length tolerance	± 0.5 m (fiber plant) ± 0.5 m (patch)
Propagation latency (faceplate to faceplate)	≤ 50 ns per chassis hop in a 4-chassis cluster
FEC latency (KP4-RS FEC)	≤ 100 ns per direction
Total cluster round-trip latency	≤ 400 ns for 4-chassis scale-up

Table 69 – Skew and Latency Requirements

Scale-out architectures connect distinct physical pods, rows, or halls via optical fabric interconnects. CPO faceplates in scale-out deployments should accommodate wavelength diversity, long-reach optics, passive plant integration, and phased-array or amplified optical distribution.

Wavelength Management

Parameter	Specification / Requirement
DWDM channel grid (Gen 2/3)	ITU-T G.694.1: 75 GHz; 37.5 GHz (future target)
Channel count per fiber pair	Up to 96 channels on C-band; up to 192 with L-band extension (future)
Wavelength accuracy (uncooled SiPh laser)	± 25 pm over 0–70°C operating range
Wavelength accuracy (thermoelectric-controlled)	± 5 pm (future requirement for 37.5 GHz grid)
OSNR budget (end-to-end, 80 km reach)	≥ 30 dB for 800G-ZR+ at BER 1e-15 post-FEC
PMD limit (total link)	≤ 3 ps for 400 Gbaud PAM4; ≤ 1.5 ps for 800 Gbaud
Chromatic dispersion compensation	Electronic DSP-based in the CPO die; no passive DCF required for ≤ 80 km

Table 70 – Wavelength Management Requirements

7.5 Power Requirements

Co-packaged optics dramatically alter the power distribution profile versus traditional pluggable optics. The faceplate assembly should accommodate the power delivery, dissipation, and monitoring demands of integrated photonic dies.

7.6 Management and OAM Interface Requirements

A well-defined management interface is critical for operational visibility, inventory tracking, predictive maintenance, and automated network operations. CPO faceplates introduce new telemetry points that should be uniformly surfaced.

Beyond CMIS, CPO faceplates should surface the following additional telemetry streams:

- Per-lane BER (pre-FEC and post-FEC), updated every 100 ms.
- Per-lane SNR and modulation error ratio (MER) for PAM4 eyes.
- Per-zone faceplate temperature (§3.3).
- Per-port optical power (Tx and Rx) with 1 dB resolution.
- Connector insertion count (lifecycle tracking).
- SiPh die operating hours and wear indicator.
- Cable plant loss estimate derived from Tx power and Rx power delta.

7.7 Technology Roadmap and Planning Horizon (2027–2035)

This section provides normative guidance for equipment planners, architects, and procurement teams on the expected technology evolution of CPO faceplates over the 2027–2035 horizon. Organizations may use this roadmap to align refresh cycles, cable plant investments, and management software capabilities.

Attribute	Year 2027	Year 2028–2030	Year 2031–2035
Switch capacity	51.2T	102.4T	204.8+T
CPO integration	SiPh flip-chip on ASIC substrate	SiPh 3DIC in ASIC package	Monolithic SiPh-CMOS
Per-port speed	400G DR4/FR4 or 800G DR8	800G DR4/FR4 or 1.6T DR8	1.6T DR4/FR4 or 3.2T DR8 or Coherent (TBD)
Connector type	MPO12, Duplex LC, or MPO16	MPO12, Duplex LC, or MPO16	MPO12, Duplex LC, or MPO16
Fiber interface	SMF G.652.D / G.657.A1 / G.657.A2	SMF G.652.D / G.657.A1 / G.657.A2 / HCF	SMF G.652.D / G.657.A1 / G.657.A2 / G.654.E / HCF
Wavelength	O-band (1310nm or CWDM4)	O-band (1310nm or CWDM4)	O-band (1310nm or CWDM4) or C+L band
Faceplate power (128p)	256 W (optical)	435 W (optical)	704 W (optical)
Thermal design	Air-cooled	Air-cooled + TIM	Liquid-assisted + TIM

Table 71 – Technology Roadmap Summary

Cable Plant Investment Guidance

Facility fiber plant is the longest-lead infrastructure investment in a CPO datacenter. The following guidance minimizes stranded investment:

- Pre-terminate MPO-12 trunk cables for deployment; plan for MPO-16 cassette upgrades at IDF/MDA panels for Years 2028+.
- Allocate duct space for G.654.E ultra-low-loss fiber in long-haul inter-hall routes where next-gen DWDM will be deployed years 2031+.
- Design patch panel density for ≥ 144 fiber-port equivalents per 1U for all new installations; legacy 96-port panels will become a bottleneck at year 2028+ densities.
- Incorporate hollow-core fiber conduit pathways in any infrastructure build starting 2027 to avoid future retrofit costs when HCF products reach commercial availability (~2030).

7.8 Compliance, Testing, and Qualification

CPO faceplate assemblies and complete chassis should pass the following qualification tests prior to product release.

Mechanical Qualification Tests

Test	Method / Standard	Pass Criterion
Insertion/extraction durability	500 cycles, IEC 61753-1 or IEC 61300-2-6	$\Delta IL \leq 0.3$ dB; no physical damage
Fiber pull-out force	IEC 61300-3-26	≥ 5 N per fiber, no fiber breakage
Faceplate flatness (temp cycling)	-40°C to $+85^{\circ}\text{C}$, 100 cycles	Warpage ≤ 0.2 mm over full width
Cable management bend radius	30 mm radius, 10 N tension, 1000 cycles	No optical power degradation >0.5 dB
Rack-mount vibration	IEC 60068-2-6, 5–200 Hz, 0.5 g	No connector disengagement; BER unchanged
Drop / shock (shipping)	IEC 60068-2-27, 40 g, 11 ms	No structural failure; IL within spec after

Table 72 – Mechanical Qualification Tests

Optical Qualification Tests

Test	Method	Pass Criterion
Insertion loss (all ports)	IEC 61300-3-4 substitution method	Within \$4.1 IL budget for 100% ports
Return loss (APC connectors)	IEC 61300-3-6	≥ 60 dB for 100% of ports
Wavelength accuracy	OSA measurement, 0°C to 70°C	Within ± 25 pm (uncooled); ± 5 pm (TEC)
BER under vibration	BERT + IEC 60068-2-6 simultaneously	$BER \leq 1e-12$ post-FEC during vibration
Optical path penalty	TDECQ method per IEEE 802.3ck	Within \$4.1 penalty budget
ALS response time	Fiber disconnect, oscilloscope measurement	Laser off within 300 ms (§5.3)
Power monitor accuracy	NIST-traceable power meter	Reading within ± 1 dB at -20 to $+3$ dBm

Table 73 – Optical Qualification Tests

Thermal Qualification Tests

Test	Conditions	Pass Criterion
Maximum junction temperature	ASHRAE A2 max inlet + full port load	$T_j \leq$ specified max (§3.1)
Thermal runaway immunity	Blocked airflow, 120% rated power	Safe shutdown within 30 s, no fire
Temperature sensor accuracy	Calibrated oven, 0–70°C in 5°C steps	Sensor reading $\pm 1^\circ\text{C}$ of calibrated reference
Differential thermal expansion	–40°C to +85°C ramp, 2°C/min	Optical alignment maintained within $\pm 1\ \mu\text{m}$

Table 74 – Thermal Qualification Tests

7.9 Requirements Traceability Matrix

The table below cross-references all normative requirements to their applicable architecture types.

Req. ID	Section	Single Chassis	Multi-Chassis (Scale-Up / Scale-Out)
MECH-OPT-001 to 007	§2.2.1	Required	Required
MECH-CAB-001 to 005	§2.3	Required	Required
THERM-001 to 004	§3.2	Required (edge-coupled only)	Required
THERM-MON-001 to 003	§3.3	Required	Required
EMC-001 to 005	§5.1	Required	Required
LASER-001 to 004	§5.3	Required	Required
SCUP-001 to 004	§6.1	Not applicable	Scale-up: Required
SCUP-FAULT-001 to 004	§6.3	Not applicable	Scale-up: Required
SCOUT-001 to 004	§7.1	Not applicable	Scale-out: Required
SCOUT-PASS-001 to 003	§7.2	Not applicable	Scale-out: Required
PWR-001 to 005	§8.2	Required	Required
MGMT-CMIS-001 to 003	§9.1	Required	Required
MGMT-TEL-001 to 002	§9.2	Required	Required
MGMT-INV-001 to 003	§9.3	Required	Required
INTER-CON-001 to 003	§10.1	Required	Required
INTER-MGMT-001 to 003	§10.2	Required	Required
INTER-COMPAT-001 to 003	§10.3	Required	Required

Table 75 – Requirements Traceability Matrix

8. Test and Measurement

This section provides the minimum architectural context necessary to apply the Test and Measurement (T&M) requirements in this document. The T&M section defines:

- Development: design verification (DVT), engineering verification (EVT), and production verification (PVT) testing performed prior to volume manufacturing.
- Manufacturing: production test performed at die, module, chassis, and multi-chassis/rack integration levels.
- Installation and Commissioning: field test performed during physical installation, fiber plant qualification, and system bring-up.
- Operation: continuous in-service telemetry, monitoring, predictive maintenance, and field diagnostics performed throughout deployed life.
- End-of-life and decommissioning: final diagnostic capture and component disposition testing performed at decommissioning.

To keep T&M requirements consistent across vendors and lifecycle stages, this document organizes tests by:

- Wafer/die level (PIC wafer, EIC wafer, laser die/bar).
- Component level (singulated laser, PIC, or EIC die; known-good-die screening).
- Optical engine module level (assembled PIC + EIC + laser interface, pre-co-package).
- Co-packaged unit level (ASIC + optical engine(s) as assembled onto a common substrate/package).
- Single-chassis level (fully populated chassis, 4 to 32 co-packaged units per the applicable generation).
- Multi-chassis/pod level (scale-up domain).
- Rack/row level (scale-out fabric segment).
- Fleet/operations level (data-center- or fleet-wide aggregate monitoring and analytics).

8.1 Multivendor Interoperability Testing Principles

Because this is a multivendor practice, the T&M program places particular emphasis on ensuring that test results generated by different vendors, using different test equipment, in different facilities, may be trusted to mean the same thing. Development-phase T&M encompasses design verification testing (DVT), engineering verification testing (EVT), and production verification testing (PVT) performed prior to and during the ramp to volume manufacturing. The objective of this phase is to characterize component- and system-level behavior across the full intended operating envelope, establish the parametric limits that will be enforced in manufacturing and qualify reliability against the target deployed life.

Installation and commissioning T&M verifies that CPO systems, once physically deployed, perform to the same standard demonstrated during manufacturing test, and that the field fiber plant, power, and cooling infrastructure supporting the system meet the design link budget and environmental requirements. Operational T&M provides continuous visibility into the health and performance of deployed CPO systems, supports early detection of degradation before it causes service-affecting failures, and provides the diagnostic data needed to isolate and repair faults efficiently.

8.2 Scale-Up vs. Scale-Out Specific Requirements

Although both scale-up and scale-out interconnect rely on the same underlying co-packaged optics technology, the two architectures impose materially different link budget, latency, and topology test requirements. This section defines reach classes and associated test considerations for each.

Scale-up reach classes:

Typical Reach	Typical Application	Notes
≤ 2 m	Intra-chassis, ASIC-to-ASIC within a single chassis	Shortest reach, tightest latency budget
≤ 7 m	Intra-rack, chassis-to-chassis within one rack	Predominant in Generation 1–2
≤ 50 m	Intra-row, multi-rack scale-up pod	Emerging as optical scale-up matures
≤ 100 m	Row-scale pod	Supports larger scale-up domains

Table 76 – Scale-up Reach

Typical Reach	Typical Application	Notes
≤ 100 m	Leaf-to-spine within a data hall	Parallel SMF, CWDM, Parallel MMF
≤ 500 m	Spine-to-core, building-scale	Parallel SMF, CWDM
≤ 2 km	Campus interconnect	Parallel or CWDM
≤ 10 km	Metro/campus extension	CWDM, Coherent-lite, DWDM C-band

Table 77 – Scale-Out Reach

Scale-up topologies are frequently richly connected (approaching all-to-all) within a pod. As ASIC count per chassis and chassis count per pod grow, the number of distinct link pairs requiring individual validation grows proportional to the square of the participant count and est plans and ATE/field test tooling should account for this scaling.

Growth in optical channel count per chassis directly drives growth in fiber count per chassis, increasing the density requirements placed on fiber shuffle assemblies, front-panel connectors, and field fiber management. The transition to mandatory liquid cooling and the corresponding growth in chassis power density increase the scope and criticality of cooling-related test activities at every lifecycle stage.

9. CPO Systems and Power

CPO System and Power Scope

- Chassis-level system requirements for single-package CPO deployments.
- Power delivery architecture, power budget, and power distribution unit (PDU) requirements.
- Thermal management including air cooling, rear-door heat exchangers (RDHx), direct liquid cooling (DLC), and two-phase immersion requirements.
- Optical interface specifications per CPO year (2027–2034).
- Scale-out implementations and super-cluster topologies spanning multiple rows, pods, halls, or data centers.
- Management, monitoring, and telemetry infrastructure recommendations.
- Reliability, availability, serviceability (RAS), environmental compliance, and safety.
- Roadmap projections and technology transition gates.

CPO chassis impose substantially different power distribution requirements compared to conventional server or networking equipment. The photonic engines require ultra-low noise, tightly regulated DC rails for laser bias, ring heater tuning, and modulator bias voltages, while the ASICs themselves require high-current, fast-transient-response delivery at voltages below 1 V in the most advanced nodes.

ASIC power delivery follows the Voltage Regulator Module (VRM) architecture with fully integrated power stages on the ASIC package substrate or adjacent interposer. APC recommends the following for ASIC power delivery:

- **Input bus:** should be 48 VDC nominal (42–58 VDC range) for all chassis classes; HVDC 380 VDC for C32/C32+ with intermediate bus converters.
- **Output voltage:** Should be processor-specified, typically 0.75–0.95 VDC core; 1.8 VDC memory; 1.2 VDC I/O.
- **Current slew rate:** Should be VRM must support ≥ 200 A/ μ s load transient response without exceeding $\pm 3\%$ voltage deviation.
- **VRM efficiency:** Should be $\geq 94\%$ at 50% load; $\geq 92\%$ at 10% and 100% load.
- **VRM current rating:** Should be sized for $1.25\times$ maximum ASIC TDP to accommodate overclock and boost states.
- **Power telemetry:** Should be per-rail current, voltage, and power at 100 ms intervals minimum; 1 ms burst capability for diagnostics.

ASIC power sequencing must comply with the ASIC vendor specifications. APC recommends that the chassis BMC controls power sequencing via PMBus commands and enforces the following general constraints:

- PGOOD signals from all rails should propagate to BMC within 10 ms of nominal voltage achieved.
- Power-on sequencing should complete within 5 seconds from AC applied to all rails stable.
- Emergency power-off (EPO) should de-energize all ASIC rails within 100 ms of EPO signal assertion.
- Power sequencing state machine should be implemented in non-volatile firmware on the BMC with revision control.

The photonic engine power domain is the most electrically sensitive portion of the CPO chassis. Laser bias currents are susceptible to power supply noise, and modulator bias voltages require millivolt-level accuracy. APC recommends the following for CPO power delivery:

- **CW laser bias rail:** Should be 3.3 VDC (Gen 1/2); 3.3 or 1.8 VDC (Gen 3); 1.2 VDC (Gen 4+); tolerance $\pm 1\%$ static, $\pm 0.5\%$ dynamic.
- **Laser bias noise:** Should be < 1 mV RMS ripple from 1 Hz to 100 MHz, measured at laser bond pad.
- **Ring heater rails:** Should be 2.5–3.3 VDC; current resolution ≤ 100 μ A for thermal tuning of ring resonators. Modulator bias: Should be ± 3 VDC differential for Mach-Zehnder; -2 to $+2$ VDC for ring; ± 10 mV accuracy.
- **SOA/EDFA pump supply:** Should be 5 VDC or 3.3 VDC; galvanically isolated from ASIC power domain.

All CPO power rails should be derived from a dedicated Linear Dropout Regulator (LDO) or ultra-low-noise switching regulator stage with measured output noise < 10 μ V RMS (100 Hz to 100 kHz). Switching regulators serving CPO domains should operate at frequencies > 5 MHz to push switching noise out of the laser modulation bandwidth.

Power Supply Unit (PSU) Requirements

- **Input:** Should be 200–240 VAC; 380–415 VAC 3-phase, universal input 100–277 VAC optional.
- **Frequency:** Should be 47–63 Hz.
- **Power factor:** Should be ≥ 0.99 at 50–100% load (active PFC mandated).
- **Hold-up time:** Should be ≥ 20 ms at full load and 200 VAC minimum input.
- **Efficiency:** Should be 80 PLUS Titanium minimum ($\geq 96\%$ at 50% load, 230 VAC); $\geq 95\%$ at 20% load.
- **Inrush current limiting:** Should be NTC thermistor or active limiting; peak inrush $< 150\%$ rated input RMS.
- **MTBF:** Should be $\geq 400,000$ hours at 25°C ; $\geq 200,000$ hours at 45°C (Telcordia SR-332).
- **PSU form factor:** should comply with OCP Power Supply Project spec version 2.0 or later for hot-swap compatibility.

PSU Redundancy and Failover

- Failover should not cause output voltage deviation exceeding $\pm 5\%$ on any rail.
- PSU status should be reported via PMBus READ_STATUS commands and IPMI SDR records.

Rack and Row Power Infrastructure

For multi-chassis rack deployments, APC recommends the following PDU specifications:

- **Rack PDU current capacity:** should be sized for sum of chassis maximum input currents plus 20% headroom.
- **Metered PDU:** should be per-outlet power metering with $\leq 0.5\%$ measurement accuracy for OPEX tracking.
- **Outlet spacing:** should be a minimum 50 mm center-to-center for C16–C32+ chassis power cords.
- **PDU input:** should be 3-phase delta or star; 400A breaker minimum for rows with multiple C32+ chassis.
- **ATS (Automatic Transfer Switch):** should be a dual-feed ATS PDUs required for all C16+ chassis deployments.
- **PDU management:** should be a SNMP v3 and Redfish API; 1-second power reporting interval.

Data Center Power Infrastructure for Super-Clusters

Super-cluster deployments require coordination with data center power infrastructure planning.

APC provides the following guidelines for facilities planning:

- Minimum power density: may plan for 30–100 kW per rack for C32/C32+ deployments.
- Power factor correction at the facility level: implement capacitor banks or active power conditioning for large CPO deployments should maintain PF > 0.95 at the utility meter.
- UPS architecture: may consider 2N or distributed redundant UPS topology; battery chemistry may support 10-minute minimum runtime at full load for graceful application shutdown.

Forced Air Cooling

Chassis may use forced air cooling as the primary cooling mechanism for ASIC TDPs up to 1,000 W per device, subject to the following constraints:

- **Fan system:** may use N+1 redundant fan modules (hot-swap); minimum airflow per design thermal model at 40°C inlet.
- **Fan speed control:** may use PMBus / I2C with PID loop; fan speed telemetry accessible via IPMI SDR.
- **Airflow direction:** may use front-to-rear; all FRU baffles should prevent recirculation.
- **Pressure differential:** chassis should maintain positive internal pressure relative to ambient to prevent hot-spot recirculation.
- **Filter:** may use electrostatic pre-filter; filter change interval ≥ 6 months at ASHRAE A2 operating conditions.
- **CPO photonic engine cooling:** may use dedicated airflow channel with velocity ≥ 3 m/s across engine face; laser temperature should not exceed 70°C junction.

Air-cooled photonic engines should require individual temperature sensors with $\leq 1^\circ\text{C}$ accuracy, and the chassis thermal controller should reduce ASIC workload the laser module approaches its junction temperature limit.

Liquid Cooling

For deployments exceeding 20 kW per chassis, designers may use rear-door heat exchangers (RDHx) as the minimum cooling enhancement. RDHx units attach to the rear of standard 19-inch rack chassis and use facility chilled water to absorb exhaust heat before it enters the data hall.

- **Coolant supply temperature:** may be 18–24°C; return temperature $\leq 40^{\circ}\text{C}$.
- **Cooling capacity:** RDHx should remove $\geq 80\%$ of chassis thermal load.
- **Leak detection:** RDHx should include optical or resistive leak sensors; alarm via SNMP trap within 5 seconds of detection.
- **Quick disconnects:** may use dry-break couplings rated for $\geq 1,000$ connect/disconnect cycles.

Chassis should support direct liquid cooling (DLC) for ASIC cold plates. DLC eliminates the air gap between the heat sink and coolant, enabling heat rejection at coolant temperatures 10–20°C above junction temperature with far lower thermal resistance than air cooling.

- **Cold plate thermal resistance:** should be $\leq 0.03^{\circ}\text{C}\cdot\text{cm}^2/\text{W}$ (ASIC contact area).
- **Cold plate material:** should be copper or aluminum with nickel plating; compatible with both propylene glycol/water (PGW) and deionized water coolants.
- **Manifold pressure:** should have an operating range 0.5–3.5 bar; burst pressure ≥ 10 bar.
- **Flow rate per cold plate:** should be ≥ 3 L/min at 3.5 bar.
- **Coolant purity:** should have resistivity $\geq 1\text{ M}\Omega\cdot\text{cm}$ for deionized water loops; inhibitor package required for PGW.
- **CPO cooling:** photonic engines on DLC chassis should have dedicated micro-channel cold plate or thermoelectric cooler (TEC) stage to hold laser modules within $\pm 0.5^{\circ}\text{C}$ for wavelength stability.




Figure 22 - Chassis Supporting Direct Liquid Cooling (DLC) for ASIC Cold

Two-phase immersion cooling is anticipated as a required option for chassis exceeding 2,000 W. In this approach, the chassis or a subset of ASICs is submerged in a dielectric fluid that boils at a designed temperature (typically 50–60°C), removing heat via phase change.

- **Dielectric fluid:** PFAS-free engineered fluid with boiling point 50–60°C; compatible with all PCB materials and optical fiber coatings.
- **Vapor recovery:** closed-loop condenser with > 99.9% fluid recovery efficiency to meet environmental regulations.
- **Photonic engine compatibility:** immersion fluid must not degrade optical fiber acrylate coatings, MPO ferrule polishes, or photonic integrated circuit facets over 10-year operating life.
- **Coolant management system (CMS):** rack-level CMS monitors fluid level, vapor pressure, and condenser outlet temperature; integrates with chassis BMC via IPMI.

Photonic components impose unique thermal stability requirements that distinguish CPO thermal management from conventional electronic cooling.

- **Laser wavelength drift:** DFB and VCSEL lasers drift approximately 0.1 nm/°C; for 50 GHz channel spacing (100 pm), temperature must be controlled to $\pm 0.5^\circ\text{C}$ at the laser junction.
- **Ring resonator tuning:** silicon ring resonators drift ~ 80 pm/°C; thermal tuning heaters must compensate for ambient temperature changes at 100 ms response time.
- **Thermo-optic isolation:** ASIC compute die thermal output must not create temperature gradients across the photonic engine exceeding 2°C; thermal isolation structures (low-k dielectric, air gap, or thermal vias) must be present in the package design.
- **Chirp and jitter from thermal cross-talk:** modulator bias voltage is sensitive to temperature; thermal cross-talk between adjacent channels must be quantified in characterization reports and compensated by the wavelength control loop.

9.1 Optical Interface Parameters

APC adopts the ITU-T G.694.1 DWDM flexible grid for wavelength allocation in CPO systems. The initial deployment uses the C-band (191.7 to 196.1 THz, 1,530–1,565 nm) and expands to L-band (184.5 to 191.6 THz, 1,565–1,625 nm) and S-band (196.1 to 205.3 THz, 1,460–1,530 nm) in later generations.

- **96 channels at 100 GHz spacing in C-band;** 400 GHz guard bands at band edges.
- **128 channels at 75 GHz; flex-grid compatible;** C+L band aggregation begins.
- **192 channels at 50 GHz;** flex-grid with 12.5 GHz granularity.
- **384 channels at 25 GHz;** S+C+L full band utilization; Nyquist WDM super-channel capable.
- **Wavelength reference:** ITU-T G.692 wavelength grid; absolute frequency accuracy $\leq \pm 1$ GHz per channel.
- **Wavelength monitoring:** each CPO chassis must provide per-channel wavelength readout via optical spectrum monitoring (OSM) function with ≤ 5 GHz resolution.

APC recommends that CPO system integrators submit link budget analyses for each supported reach configuration. The standard link budget components for a CPO-to-CPO intra-rack path are:

- **Laser output power:** should be +2 to +10 dBm CW (before modulator).
- **Modulator insertion loss:** should be –3 to –5 dB (ring or MZM).
- **On-chip waveguide loss:** should be –0.5 to –1 dB (per 1 cm Si-Ph waveguide).
- **Fiber coupling loss (chip-to-fiber):** should be –0.5 to –1.5 dB (per facet).
- **Connector insertion loss:** should be –0.1 to –0.3 dB (per MPO-12).
- **Fiber path loss:** should be –0.2 dB/km $\times$ L km (OS2 at 1,550 nm).
- **Splitter loss (if applicable):** should be –3.5 dB (1 $\times$ 2 PLC splitter).
- **Receiver sensitivity:** should be –18 to –24 dBm (at target BER pre-FEC for PAM-4 at rated lane rate)
- **Required OSNR:** should be ≥ 28 dB for Year 2026; ≥ 35 dB for later.

Total system optical margin should be ≥ 3 dB at end-of-life conditions including connector cleaning degradation, fiber aging, and laser power degradation. Designs with < 3 dB margin should include a semiconductor optical amplifier (SOA) booster stage. All CPO lanes operating at ≥ 112 G should implement hardware FEC. APC suggests the following FEC requirements:

- **Pre-FEC BER threshold:** should be $\leq 2 \times 10^{-2}$ for all standard FEC codes.
- **Post-FEC BER target:** should be $\leq 10^{-15}$ for C32/C32+ chassis (error-free for 10-year life at full traffic).
- **FEC algorithm:** should be KP4 Reed-Solomon (IEEE 802.3 Clause 91) minimum for Year 2026; concatenated FEC or LDPC for later years.
- **FEC latency budget:** should be ≤ 100 ns for intra-chassis CPO fabric (FEC contributes < 50 ns).
- **FEC counters:** should be per-lane corrected and uncorrected symbol error counts; accessible via chassis management interface.

9.2 Multi-Rack Scale-Up System Implementations

Scale-up CPO implementations extend the high-bandwidth, low-latency optical fabric beyond the single chassis to encompass multiple racks forming a coherent computational unit. The fundamental design principle is to maintain near-chassis-level bandwidth density and latency at rack-to-rack and row-to-row distances, replacing all intra-pod electrical SerDes interconnects with optical or co-packaged photonic fabric. Aggregate BW is total bisection bandwidth of the pod fabric. Fab Interconnect is the per-port speed of the fabric switch tier.

Configurations from MR-64 to MR-256 use a single-tier fabric architecture where all CPO chassis in a 4-rack pod connect directly to a Top-of-Rack (ToR) optical switch. APC defines the following recommendations for this configuration:

- **ToR switch:** CPO-integrated optical switch with ≥ 32 ports at 800G per port; total switching capacity ≥ 25.6 T per switch.
- **Fabric latency:** ToR cut-through latency ≤ 100 ns; optical path latency ≤ 200 ns for intra-rack communication.
- **Load balancing:** Equal-Cost Multi-Path (ECMP) with per-flow hashing; adaptive load balancing via RDMA-aware congestion control.
- **QoS:** minimum 8 traffic classes; strict priority queuing for control plane, weighted fair queuing for data plane.
- **Switch management:** same APC management stack (OpenConfig, gNMI, Redfish) as chassis.
- **Fiber plant:** all ToR switch uplinks and downlinks use pre-terminated MPO-12 OS2 SMF trunk cables; maximum 12 fibers per cable assembly to maintain bend radius compliance.

Configurations MR-512 and above require a 3-tier Clos fabric: Leaf (ToR), Spine, and Super-Spine. APC recommends that all fabric tiers be optically interconnected, eliminating electrical retimers at tier boundaries.

- **Leaf (Tier 1):** CPO-integrated ToR switch; 48 downlink ports to ASIC chassis; 6 uplink ports to Spine.
- **Spine (Tier 2):** High-radix optical switch; 128–512 ports at 1.6 Tb/s per port; optical bypass capable.
- **Super-Spine (Tier 3):** Required for MR-2K+; ROADM (Reconfigurable Optical Add-Drop Multiplexer) fabric enabling dynamic wavelength assignment across pods.
- **Oversubscription:** Leaf:Spine ratio $\leq 3:1$ for AI training workloads; $\leq 8:1$ for inference and storage.
- **Optical amplification:** EDFA booster amplifiers at Spine and Super-Spine output ports for reach > 100 m; gain flatness ≤ 1 dB across C-band.
- **Dispersion compensation:** DCF or FBG-based dispersion compensators required for fiber lengths > 300 m at 224 Gb/s lane rates.

Rack Power and Physical Infrastructure for Scale-Up Systems

- **Rack selection:** all scale-up implementations must use seismically-rated 48U or 52U open-frame racks (RETMA dimension compliant).
- **Aisle containment:** hot-aisle/cold-aisle containment (HACA) required for all pod deployments.
- **Overhead busway:** 400A or 800A overhead busway with tap boxes per rack for ease of reconfiguration.
- **Cable management:** all optical fiber in overhead cable trays with individual cable management arms; minimum tray fill ratio $< 40\%$ to allow future expansion.
- **Power density:** plan for 30–100 kW per rack; data center infrastructure must support ≥ 15 kW/m² floor loading.
- **Physical security:** pod perimeter locking, RFID access logging, and optical fiber tamper detection required for future deployments.

9.3 OCS Control Plane

A CPO scale-up system may operate two logically and physically distinct control domains, and this document treats them separately throughout. The management / OAM plane handles configuration, inventory, telemetry, and firmware over CMIS, Redfish, gNMI, and equivalent transports. It operates on human- and software-management timescales (milliseconds to seconds) and is not latency-deterministic.

The OCS control plane (CP) schedules and operates the optical fabric itself. Where the fabric is a slot-scheduled or fast optical circuit switch, the CP should compute and commit switch configurations, admission, and flow control within the fabric's switching period. Because of the different timescales the CP should not be carried over the management-plane transports, and should be specified as an independent plane.

A conformant fast-OCS CP should provide the following:

- Fabric-coordination messaging
- Grant / demand / credit signaling for scheduling and flow control
- Time-synchronization maintenance

9.4 Scale-Out System Implementations

Scale-out CPO implementations connect multiple pods, rows, halls, or geographically separated data centers into a unified computational cluster. Unlike scale-up systems, which focuses on maximizing bandwidth density and minimizing latency within a pod, scale-out prioritizes fault isolation, dynamic resource allocation, and resilience. APC defines scale-out fabric tiers based on physical distance and latency constraints.

Tier	Protocol	Latency (p50)	Latency (p99)	BW per Port	Topology	Year
Intra-chassis	CPO Fabric	<5 ns	<10 ns	25.6T	All-to-all	2025+
Intra-rack	CPO-Fabric / NVLink	<50 ns	<100 ns	800G	Full mesh	2025+
Rack-to-Rack	Optical Ethernet	<200 ns	<500 ns	400G	Fat-tree	2026+
Row-to-Row	Optical Ethernet	<1 μ s	<2 μ s	800G	Clos	2027+
Pod-to-Pod	WDM Optical	<5 μ s	<10 μ s	1.6T	Butterfly	2028+
DC-to-DC	Coherent DWDM	<100 μ s	<200 μ s	400T	Ring/Mesh	2029+
Super-cluster	APC-Fabric v2	<50 μ s	<100 μ s	3.2T	Dragonfly+	2030+

Table 78 – Scale-Out Network Protocols by Tier

Latency figures include fiber propagation, switch forwarding, and serialization delay. All links are bidirectional.

Intra-Pod to Inter-Pod Scale-Out

The transition from intra-pod to inter-pod scale-out introduces the first long-haul optical spans and the need for coherent transmission. APC recommends the following for pod-to-pod connectivity:

- **Coherent DWDM for spans > 300 m:** DP-QPSK minimum for reaches to 10 km; DP-16QAM for reaches to 2 km; DP-64QAM for intra-data-center reaches < 300 m.
- **Optical amplification:** EDFA with automatic gain control (AGC); OSNR ≥ 28 dB at receiver after all amplification stages.
- **Channel monitoring:** inline optical performance monitors (OPM) at each amplifier site with < 5 ms monitoring update interval.
- **Protection switching:** 1+1 path protection at the optical layer; switchover ≤ 50 ms for link failures (OTN G.873.1 linear protection).
- **Latency awareness:** all scale-out switches must export per-flow latency statistics via OpenConfig telemetry for scheduling feedback to workload orchestration layer.

Super-Cluster Architecture

A super-cluster is defined as a construct interconnected with an all-photonic fabric, achieving petabit-scale aggregate bandwidth. Super-clusters are the target deployment model for frontier AI training, climate modeling, drug discovery, and scientific computing workloads requiring sustained exaflop-scale compute.

- **Physical footprint:** super-cluster requires approximately 8–16 data hall rows; 40–80 standard 48U racks depending on chassis class.
- **Fabric:** Fabric uses a 5-stage Dragonfly+ topology with direct optical links between all group pairs (no store-and-forward hops across groups).
- **All-reduce performance:** APC recommends that a super-cluster must sustain all-reduce at $\geq 95\%$ of theoretical bisection bandwidth for 1 GB message size.
- **Collective communication:** support for ring-allreduce, tree-allreduce, and recursive halving/doubling algorithms; hardware offload for MPI collectives optional but strongly recommended for 2028+.
- **Congestion control:** protocol-level congestion control using Explicit Congestion Notification (ECN) and Back-Pressure with ≤ 2 μ s control loop; RDMA-aware at 2028+.
- **Job scheduling:** resource manager integration (SLURM, Kubernetes, proprietary) via OpenAPI; gang scheduling with fabric-aware placement for topology-sensitive workloads.

Multi-Datacenter CPO Connectivity

For scale-out implementations spanning multiple data centers (DC-to-DC), APC provides the following guidance:

- **Preferred technology:** coherent DWDM with flexible-rate transponders (100G–800G adjustable); ROADM at each DC edge for optical bypass.
- **Latency budget:** 1 ms per 200 km of fiber (propagation); total control-plane latency including buffer time ≤ 10 ms for synchronous distributed training.
- **Synchronization:** IEEE 1588v2 PTP boundary clocks at each DC edge; SYNC-E for physical layer frequency; time accuracy ≤ 100 ns between DCs.
- **Security:** MACsec (IEEE 802.1AE) on all inter-DC optical paths; key rotation every 24 hours; hardware key management module.
- **Traffic engineering:** RSVP-TE or Segment Routing with TE (SR-TE) for bandwidth reservation and path computation; integration with AI workload scheduler for pre-provisioned circuits.

9.5 Management, Monitoring, and Telemetry

The Management Architecture is a hierarchical, vendor-neutral system built on open standards. The architecture consists of four layers: Device Layer (individual chassis BMC/CMC), Rack Layer (rack management controller), Pod Layer (pod orchestration system), and Cluster Layer (super-cluster resource manager). Each layer exposes standardized northbound APIs and communicates with subordinate layers via southbound interfaces.

Management and Monitoring Protocols

Function	Protocol	Minimum Requirement	Notes
Chassis Management	IPMI 2.0 / Redfish	Full IPMI v2 support	OCP Baseline HW Mgmt
Optical Monitoring	OIF MEF 30.1	Per-lane BER, Tx/Rx Pwr	Real-time, <1s update
Power Telemetry	PMBus Rev 1.3.1	All rail V, I, P	100 ms sampling
Thermal Telemetry	MCTP over SMBus	All zone temperatures	10 ms sampling
Firmware Update	PLDM for FW Update	Atomic, rollback safe	MCTP transport
Security	SPDM 1.1 / PCIe IDE	Mutual auth, encryption	RoT required
Time Synchronization	IEEE 1588v2 PTP	<10 ns accuracy	Boundary clock in chassis
Fabric time base (slot-scheduled fabrics)	IEEE 1588-2019 HA	<1 ns accuracy	Edge-devices scheduling for OCS communication. Distinct from management PTP.
Network Mgmt	OpenConfig + gNMI	Full yang model	Streaming telemetry
Inventory	DMTF SMBIOS 3.5	FRU per assembly	UUID per photonic engine

Table 79 – Management and Monitoring Standards.

All standards are normative unless marked optional.

Optical Monitoring Requirements

CPO systems require substantially more detailed optical telemetry than conventional networked systems. APC defines Optical Monitoring parameters, which proposes the following per-lane measurements accessible via the chassis management interface:

- **Tx optical power:** ± 0.5 dB accuracy; 1-second update interval; alarm thresholds configurable.
- **Rx optical power:** ± 0.5 dB accuracy; 1-second update interval.
- **Laser bias current:** $\pm 1\%$ accuracy; used for laser aging prediction.
- **Wavelength (center frequency):** ± 1 GHz accuracy; per-channel.
- **Pre-FEC BER:** per lane; 100 ms update interval; historical trending for predictive failure analysis.
- **Post-FEC BER:** per lane; 1-second update interval; alert at $> 10^{-12}$.
- **OSNR:** ± 0.5 dB accuracy; per DWDM channel.
- **Chromatic dispersion (CD):** ± 100 ps/nm accuracy; trend over 24 hours.
- **Polarization mode dispersion (PMD):** relevant for coherent links > 1 km.
- **Eye diagram quality metrics:** eye width and height percentages relative to UI; accessible via PRBS loopback diagnostic mode.

All optical telemetry must be accessible in real time via gNMI streaming with 100 ms granularity and stored in a time-series database (TSDB) for 90-day retention minimum. Anomaly detection algorithms should flag pre-FEC BER trending upward at a rate > 1 decade per day as a predictive maintenance event.

Time synchronization for slot-scheduled fabrics

The PTP service specified for chassis management (Table 69, boundary clock, < 10 ns) and for inter-datacenter synchronization (≤ 100 ns) is a management-plane facility and is not sufficient to operate a slot-scheduled optical fabric. A fabric that aligns endpoint bursts into shared time frames requires every endpoint and switching communications to agree on the time-frame boundary to better than 1 ns, so that bursts fall within a short guard interval (few ns). This sub-nanosecond timing should be applied throughout applicable optical network communications.

Security and Root of Trust

CPO systems operate at the intersection of computation and communications infrastructure, making security critical. APC recommends the following security requirements for all chassis classes:

- **Hardware Root of Trust (RoT):** NIST SP 800-193 compliant RoT chip on every chassis; TPM 2.0 or equivalent.
- **Secure boot:** all firmware images must be cryptographically signed; unsigned firmware must be rejected.
- **Mutual authentication:** SPDH v1.1 for component authentication between chassis subsystems.
- **Encryption in transit:** PCIe Integrity and Data Encryption for ASIC-to-fabric communication; MACsec for all Ethernet management interfaces.
- **Audit logging:** immutable audit log in RoT-protected storage; 10,000-event minimum capacity.
- **Vulnerability management:** chassis firmware must support CVSS-based severity classification and automated patch notification via PLDM FW Update mechanism.

9.6 Reliability, Availability, and Serviceability Compliance

Reliability and Safety Targets

Requirement Category	Target	Details
MTBF (Chassis)	≥200,000 hours (APC-C4/C8); ≥150,000 hours (C16–C32+)	Telcordia SR-332 Issue 4, Method 1, per field condition
MTBF (Photonic Engine)	≥500,000 hours	Laser lifetime from accelerated aging per IEC 62007-2
MTTR	≤30 minutes (FRU swap)	Hot-swap for PSU, fan, photonic engine; cold-swap for ASIC
Availability	≥99.9999%	Six-9s at chassis level with 2N redundancy active
Vibration	MIL-STD-810H Method 514.8	Operational and survival; sine sweep 5–500 Hz
Shock	MIL-STD-810H Method 516.8	11 ms half-sine, 40 G operational
ESD	IEC 61000-4-2 Level 4	±8 kV contact, ±15 kV air discharge
EMI/EMC (Emissions)	FCC Part 15 Class A; CISPR 32	Radiated and conducted; applies to all SKUs
EMI/EMC (Immunity)	IEC 61000-4 series	Levels 3–4 for industrial environments
Safety	UL 62368-1 / IEC 62368-1 Ed.3	Audio/video, IT and communications equipment
Laser Safety	IEC 60825-1 Ed.3 / FDA 21 CFR 1040	Class 1M maximum at external fiber ports
Thermal Cycling	IEC 60068-2-14 Nb: -40°C to +85°C	250 cycles, 2°C/min ramp, non-operational
Humidity	IEC 60068-2-78; 85°C / 85% RH	1000 hours damp heat, non-condensing operational
RoHS / REACH	EU Directive 2011/65/EU (RoHS 3)	All substances compliant including Annex II

Table 80 – Reliability, Safety, and Environmental Compliance Requirements.

CPO systems introduce new serviceability challenges due to the tight coupling of optical and electronic subsystems. APC recommends the following serviceability requirements to maintain system availability:

- **Hot-swap FRUs:** PSU, fan module, BMC card, and optical amplifier modules must all support hot-swap without chassis power-down.
- **Warm-swap FRUs:** photonic engine modules in chassis generations supporting module-level CPO may support warm-swap with < 1-second traffic interruption using optical protection switching.
- **Cold-swap FRUs:** ASIC die and package; chassis backplane; main switching fabric ASIC.
- **FRU labeling:** each FRU must carry an RFID tag or 2D barcode with part number, serial number, manufacturing date, and firmware version accessible without disassembly.
- **Fiber cleaning:** MPO connectors on FRUs must be accessible with standard interferometric cleaner without tools; IEC 61300-3-35 contamination limits required before installation.
- **ASIC removal procedure:** documented maximum fiber bend radius of 15 mm must be maintained during all ASIC servicing procedures; fiber slack management tray required in C16+ chassis.
- **MTTR target:** ≤ 15 minutes for PSU/fan swap; ≤ 30 minutes for photonic engine; ≤ 60 minutes for ASIC.

Laser Safety

CPO systems by nature aggregate substantial optical power at the external fiber ports. APC recommends the following laser safety measures:

- All external fiber ports must comply with IEC 60825-1 Ed. 3 Class 1M or better; Class 3B and 4 classifications are not permitted at chassis boundary.
- **Automatic laser shutdown (ALS):** APC recommends IEC 60825-2 ALS implementation on all chassis external fiber ports; laser must shut down within 300 ms of detecting fiber break (via OTDR pulse or power drop ≥ 10 dB).
- **Interlock:** chassis rear panel must include an optical fiber panel interlock that removes laser power from all external ports within 100 ms if the fiber panel is opened during operation.
- **Warning labels:** IEC 60825-1 Annex D format labels in English and the language of the target market; minimum 1 label per fiber panel.
- **Aperture controls:** all fiber routing channels within the chassis must be enclosed to prevent inadvertent personnel exposure to stray optical power > 1 mW.

Environmental and Sustainability

- **RoHS 3 (EU 2011/65/EU):** full compliance including Annex II substances; certificate of conformity required.
- **REACH SVHC:** candidate list substances below 0.1% w/w threshold; annual supply chain declaration.
- **WEEE Directive:** chassis must be designed for disassembly within 30 minutes for material stream separation.
- **Packaging:** $\geq 80\%$ recyclable or reusable packaging materials; EPD (Environmental Product Declaration) for C16+ chassis within 24 months of product launch.
- **Energy efficiency:** EU Ecodesign Regulation 2019/424 compliance for server equipment; EPEAT Gold registration target for all chassis classes.
- **Carbon reporting:** APC recommends Product Carbon Footprint (PCF) reporting per ISO 14067 for chassis supply chains, recognizing that optical interconnects offer significant lifecycle carbon reduction vs. electrical alternatives.

9.7 Technology Roadmap 2026–2034

Near-Term Technology Transitions (2025–2027)

The 2025–2027 period is characterized by the transition from near-packaged optics (NPO) and pluggable modules to true co-packaged photonic engines. Key technology transitions include:

- **112G to 224G lane rate transition:** driven by silicon photonics modulator bandwidth improvements and introduction of TIA/DSP chiplets within the CPO package.
- **External cavity laser (ECL) adoption:** superior wavelength stability and power efficiency compared to directly modulated lasers; enables 75 GHz and 50 GHz channel spacing.
- **Advanced packaging adoption:** wafer-scale integration of Si-Ph PICs with CMOS drivers using hybrid bonding ($< 1 \mu\text{m}$ pitch); reduces electrical parasitics and enables multi-layer photonic routing.
- **48V DLC deployment:** 48V direct liquid cooling systems achieving PUE < 1.15 at rack level; transition to HVDC (380V) bus for efficiency gains in chassis.
- **CPO socket:** Mechanical and electrical interface recommendations for interchangeable photonic engines from multiple vendors (similar to DIMM standardization for memory).

Mid-Term Technology Transitions (2027–2030)

- **448G lane rates and QAM-16:** enabled by InP-Si hybrid photonic platforms with electro-optic bandwidths $> 100 \text{ GHz}$; coherent DSP co-packaged with modulator die.
- **Optical memory interfaces:** CPO-extended memory fabrics using optical links to photonic memory interfaces (PMI) for AI model weight distribution at 100T per ASIC to memory pool.
- **Optical switching fabric integration:** optical circuit switches (OCS) with sub-microsecond reconfiguration time replacing static MPO fiber patch panels in scale-up pods.
- **All-optical 3-stage Clos fabric:** elimination of electrical-optical-electrical (EOE) conversions at fabric nodes; pure photonic routing for intra-pod fabric with latency $< 10 \text{ ns}$ per hop.
- **Silicon nitride (Si₃N₄) photonics:** ultra-low-loss waveguides ($< 0.1 \text{ dB/m}$) enabling complex on-chip photonic signal processing including filtering, spectral monitoring, and optical equalization.
- **Micro-EDFA:** on-chip erbium-doped waveguide amplifiers providing 10–15 dB gain for loss compensation without external discrete EDFA modules.

9.8 Power Consumption Estimation Worksheet

System integrators should use the following methodology to estimate total power consumption for APC CPO deployments:

- Identify chassis class and quantity.
- Obtain per-ASIC TDP from ASIC vendor datasheet (specify boost TDP, not base TDP).
- Apply CPO photonic engine overhead (add to per-ASIC power).
- Calculate chassis total IT load: $(\text{ASIC TDP} + \text{CPO overhead}) \times \text{ASIC count per chassis}$.
- Add chassis infrastructure overhead: memory, storage NVMe, BMC, fans, and board-level power = typically 5–10% of ASIC TDP.
- Apply PSU efficiency factor: divide chassis IT load by PSU efficiency.
- Sum all chassis in deployment to obtain total rack IT load.
- Add networking/switching power: fabric switches = 10–20% of chassis IT load for balanced topologies.
- Apply PUE factor for cooling and facility overhead
- Total facility power = Step 9 output $\times$ facility PUE.

10. Cooling Systems

Section 10 is to establish recommended requirements for cooling systems deployed in hyperscale, enterprise, colocation, and AI-focused data center facilities. It covers the full spectrum of thermal management technologies required to sustain safe, efficient, and reliable operations across a ten-year planning horizon from 2027 through 2035.

The Recommended Practices document addresses a period of unprecedented growth in compute density, driven principally by large-scale artificial intelligence training clusters, high-performance computing (HPC) workloads, and the broad proliferation of GPU-accelerated infrastructure. Cooling requirements that were considered extreme edge cases in prior generations are becoming routine operational baselines, and this document codifies requirements accordingly.

This revision (October 2026) adds comprehensive requirements for Co-Packaged Optics (CPO) systems, covering single-chassis configurations, scale-up multi-rack installations, and scale-out distributed multi-rack architectures. CPO thermal management represents a fundamentally distinct engineering challenge from conventional discrete-component cooling and is addressed in dedicated sections of this document.

This section applies to:

- All new data center construction with a planned capacity exceeding 1 MW of critical IT load.
- Existing facilities undergoing major capital refresh or expansion of 500 kW or greater.
- Colocation operators serving tenants with rack densities of 10 kW per rack or higher.
- Equipment manufacturers supplying cooling infrastructure.
- AI/HPC cluster deployments where aggregate thermal load per deployment zone exceeds 50 kW.
- CPO systems and CPO-integrated chassis from single-node to multi-rack scale.

This document does not apply to edge computing nodes, telecommunications central offices, or facilities with total critical IT load under 500 kW unless those facilities voluntarily adopt these recommendation

10.1 Market Drivers

AI model training and inference workloads have become the dominant growth vector for data center compute. Large language models (LLMs), multimodal foundation models, and scientific AI applications require dense GPU or AI accelerator clusters that dissipate 700W to 1,200W per chip as of 2025, with roadmap projections indicating 1,500W to 2,500W per chip by 2030. A single AI training pod of 512 to 1,024 GPUs may represent a thermal cluster load of 400 kW to 1.2 MW within a contiguous footprint of 20 to 40 rack positions. Next-generation accelerators from leading semiconductor manufacturers confirm continued TDP escalation. Packaging innovations including 3D chip stacking, chiplet architectures, and co-packaged optics further concentrate thermal energy at increasingly small spatial scales, fundamentally challenging traditional heat spreading mechanisms and air-cooling effectiveness. Global hyperscale operators are projected to commission over 50 GW of new data center capacity between 2026 and 2035. Even edge and near-edge deployments are incorporating AI inference hardware, pushing dense liquid cooling requirements into environments historically served only by air cooling.

10.2 Thermal Load Growth Projections

Year	Avg Rack Density	Leading Edge Rack	AI Cluster Pod Peak	Global DC Capacity
2026	10–15 kW	60–100 kW	600 kW – 1.5 MW	~870 GW
2027	12–18 kW	80–120 kW	800 kW – 2.0 MW	~950 GW
2028	15–22 kW	100–150 kW	1.0 – 2.5 MW	~1,050 GW
2029	18–28 kW	120–175 kW	1.2 – 3.0 MW	~1,150 GW
2030	20–35 kW	140–200 kW	1.5 – 4.0 MW	~1,280 GW
2031	24–42 kW	160–250 kW	1.8 – 5.0 MW	~1,420 GW
2032	28–50 kW	200–300 kW	2.0 – 6.0 MW	~1,580 GW
2033	32–60 kW	240–350 kW	2.5 – 7.5 MW	~1,750 GW
2034	38–70 kW	280–400 kW	3.0 – 9.0 MW	~1,940 GW
2035	42–80 kW	320–500 kW	3.5 – 12 MW	~2,150 GW

Table 81 – Thermal Load Growth.

Projections are based on APC analysis of semiconductor roadmaps, hyperscaler capex disclosures, and industry analyst data as of Q1 2026. Actual deployment densities may vary.

For data center cooling system designers and operators, CPO introduces three novel thermal challenges that are distinct from conventional ASIC cooling:

- **Concentrated mixed thermal loads:** A single CPO package dissipates both electronic (ASIC die) and photonic (optical engine) heat, with fundamentally different temperature sensitivities for each component type. ASIC dies tolerate junction temperatures up to 90–105°C; silicon photonics optical engines require tight thermal control in the range of 25–70°C with ±2°C stability for reliable wave-length-division multiplexing (WDM) operation.
- **Optical thermal sensitivity:** Laser threshold current, photodetector responsivity, optical modulator chirp, and silicon waveguide effective index are all temperature-dependent. Thermal excursions beyond ±5°C of the optical engine setpoint may cause bit error rate (BER) degradation, link failure, or permanent damage to photonic structures.
- **Scale concentration:** CPO systems are designed to eliminate external pluggable optics, concentrating all optical I/O thermal load onto the package and chassis cold plate. In a 16-ASIC chassis at 2030 TDPs, this concentrates up to 8 kW of optical engine heat within a single 2–4U chassis footprint.

APC establishes three CPO-specific thermal sectors that overlay the existing T1–T5 rack classification framework. CPO sectors classify the optical engine thermal management challenge independently of the ASIC TDP, recognizing that the optical cooling requirement is often the more demanding constraint.

OE Heat per Chassis	OE Temp Setpoint	Required Stability	Cooling Method	Era
< 1 kW	45–65°C	±3°C	Air or shared DLC loop	2025–2027
1–5 kW	35–55°C	±2°C	Dedicated CPO-CDU; isolated loop	2026–2030
5–15 kW	25–50°C	±1°C	Dual-loop CPO-CDU; photonic temp control	2028–2032
> 15 kW	20–45°C	±0.5°C	Active TEC + liquid; cryogenic assist may be required	2031–2035

Table 82 – CPO Thermal Classification

Regulatory and Standards Landscape for CPO Cooling

As of 2026, no single international standard governs CPO cooling infrastructure. The APC CPO recommendations of this document are intended to bridge this gap and will be submitted to IEEE P802.3dk, the OIF CPO Working Group, and ASHRAE TC 9.9 for harmonization. APC members and the industry at large should treat these requirements as normative pending adoption by international bodies:

- **OIF CPO Implementation Agreement (IA) series:** Addresses electrical and optical interface requirements but does not specify cooling infrastructure.
- **IEEE 802.3ck/df:** Governs 100G lane electrical interfaces relevant to CPO SerDes elimination but does not address thermal management.
- **SEMI E187:** Semiconductor equipment environmental and safety standards — applicable to CPO manufacturing but not data center deployment.
- **ASHRAE TC 9.9:** Working group on liquid cooling for CPO-integrated systems; APC is an active contributor; draft guidance expected 2027.

10.3 Single-Chassis CPO Cooling Requirements

This section specifies detailed thermal and cooling requirements for single-chassis deployments of CPO-equipped ASIC systems across the 4-ASIC configurations. Single-chassis deployments are defined as all ASIC and CPO components within a single mechanical enclosure (1U to 12U) sharing a common power and cooling infrastructure.

4-ASIC CPO Chassis and Cooling Requirements

The 4-ASIC chassis is the CPO deployment unit targeted at inference servers, edge AI acceleration, and research workloads. As of 2025–2027, 4-ASIC CPO chassis are the first commercially available CPO-integrated systems. Representative form factors are 4U–6U with front-access CPO optical fiber management.

Configuration Profile

Parameter	2027 Requirement	2028–2035 Requirement
ASIC TDP per die (target)	700–1,000 W	1,200–2,000 W
Total ASIC TDP (4 dies)	2.8–4.0 kW	4.8–8.0 kW
CPO Optical Engine count	8–32 OE total (2–8 OE/ASIC)	64–256 OE total
CPO OE aggregate heat	0.08–0.48 kW	0.32–5.1 kW
Total chassis thermal load	2.9–4.5 kW	5.1–13.1 kW
Chassis cooling method	Shared cold plate DLC	Dual-loop DLC (CPO + ASIC separate)
ASIC cold plate coolant inlet	$\leq 22^{\circ}\text{C}$	$\leq 25^{\circ}\text{C}$ (elevated coolant permitted)
CPO OE coolant inlet (dedicated)	$35\text{--}55^{\circ}\text{C} \pm 2^{\circ}\text{C}$	$25\text{--}50^{\circ}\text{C} \pm 1^{\circ}\text{C}$
Coolant flow rate (total chassis)	15–25 L/min	25–60 L/min
CDU capacity required	≥ 10 kW	≥ 20 kW
Chassis form factor	4U–6U, 19-inch EIA or OCP	4U–6U; OCP preferred
Rack slots consumed	4–6U	4–6U
APC CPO Tier	CPO-T1 to CPO-T2	CPO-T2 to CPO-T3

Table 83 – Configuration Profile

- Coolant supply to CPO optical engines should be independently adjustable from ASIC cold plate supply temperature, with a minimum $\pm 15^{\circ}\text{C}$ range relative to facility CDU supply.
- A dedicated CPO temperature control valve or secondary heat exchanger with TEC assist should be provided to maintain optical engine temperature within $\pm 2^{\circ}\text{C}$ of setpoint regardless of ASIC thermal load variation.
- Chassis should include per-OE or per-optical-group thermal monitoring with a minimum resolution of 0.1°C and a sampling interval of ≤ 1 second.
- In-chassis leak detection should generate an alarm within 30 seconds and should integrate with facility CDU to isolate coolant flow within 60 seconds of confirmed leak.
- Minimum fiber management clearance of 150mm should be maintained at the chassis optical I/O panel; fiber bend radius should not be violated during standard rack installation and maintenance operations.

Universal Single-Chassis CPO Recommendations

The following recommendations apply to all CPO-equipped single-chassis deployments regardless of ASIC count:

- All CPO-equipped chassis should implement hardware-level OE over-temperature protection that reduces laser drive current within 100ms of OE temperature exceeding the upper design limit.
- OE temperature should be logged at minimum 1-second intervals with timestamps synchronized to facility time reference (NTP or PTP, ≤ 1ms accuracy).
- Chassis should not boot to operational state if CPO coolant flow is absent or below minimum flow threshold. Pre-boot CPO cooling check should be integrated into BIOS/UEFI/BMC firmware.

Coolant Specifications for CPO Systems

Parameter	ASIC Cold Plate Loop	CPO Optical Engine Loop
Coolant type	DI water + inhibitors or engineered fluid	DI water + inhibitors (no corrosive additives)
Conductivity limit	< 10 µS/cm alarm	< 2 µS/cm alarm (stricter for photonic contact)
pH range	7.0–9.0	6.8–7.5 (neutral to protect SiPh waveguides)
Particulate filtration	10 µm filter	5 µm filter minimum; 1 µm preferred
Dissolved oxygen	< 50 ppb preferred	< 10 ppb (prevents SiPh surface oxidation)
Chloride concentration	< 50 ppm	< 5 ppm (SiPh corrosion risk)
Coolant inlet temp (chassis)	≤ 22–30°C (see tier tables)	25–55°C (see tier tables)
Flow stability requirement	± 5% of setpoint	± 2% of setpoint (optical stability)
Minimum operating pressure	0.15 MPa	0.10 MPa
Maximum operating pressure	0.5 MPa	0.35 MPa (SiPh substrate stress limit)

Table 84 – Coolant Specifications for CPO Systems

Fiber and Optical Interface Management

All CPO chassis should ship with factory-tested, pre-terminated optical fiber assemblies. Field-terminated CPO fiber connections are not permitted for in-warranty operation. Optical fiber routes between CPO chassis and optical switching fabric should be managed in dedicated overhead or underfloor fiber trays with a minimum bend radius of 15mm enforced by fiber guides at all routing transitions. Fiber connections at CPO chassis face plates should use MPO/MTP connectors with IEC 61754-7 grade B or better angular physical contact polish to minimize return loss-induced OE temperature variation.

10.4 CPO Scale-Up Installations: Cooling Requirements

Scale-up CPO installations should implement a hierarchical coolant distribution architecture with the following functional layers:

- **Facility Cooling Plant (FCP):** The primary chiller or dry cooler plant serving the data hall. For scale-up CPO zones, FCP should provide a minimum of N+1 capacity redundancy.
- **Zone Coolant Distribution Unit (Zone-CDU):** A rack-row-level CDU serving 2–8 racks. Zone-CDU manages pressure, temperature, and flow to individual chassis-level connectors. Zone-CDU should be rated for 120% of the maximum connected chassis thermal load.
- **Chassis Coolant Distribution Module (CCDM):** The chassis-integrated cooling module that distributes coolant to individual ASIC cold plates and CPO optical engine loops within each chassis.
- **CPO-Dedicated Secondary Loop:** A dedicated secondary loop within the Zone-CDU serving only CPO optical engines at independently controlled temperature, isolated from the ASIC primary coolant circuit.

Parameter	SU-2R / SU-4R Zone-CDU	SU-8R to SU-32R Zone-CDU
Rated cooling capacity	200–800 kW	800 kW–3.2 MW (modular expansion)
CPO dedicated loop capacity	20–80 kW	80–320 kW
Number of rack connections	2–4 racks	8–32 racks
CPO loop temperature control	35–55°C ± 1°C	25–55°C ± 0.5°C
ASIC loop temperature control	18–28°C ± 1°C	18–30°C ± 0.5°C
Primary coolant flow rate	200–800 L/min	800–3,200 L/min
CPO secondary loop flow	20–80 L/min	80–320 L/min
Pump redundancy	N+1 per loop	N+2 per loop; auto-failover < 5 sec
Heat exchanger approach	≤ 2°C at design conditions	≤ 1.5°C at design conditions
BMS integration	Modbus TCP + REST API	REST API + real-time streaming (Kafka/MQTT)
Supply/return header size	2-inch (DN50) minimum	4-inch (DN100); 6-inch for SU-32R
Leak detection	Zone-level; alarm ≤ 60 sec	Zone + rack-level; alarm ≤ 15 sec; auto-isolation ≤ 30 sec
CPO conductivity monitoring	< 5 µS/cm alarm (continuous)	< 2 µS/cm alarm (continuous)
Zone-CDU footprint	≤ 0.5 rack bays per 200 kW capacity	≤ 0.4 rack bays per 300 kW; in-row or top-of-row

Table 85 – Zone-CDU Requirements for Scale-Up Installations

Optical Fabric Thermal Management in Scale-Up

Scale-up CPO installations include one or more optical fabric switches that provide all-to-all ASIC connectivity within the cluster. These switches combine high-radix optical switching with their own CPO interfaces and represent a significant thermal load that should be explicitly addressed. Optical fabric switches in scale-up clusters should be liquid-cooled using the CPO secondary coolant loop or a dedicated third loop at a temperature compatible with both switch ASIC TDP and switch CPO OE operating requirements.

Optical fabric switch thermal load should be included in Zone-CDU capacity calculations at 110% of rated switch TDP to account for optical switching loss conversion to heat. A minimum of 50mm air gap should be maintained between optical fabric switch chassis and adjacent CPO server chassis to enable fiber routing without bend radius violation.

Scale-Up Electrical Power Distribution Requirements

- Scale-up CPO installations impose power distribution requirements that exceed standard data center busway systems. The following requirements apply:
- Total zone power for SU-8R and larger should be distributed via 480V 3-phase busway at a minimum ampacity of 2,000A per phase per rack row; 4,000A per phase is recommended for SU-16R and SU-32R installations.
- Power factor correction should be implemented at the zone PDU level to maintain a minimum power factor of 0.97 under all load conditions, reducing reactive current demand on facility transformers.
- Each rack should have independent A+B dual power feeds with automatic transfer switching rated for the full rack load; transfer time should not exceed 20ms to avoid CPO laser thermal transient from power interruption.
- CPO optical engine power supplies within chassis should be independently redundant from ASIC processor power supplies, enabling CPO thermal stabilization to continue during ASIC power fault events.

Scale-Up Thermal Monitoring and Control

APC Recommended Monitoring Points for Scale-Up CPO Zones:

- Per-chassis: ASIC cold plate inlet/outlet temperature, CPO OE loop inlet/outlet temperature, coolant flow rate on each loop, leak sensor status — minimum 1-second resolution.
- Per Zone-CDU: Primary loop supply/return temperature, CPO loop supply/return temperature, all pump speeds and currents, heat exchanger approach temperature, conductivity readings — minimum 5-second resolution.
- Per rack: Rack-level power consumption (A+B feeds), ambient inlet temperature (3 vertical positions), optical fiber temperature (at dense fiber routing points) — minimum 30-second resolution.
- Cluster-level: Aggregate thermal load, PUE and partial-PUE for the CPO zone, CPO optical link BER correlation with temperature events — minimum 1-minute resolution.

Thermal Control System Integration

Scale-up CPO zones should integrate their cooling control systems with facility DCIM within 90 days of commissioning. DCIM integration should provide read access to all required monitoring points and write access to Zone-CDU temperature setpoints. CPO optical engine temperature setpoint changes should be rate-limited to 1°C per minute to prevent rapid thermal cycling of photonic structures; faster changes require explicit override with audit logging. Scale-up cooling control should implement feedforward control based on scheduled HPC/AI job queue data; predictive CDU pre-cooling should be initiated ≥ 5 minutes before a scheduled high-load job.

10.5 CPO Scale-Out Installations: Distributed Multi-Rack Requirements

Scale-out CPO zones impose facility-level cooling plant requirements that should be addressed in the data center master infrastructure plan well before CPO chassis deployment. The following recommendations apply to facilities hosting SO-1 or larger scale-out CPO deployments:

- Chilled water plant capacity should be sized for 110% of the projected peak CPO zone thermal load at full build-out, with provision for expansion to 150% of initial design within 24 months without plant replacement.
- Facilities hosting SO-2 or larger should implement a dedicated CPO secondary cooling loop fed from a separate heat exchanger circuit within the chilled water plant. This CPO secondary loop should maintain supply temperature independently of the main ASIC cooling loop.
- CPO secondary loop supply temperature should be controllable in the range of 20°C to 60°C with $\pm 1^\circ\text{C}$ accuracy. This range accommodates both near-future CPO-T2 systems and long-horizon CPO-T4 elevated-temperature operations.
- Heat rejection systems (cooling towers, dry coolers) should be provisioned for 100% of scale-out CPO zone peak load on the CPO secondary loop independently, as CPO optical engines may require active cooling during periods when ASIC loads are low and heat rejection from the primary loop is insufficient to maintain plant operation.

Distribution Piping for Scale-Out CPO

Distribution Level	ASIC Primary Loop	CPO Secondary Loop	Notes
Facility supply/return main	≥ 8 -inch (DN200)	≥ 4 -inch (DN100)	Insulated CPVC or stainless; pressure rated ≥ 1.6 MPa
Zone branch (per SU-pod)	4-inch (DN100)	2-inch (DN50)	Zone isolation valve; pressure gauge; flow meter
Row sub-branch (per 4 racks)	2-inch (DN50)	1-inch (DN25)	APC-CS-01 compliant quick-disconnect headers
Rack drop (per rack)	1-inch (DN25)	0.75-inch (DN20)	Flexible hose last 500mm; drip loop mandatory
Chassis connection	Quick-disconnect	Quick-disconnect	Color coded: ASIC=blue/red; CPO=green/yellow
Pipe material	CPVC, 316 SS, or carbon steel with inhibitors	316 SS or CPVC preferred	No copper in CPO secondary loop (SiPh compat.)
Velocity limit	0.9–2.5 m/s	0.5–1.5 m/s (lower for CPO precision)	Erosion and noise limits
Expansion provisions	Expansion joints every 30m	Expansion joints every 20m	Thermal growth accommodation

Table 86 – Distribution Piping for Scale-Out CPO

Scale-Out Optical Fabric Infrastructure Recommendations

Scale-out CPO deployments are distinguished from scale-up by the presence of multi-tier optical switching fabric that itself contains CPO interfaces and should be thermally managed as a first-class CPO installation. These optical fabric switches ('spine switches' or 'leaf-spine switches') are often the highest-density CPO systems in the deployment.

Optical spine switches in SO-2 and larger deployments should have CPO-integrated optical I/O and should be liquid-cooled. Air-cooled optical spine switches should not be utilized in SO-2 or larger deployments after 2027. Optical spine switches should be installed in dedicated switch racks or spine rows separate from compute racks. Minimum 1,200mm aisle width should be maintained on both sides of spine switch rows for fiber management. Spine switch rows should be served by the CPO secondary cooling loop at the same temperature as the connected compute CPO optical engines, ensuring thermal consistency across the optical path. Optical spine switch racks should be equipped with overhead fiber management trays of minimum 600mm × 300mm cross-section to accommodate the extreme fiber density of multi-pod spine connections. Spine switch power density should be included in zone thermal calculations; representative spine switch TDPs for 2025–2030 deployments are 10–40 kW per 1U switch chassis.

All optical fiber links between scale-out CPO pods should use OS2 single-mode fiber (ITU-T G.652D or G.657A1) with MPO/MTP pre-terminated trunk cables factory tested to IEC 61753-1 performance standards. Optical fiber pathways between compute pods and spine switches should be installed in dedicated overhead cable trays with no shared routing with power, cooling hoses, or copper data cables. Minimum 200mm separation from power cables.

All inter-pod fiber links should be labeled at both endpoints. Labels should be machine-readable (QR or RFID) and should encode link ID, source/destination chassis, wavelength plan, and installation date. Fiber routing should maintain bend radius compliance ($\geq 15\text{mm}$ for OS2, $\geq 10\text{mm}$ for bend-insensitive G.657A2) throughout all routing transitions including cable tray turns, rack entry, and chassis attachment points.

A fundamental requirement for scale-out CPO deployments is thermal independence between pods, enabling individual pods to be brought up, taken down, or thermally serviced without affecting adjacent pods. This independence is critical for continuous AI training operations where pod failures should be quickly isolated.

Each SU-pod within a scale-out deployment should have independent Zone-CDU isolation valves on both ASIC primary and CPO secondary loops, operable without draining the facility distribution loop. Isolation valve actuation time should not exceed 30 seconds for manual operation or 5 seconds for automated emergency isolation triggered by leak detection or thermal alarm. Pod thermal isolation should not affect coolant supply temperature or flow rate to any other pod in the scale-out cluster by more than 5% for more than 60 seconds. Each pod should maintain its own CDU flow and temperature telemetry independent of adjacent pods, enabling granular thermal diagnostics without cluster-wide correlation.

Scale-Out Thermal Zone Layout

Cooling Zone Structure	Zone-CDU Count	Facility Loop Connection
1 zone per SU-pod; single Zone-CDU	1 per SU-pod	Direct facility loop connection with pod isolation valve
1 zone per 2 SU-pods; row-level Zone-CDU	1 per 2 pods	Row distribution manifold; pod isolation within row
Row-level zones; hall-level secondary plant	1 per 4 pods + 1 hall-level CDU	Hall distribution loop; row sub-loop; pod isolation
Hall-level zones; campus cooling plant	Distributed; campus-level chiller	Campus loop with inter-hall CPO cooling coordination

Table 87 – Scale-Out Thermal Zone Layout

Scale-Out DCIM and Thermal Telemetry Requirements

The scale and complexity of scale-out CPO deployments exceed the capacity of traditional polling-based DCIM systems. APC recommends streaming telemetry architectures for SO-2 and larger deployments. All Zone-CDUs, optical spine switches, and CPO chassis should publish thermal telemetry to a facility-wide message broker (MQTT, Kafka, or equivalent) at configurable intervals of 1–60 seconds, defaulting to 5 seconds. Telemetry data should include: timestamp (PTP-synchronized, $\leq 1\text{ms}$ accuracy), chassis/device ID, ASIC cold plate temperature (per-die), CPO OE temperature (per OE or OE group), coolant flow rate, coolant conductivity, pump status, leak sensor status, and optical link BER per port. Facilities hosting SO-3 or larger deployments should implement a dedicated thermal digital twin updated in real-time from telemetry streams, capable of predicting thermal exceedances ≥ 5 minutes in advance with $\geq 85\%$ accuracy. The digital twin should be capable of modeling the thermal effect of planned AI workload migrations between pods to predict cooling infrastructure impact before execution.

Scale-Out Resilience and Redundancy

Scale-out large deployments should implement N+1 Zone-CDU redundancy at the row level; loss of one Zone-CDU should not cause thermal exceedance in any pod served by that CDU within 5 minutes. CPO secondary loop redundancy should be treated as more critical than ASIC primary loop redundancy, as OE temperature exceedances have shorter time-to-damage characteristics than ASIC silicon. CPO secondary loop should have ≤ 30 -second failover; ASIC primary loop should have ≤ 120 -second failover. Larger deployments should implement geographic redundancy for cooling plant. No single mechanical room failure should affect more than 25% of the scale-out cluster's cooling capacity.

Scale-out CPO deployments should implement a cluster-level emergency thermal management controller (ETMC) separate from normal DCIM and building automation systems. The ETMC should monitor all pod thermal states and should be capable of initiating emergency ASIC load reduction via out-of-band management network within 90 seconds of a predicted multi-pod thermal event. The ETMC should test its pod thermal isolation capabilities quarterly in coordination with facility operations, with test results reported in compliance documentation. AI workload schedulers used in scale-out CPO clusters should expose a thermal headroom API that the ETMC may query and use to redistribute compute load away from thermally constrained pods without interrupting in-flight training jobs.

10.6 CPO Cooling Compliance, Testing, and Future Roadmap

All CPO cooling installations of single-chassis, scale-up, and scale-out systems should undergo a structured commissioning process before being placed into production operation. The following commissioning tests are recommended:

- **Coolant flow verification:** Confirm design flow rate ($\pm 5\%$) on all cooling loops with calibrated flow meter independent of chassis instrumentation.
- **CPO temperature setpoint accuracy:** Load chassis to 100% ASIC TDP and verify CPO OE temperature remains within $\pm 2^\circ\text{C}$ of setpoint for 30 continuous minutes.
- **Thermal transient test:** Execute 0% to 100% ASIC TDP step change and verify CPO OE temperature recovers to within $\pm 1^\circ\text{C}$ of setpoint within 120 seconds.
- **Coolant chemistry baseline:** Obtain coolant sample from CPO secondary loop and verify conductivity, pH, chlorides, and particulate count comply with Section 17.5.2 requirements.
- **Leak detection functional test:** Simulate a 50mL coolant release at each chassis leak sensor location and verify alarm generation within 60 seconds.
- **Optical link integrity correlation test:** Log CPO OE temperature and optical link BER simultaneously through a thermal transient; verify BER remains below 10^{-12} during all thermal events.

Scale-Up and Scale-Out Commissioning Tests

In addition to per-chassis commissioning tests, scale-up and scale-out installations should complete the following cluster-level commissioning tests:

- **Simultaneous multi-chassis thermal load test:** Load all chassis to 100% TDP simultaneously; verify Zone-CDU supply temperatures remain within 2°C of setpoints for 60 continuous minutes.
- **Pod isolation test:** Isolate one pod's coolant supply and verify remaining pods maintain thermal setpoints within 5% flow deviation for 120 seconds.
- **DCIM telemetry completeness verification:** Confirm all required monitoring points (Section 18.5.1 or 19.5) are reporting to DCIM at required intervals with $\leq 5\%$ data loss over a 24-hour test period.
- **Spine switch thermal qualification:** Load optical fabric spine switches to rated throughput and verify CPO OE temperatures comply with CPO requirements while chassis chassis are simultaneously at full ASIC TDP.

Ongoing CPO Cooling Compliance Monitoring

Parameter	Monitoring Frequency	Alarm Threshold	Reporting
CPO OE temperature (per chassis)	Continuous (≤ 1 sec)	Setpoint $\pm 2^\circ\text{C}$ warning; $\pm 5^\circ\text{C}$ critical	Monthly summary; all exceedances logged
CPO secondary loop conductivity	Continuous	$< 2 \mu\text{S/cm}$ alarm	Quarterly trend report; chemistry change events
Coolant pH (CPO loop)	Weekly (inline sensor) + monthly lab	< 6.8 or > 7.5	Monthly lab report
Optical link BER vs. temperature	Continuous correlation	$\text{BER} > 10^{-10}$ with OE temp anomaly	Incident report within 24 hours
CDU approach temperature	Continuous (5 min average)	$\geq 2^\circ\text{C}$ above specification	Weekly efficiency report
CPO loop flow rate	Continuous (1 min)	$\pm 5\%$ of setpoint sustained > 2 min	Alert to DCIM; log for maintenance review
Fiber bend radius compliance	Quarterly visual inspection	Any visible kink or radius violation	Immediate corrective action; compliance log

Table 88 – Ongoing CPO Cooling Compliance Monitoring

11. CPO Management Systems and Control

This section defines the software management architecture governing CPO-enabled chassis, both at the individual chassis level and across multi-chassis scale-out topologies. As AI workloads grow in scale, the interconnect fabric connecting accelerators, memory, and storage should evolve beyond pluggable optics toward integrated co-packaged solutions. CPO eliminates the electrical chip-to-module interface, dramatically reducing power consumption and signal latency. However, this architectural shift introduces new requirements for initialization, monitoring, firmware management, thermal control, and fault isolation that existing management frameworks do not fully address.

This document covers the CPO software management reference architecture, anchored by the Common Management Interface Specification (CMIS) as the authoritative chassis-level interface, and extends it to support distributed multi-chassis environments. This section reviews the following domains:

- **Software management stack:** layer definitions from physical optical hardware through to AI workload orchestration APIs.
- **CMIS integration:** use of CMIS 5.x and extensions for CPO-specific management registers, telemetry, and control planes.
- **Lifecycle management:** firmware update, provisioning, diagnostics, and decommissioning workflows.
- **Security:** chain of trust, attestation, and encrypted management channel requirements.

11.1 CMIS as the Chassis-Level Management Interface

The Common Management Interface Specification maintained by the Optical Internetworking Forum, defines a standardized register map, state machine, and communication protocol for managing optical modules. Originally designed for pluggable transceivers, CMIS provides a rich framework for optical component initialization, monitoring, and control that APC adopts and extends for the CPO context. CMIS is designated the mandatory chassis-level management interface in all CPO implementations. The Chassis Management Controller (CMC) implements the CMIS server role and exposes CMIS register pages to upstream management software. CMIS is important for CPO because:

- **Proven ecosystem:** CMIS is already implemented in thousands of datacenter deployments and CPO vendors may leverage existing toolchains.
- **Rich state machine:** CMIS Module State Machine (MSM) and Data Path State Machine (DPSM) map naturally to CPO optical engine initialization sequences.
- **Extensible register space:** The Upper Memory Page architecture allows CPO-specific registers without breaking backward compatibility.
- **Vendor-neutral:** CMIS is an OIF specification and APC follows the same open process.
- **Tooling availability:** CMIS diagnostic and management tools from multiple vendors may be extended to support the industry.

The standard CMIS Module State Machine (MSM) defines states: ModuleLowPwr, ModulePwrUp, ModuleReady, ModulePwrDn, Fault, and ModuleUnplugged. For CPO, the module is never physically unplugged, and the initialization sequence is substantially more complex. Standard CMIS uses I2C as its primary host interface. CPO introduces additional interface options to accommodate the higher management bandwidth and lower latency requirements:

- **I2C (400 kHz or 1 MHz Fast Mode Plus):** Baseline interface for CMC-to-host management communication.
- **UART/SPI:** High-speed debug and diagnostics interface. Optional; recommended for optical engine DSP access during characterization.
- **PCIe MMIO (Management Bar):** For CPO integrated into compute platforms, a PCIe Memory-Mapped I/O management BAR may expose CMIS registers directly to host software. Required for compute-side CPO implementations.
- **Ethernet Management Channel:** For chassis-level management over the datacenter out-of-band network. The CMC implements via gRPC over TLS on a dedicated management VLAN.

11.2 Multi-year CPO Deployment

Firmware update management is a critical lifecycle function given the multi-year deployment horizon of CPO chassis. APC recommends the following firmware update capabilities:

- **In-service firmware update (ISFU):** Optical engine firmware may be updated without interrupting traffic on non-updated engines. The CMC manages a staged rollout within the chassis.
- **Dual-image storage:** Each optical engine DSP stores two firmware images (active and standby). Updates are written to the standby partition and activated on a coordinated cutover event.
- **Rollback capability:** If the new firmware image fails health checks within a configurable observation window (default: 30 minutes), the CMC automatically rolls back to the previous image.
- **Fleet-wide update orchestration:** Rolling updates across all chassis in a domain, ensuring no more than a configurable percentage (default: 10%) of chassis are simultaneously in update state.
- **Cryptographic verification:** All firmware images are signed using Ed25519. The CMC verifies image signatures before writing to flash. Images with invalid signatures are rejected.

Optical Recalibration

CPO optical engines require periodic recalibration to compensate for component aging, thermal drift, and laser wavelength drift. APC recommends the following recalibration framework:

- **Continuous background calibration:** Minor alignment adjustments and modulator bias optimization are performed continuously without traffic interruption, using the CMIS Page 80h vernier controls.
- **Scheduled deep calibration:** Full optical alignment and DSP coefficient optimization is performed during low-traffic periods. Deep calibration requires a 30-60 second data path interruption per optical engine.
- **Predictive calibration triggers:** Monitors per-lane SNR trends and triggers preemptive recalibration when degradation trends indicate impending service impact.

APC recommends a secure decommissioning workflow to ensure that chassis removed from service do not retain sensitive configuration data:

- **Decommission command:** Issues a decommission request to the chassis CMC, which gracefully quiesces traffic and removes the chassis from the topology.
- **Cryptographic erasure:** CMC performs cryptographic erasure of all management keys and device identity certificates.
- **Configuration reset:** All non-factory configuration is erased and chassis returns to factory default state.
- **Attestation revocation:** Notifies the PKI system to revoke the chassis device certificate.

11.3 Security Architecture and Security Design Principles

The security architecture should be designed around five core principles applicable to the full five-to-ten year deployment lifecycle:

- **Zero Trust:** No chassis or management component is implicitly trusted. All management interactions require mutual authentication and authorization.
- **Defense in Depth:** Multiple independent security controls at the hardware, firmware, OS, and application layers.
- **Cryptographic Agility:** Management interfaces support algorithm negotiation to enable migration from current (2027) algorithms to post-quantum cryptography as standards mature.
- **Minimal Attack Surface:** CMC management software runs on a minimal OS footprint with no unnecessary services. All network-facing services are subject to strict input validation.
- **Auditability:** All management operations are logged with cryptographic integrity protection; audit logs may not be deleted by management software.

Each chassis is provisioned at manufacturing time with a unique Device Identity (DevID) conforming to IEEE 802.1AR. The DevID consists of:

- **Initial Device Identity (IDeVID):** Factory-provisioned X.509 certificate signed by the vendor's Certificate Authority. Stored in hardware-protected storage (TPM or HSM equivalent) on the CMC.
- **Locally Device Identity (LDeVID):** Operator-issued certificate provisioned during ZTP enrollment, used for ongoing management plane authentication.
- **Conducts periodic remote attestation of all chassis using the IETF RATS** (Remote ATtestation procedureS) framework. Attestation verifies firmware measurement values against known-good golden measurements published by the optical engine firmware vendor.

All AMI communications use mutual TLS (mTLS) 1.3 or later. gRPC channels carrying telemetry streams use TLS with session resumption to minimize reconnection overhead. CMIS register access via I2C is protected by physical access controls (chassis access logging) and logical access controls (CMC enforces CMIS page access permissions). In-band management traffic is isolated to a dedicated management VLAN with ACL enforcement on all ToR switches.

11.4. Sustaining CPO Management for Five to Ten Years

The five-to-ten year deployment horizon of AI datacenter CPO infrastructure demands that the CPO management architecture be explicitly designed for longevity, not just current requirements. The following architectural decisions have been made specifically to ensure management software does not become an obstacle to hardware evolution:

Architecture Longevity Decisions

- **CMIS extension page architecture:** Use well-defined extension pages rather than modifying CMIS core registers to ensure forward compatibility with future CMIS baseline revisions.
- **HAL abstraction:** The Hardware Abstraction Layer isolates CMC from vendor-specific hardware details, enabling CMC software to run unchanged on next-generation CPO hardware.
- **Versioned APIs with long deprecation timelines:** Northbound API versions are supported for a minimum of three years after deprecation announcement.
- **Pluggable protocol support:** The CMC Northbound Agent supports multiple transport protocols (gRPC, NETCONF, RESTCONF) to accommodate evolution in orchestration system preferences.
- **Stateless control plane:** Services may be upgraded independently without migration of persistent state.
- **OpenConfig model extensions:** Using OpenConfig as the base data model ensures interoperability with the broader network management ecosystem as it evolves.

11.5 AI-Driven Management Extensions

APC acknowledges that AI-driven management will become an increasingly important capability over the deployment horizon of this document. The following framework is provided for AI-assisted CPO management:

- **Telemetry Data Lake:** The Telemetry Service is designed to export raw telemetry to external data lakes (S3-compatible object storage) for offline ML model training.
- **Inference Integration API:** Exposes an Inference Integration API allowing external ML inference services to push policy recommendations and receive feedback on outcomes.
- **Anomaly Detection Hooks:** The Fault Correlation Engine provides structured hooks for external anomaly detection models to inject fault predictions, which are evaluated alongside rule-based fault detection.
- **Closed-Loop Automation:** Policy Engine supports automated action execution triggered by AI model outputs, subject to operator-defined safety guardrails (maximum action rate, human approval thresholds).

Recommended Requirements

APC recommends CPO chassis may implement:

- PCIe MMIO management BAR for compute-side host integration.
- Hardware-protected storage (TPM 2.0 or equivalent) for IDevID and firmware measurement registers.
- Streaming telemetry via gRPC with configurable subscription filters.

Glossary of Terms and Definitions

Adiabatic Taper	A waveguide taper sufficiently long that the optical mode transforms adiabatically between two waveguide geometries without coupling to higher-order modes.
AI Edge Data Centers	Smaller AI edge and remote datacenters are commonly referred to as edge data centers, micro data centers (MDCs), or modular units. These facilities are characterized by their smaller size, often ranging from a single rack to container-sized deployments and are strategically located closer to end-users and data sources to reduce latency and bandwidth usage. They are also known as edge sites or micro facilities, particularly when deployed in urban environments, factory floors, or other remote locations. These smaller data centers are designed to operate autonomously with minimal intervention, often utilizing advanced management software for remote monitoring and maintenance.
Application-Specific Integrated Circuit (ASIC)	Accelerator Processing Unit (encompasses GPU, TPU, NPU, IPU, and future AI accelerators)
Arrayed Waveguide Grating	A wavelength multiplexer/demultiplexer based on an array of optical waveguides with linearly increasing lengths, producing a cyclic wavelength routing function.
Arrhenius Model	An empirical model relating failure rate to temperature via E_a (activation energy): $AF = \exp[E_a/k \times (1/T_{use} - 1/T_{stress})]$
Auto Laser Shutdown	A safety mechanism that quenches the laser output when an open fiber condition is detected, preventing unintended laser exposure
Bow-Tie Fiber	A polarization-maintaining fiber type using boron-doped stress-applying members to induce birefringence.
Chirp	Instantaneous frequency modulation of an optical signal; causes spectral broadening and dispersion-induced penalty in high-speed transmission.
Chip on Board	Optical engine die placed on substrate adjacent to ASIC; short wire-bond or flip-chip electrical interface.
Coherent Detection	Optical detection technique using a local oscillator laser to enable phase and amplitude demodulation; enables QAM formats and high OSNR sensitivity.
Comb Repetition Rate	The spacing between adjacent lines of an optical frequency comb, equal to the free spectral range of the comb-generating resonator.
Common Management Interface Specification	Defines the I2C/MDIO register map for optical transceiver modules.
Computer Room Air Conditioner	A precision air conditioning unit deployed within a data hall
Computer Room Air Handler	A unit that uses chilled water from a central plant to cool supply air.

Glossary of Terms and Definitions (con't)

Co-Packaged Optics	Integration of photonic transceiver components within the same package or in direct proximity (< 20 mm electrical trace) to the host compute die, eliminating faceplate SerDes reach losses.
Co-Packaged Optics Coolant Loop	A dedicated liquid cooling circuit serving only the CPO thermal zone, segregated from ASIC processor cooling loops to enable independent temperature setpoint control.
Co-Packaged Optics Thermal Zone	A thermally isolated volume within a chassis or rack designed to manage the combined heat flux from co-packaged optical engines and their associated processing dies.
Coolant Distribution Unit	A device that conditions and distributes liquid coolant to server-level cold plates or manifolds.
Critical IT Load	The aggregate rated power of all information technology equipment within a data center, excluding facility support systems.
Detachable Fiber Array Unit	Combines the precision of a traditional FAU with the flexibility of modular design, allowing for easier assembly, testing, and maintenance in high-density photonic and optical interconnect systems. It's especially valuable in emerging optical I/O architectures like co-packaged optics.
Direct Liquid Cooling	A cooling method where liquid coolant is circulated in direct contact with or in close proximity to heat-generating components via cold plates or manifolds.
Dragonfly Topology	A network topology for high-radix interconnects that minimizes diameter and bisection bandwidth cost by connecting groups of switches through all-to-all connections.
Economization	The use of ambient environmental conditions (free cooling) to reduce or eliminate mechanical refrigeration energy.
Electro-Optic Modulator	A device using the electro-optic effect (Pockels effect in LiNbO ₃ , plasma dispersion in silicon) to impress data on a CW optical carrier.
Etalon	A Fabry-Pérot cavity used as a wavelength reference filter; in AWC systems, the transmission fringe slope provides the wavelength error signal.
Evanescent Coupling	Optical power transfer between adjacent waveguides via the exponentially decaying field outside the waveguide core; used in directional couplers and ring resonators.
Extinction Ratio	Ratio of optical power level '1' to optical power level '0' in modulated signal; higher ER gives better SNR at the receiver.
Faceplate Interface Plane	The physical reference plane at the front panel of the chassis where optical connectors mate; all loss budgets and alignment specs are defined at this plane
Fiber Array Unit	A precision-aligned multi-fiber assembly that couples an array of optical fibers to the waveguide input/output facets of a PIC.

Glossary of Terms and Definitions (con't)

Fiber Array Unit Receptacle	A mechanical and optical interface designed to receive and securely align a Fiber Array Unit (FAU) with a photonic device, such as a Photonic Integrated Circuit (PIC) or optical engine and is the precision docking interface for a fiber array unit, enabling reliable, low-loss optical coupling between fibers and integrated photonic devices. It's essential in systems requiring modular, scalable, or serviceable optical connections, such as in co-packaged optics and high-density optical interconnects. It ensures precise alignment between the fibers in the FAU and the optical waveguides or grating couplers on the chip, enabling low-loss, high-performance optical coupling.
Flip-Chip Bonding	A die-attachment technique in which the active face of a die is bonded face-down to a substrate using solder bumps or micro-pillars.
Forward Error Correction	Error correction coding applied to the optical data stream.
Free Space Range	The wavelength or frequency spacing between adjacent resonance modes of a resonator: $FSR = c/(n_g \times L)$.
Group Refractive Index	The effective refractive index experienced by a pulse of light in a waveguide, accounting for waveguide dispersion; $n_g = n - \lambda(dn/d\lambda)$.
Heterogeneous Integration	The co-integration of dissimilar materials (e.g., III-V on Si) or separately manufactured chiplets within a single package or on a common substrate.
Hollow Core Fiber	An optical fiber design where light propagates predominantly through an air or vacuum core rather than solid glass, achieving lower latency and non-linearity.
Hot Aisle / Cold Aisle	A rack arrangement methodology that segregates cold supply air from hot exhaust air.
Hybrid Bonding	Wafer-to-wafer or die-to-wafer direct bonding without solder bumps; enables sub-micron interconnect pitch for dense photonic-electronic integration.
Immersion Cooling	A cooling technique in which IT hardware is submerged in a dielectric fluid that absorbs and transports heat.
In-package Optics	Optical engine co-packaged in same multi-chip package (MCP) as ASIC.
Infant Mortality	The early failure period in a component's lifetime, modeled by the decreasing-hazard region of a Weibull distribution; burn-in screens remove infant mortality failures.
Insertion Loss	Optical power loss introduced by a device or connection, expressed in dB. $IL = 10 \log_{10}(P_{out}/P_{in})$.
Integrated Silicon Photonics	Photonics integrated on same die as ASIC or in heterogeneous 2.5D/3D stack.
Kerr Effect	A nonlinear optical effect in which the refractive index of a medium changes proportionally to the squared electric field amplitude; exploited in micro-ring frequency comb generators.
Mach-Zehnder Interferometer	Optical device splitting light into two paths, applying a phase shift to one, and recombining; used as an optical modulator or switch.

Glossary of Terms and Definitions (con't)

Micro-Ring Resonator	A circular optical waveguide resonating at specific wavelengths; used as a wave-length-selective filter, modulator, or switch in Si-Ph PICs.
Micro-Transfer Printing	A wafer-scale technique using elastomeric stamps to pick and place pre-fabricated semiconductor micro-structures from source wafers to target wafers with mi-cron-scale precision.
OpenConfig	A vendor-neutral, community-driven initiative to create open, standards-based YANG data models for network device configuration and telemetry.
Optical Engine	A specialized, compact module that integrates photonic and electronic components to perform electro-optical conversion with high efficiency and low footprint.
Optical Frequency Comb	A laser source generating a spectrum of equidistant frequency lines; provides multi-ple WDM channel seeds from a single source device.
Optical Signal-to-Noise Ratio	Ratio of signal power to optical noise power in a defined bandwidth; primary metric for coherent link performance.
Optical Thermal Interface Material	A thermally conductive, optically transparent or optically isolated material used at the interface between an optical engine and its cooling structure, without introduc-ing optical signal degradation.
Photonic Integrated Circuit	A monolithic chip integrating multiple photonic functions (modulators, detectors, waveguides, filters, couplers) analogous to an electronic integrated circuit but oper-ating on photons.
Photonic Wire Bond	A laser-written 3D polymer waveguide connecting optical chips with sub-microme-ter alignment and <0.5 dB insertion loss per bond.
Polarization-Maintaining Fiber	An optical fiber designed to preserve the polarization state of the launched light, using geometric or stress-induced birefringence.
Power Usage Effectiveness	Ratio of total facility power to critical IT load; lower values indicate higher efficiency.
Probability of Failure on Demand	A reliability metric for safety-instrumented systems
Quantum Dot	A semiconductor nanostructure that confines carriers in all three spatial dimensions, providing delta-function-like density of states and reduced temperature sensitivity of threshold current.
Rack Power Density	The average or peak power dissipation of a single equipment rack, expressed in kilowatts (kW).
Rear-Door Heat Exchanger	A heat exchanger mounted on the rear door of a server rack that captures exhaust heat before it enters the hot aisle.
Relative Intensity Noise	Normalized measure of laser intensity noise; expressed in dB/Hz; directly contrib-utes to receiver SNR floor.
Return Loss	The ratio of incident to reflected optical power at a connection; high RL (> 30 dB) indicates low back-reflection, critical for laser stability.

Glossary of Terms and Definitions (con't)

Semiconductor Optical Amplifier	A forward-biased semiconductor p-n junction that amplifies optical signals via stimulated emission; used as a booster or gate in PICs.
Silicon Photonics	Integration of photonic devices on silicon substrates using CMOS-compatible fabrication; dominant platform for CPO PICs due to wafer-scale manufacturability.
Spot-Size Converter	A tapered waveguide structure that adiabatically transforms the mode field diameter between a laser output ($\sim 3\ \mu\text{m}$) and a silicon photonic waveguide ($\sim 450\ \text{nm}$).
Stranded Capacity	Cooling or power capacity that may not be utilized due to infrastructure mismatches or bottlenecks.
Surface-Activated Bonding	A room-temperature wafer bonding technique in which surfaces are activated by ion/atom beam bombardment, achieving covalent bond strengths without elevated temperature.
Thermal Design Power	The maximum sustained power a processor or accelerator is rated to dissipate under full computational load.
Thermal Resistance	Opposition to heat flow between two points, measured in degrees Celsius per watt ($^{\circ}\text{C/W}$).
Thermoelectric Cooler	Solid-state heat pump using the Peltier effect to actively cool a laser die below ambient temperature; critical for laser temperature stabilization in CPO.
Transimpedance Amplifier	An amplifier converting photocurrent from a photodetector to a voltage signal with high bandwidth and low noise; integral to CPO receiver designs.
Transmitter and Dispersion Eye Closure	A composite penalty metric accounting for modulation non-ideality, chromatic dispersion, and bandwidth limitations; measured in dB.
Water Usage Effectiveness	Annual site water usage divided by IT energy consumption, expressed in liters per kilowatt-hour (L/kWh).
Wavelength Division Multiplexing	Transmission of multiple optical signals on different wavelengths over a single fiber; DWDM uses channel spacings of 100 GHz or less.
Wavelength Selective Switch	A reconfigurable optical filter capable of routing individual DWDM wavelengths to different output ports; key element in scale-out optical switching fabrics.
Weibull Distribution	A statistical distribution commonly used in reliability engineering to model failure time distributions; characterized by shape parameter β and scale parameter η .

Acronyms

AI	Artificial Intelligence
ALS	Auto Laser Shutdown
APC	Advanced Photonics Coalition (the standards body publishing this specification)
APS	Automatic Protection Switching (the optical layer mechanism for rapid traffic restoration on link failure)
ASHRAE	America Society of Heating, Refrigerating and Air-Conditioning Engineers
ASIC	Application-Specific Integrated Circuit
BER	Bit Error Ratio
BGA	Ball Grid Array
BW	Bandwidth
CAPEX	Capital Expenditure
CDR	Clock and Data Recovery
CDU	Coolant Distribution Unit
CLEC	Coupling Loss Error Coefficient
CMC	Chassis Management Controller
CMIS	Common Management Interface Specification
CML	Current Mode Logic
CoSWID	Concise Software Identification Tags
CPO	Co-Packaged Optics
CRAC	Computer Room Air Conditioner
CRAH	Computer Room Air Handler
CTE	Coefficient of Thermal Expansion
DAC	Direct Attach Copper
DC	Data Center
DLC	Direct Liquid Cooling
DLHM	Datacenter Scale Laser Health Monitoring
DSP	Digital Signal Processor
DUT	Device Under Test
DWDM	Dense Wavelength Division Multiplexing
EER	Energy Efficiency Ratio
ELI	External Laser Input
EMI	Electromagnetic Interference
EML	Electro-Absorption Modulated Laser
EOM	Electro-Optic Modulator
ER	Extinction Ratio

Acronyms (cont')

FAU	Fiber Array Unit
FEC	Forward Error Correction
FIP	Faceplate Interface Plane
FRU	Field Replaceable Unit
FSR	Free Space Range
FWHM	Full Width at Half Maximum
GI	Glass Interposer
HACA	Hot Aisle / Cold Aisle
HAL	Hardware Abstraction Layer
HAST	Highly Accelerated Stress Test
HCF	Hollow Core Fiber
ICR	Integrated Coherent Receiver
IL	Insertion Loss
IMF	Intra-Module Fiber
InP	Indium Phosphide (III-V photonic material)
IPO	In-Package Optics
IPU	Infrastructure Processing Unit
ISFU	In-Service Firmware Update
ISP	Integrated Silicon Photonics
KP4 FEC	Clause 91 KP4 Reed-Solomon FEC
kW	Kilowatt
LFSR	Linear Feedback Shift Register
LoC	Laser on Chip
LPO	Linear-drive Pluggable Optics
MAC	Media Access Control Address
MCM	Multi-Chip Module
MCTP	Management Component Transport Protocol
MDC	Micro data centers
MFD	Mode Field Diameter
MPI	Message Passing Interface
MPO	Multi-fiber Push-On connector
MR	Medium Reach
mTLS	Mutual Transport Layer Security
MTP	Micro-Transfer Printing

Acronyms (cont')

MW	Megawatt
NPI	New Product Introduction
NRZ	Non-Return-to-Zero modulation
OBPF	Optical Band-Pass Filter
OCP	Open Compute Project
OE	Optical Engine
OFC	Optical Frequency Comb
OIF	Optical Internetworking Forum
OMA	Optical Modulation Amplitude
OMC	Optical Mezzanine Card
OPEX	Operational Expenditure
OPM	Optical Performance Monitoring
OSNR	Optical Signal-to-Noise Ratio
OTIM	Optical Thermal Interface Material
PAM-4	4-level Pulse Amplitude Modulation
PDN	Power Delivery Network
PFD	Probability of Failure on Demand
PHY	Physical Layer
PIC	Photonic Integrated Circuit
PKI	Public Key Infrastructure
PLC	Planar Lightwave Circuit
PLD	Programmable Logic Device
PMBus	Power Management Bus
PMD	Polarization Mode Dispersion
PMF	Polarization-Maintaining Fiber
PSRR	Power Supply Rejection Ratio
PSU	Power Supply Unit
PTM	Photonic Transceiver Module
PUE	Power Usage Effectiveness
PWB	Photonic Wire Bond
QD	Quantum Dot
QSFP-DD	Quad Small Form-factor Pluggable Double Density
RATS	Remote ATtestation procedureS
RDHx	Rear-Door Heat Exchanger

Acronyms (cont')

RDMA	Remote Direct Memory Access
RIN	Relative Intensity Noise
RL	Return Loss
RMS	Root Mean Square
RoHS	Restriction of Hazardous Substances
RoT	Root of Trust
RSSI	Received Signal Strength Indicator
Rth	Thermal Resistance
SAB	Surface-Activated Bonding
SerDes	Serializer/Deserializer
SFE	Structural Fabrication Error
SI	Signal Integrity
SiPh or SIP	Silicon Photonics
SMF	Single-Mode Fiber
SMF28	Standard Single-Mode Fiber (ITU-T G.652.D)
SNR	Signal-to-Noise Ratio
SOC	System-on-Chip
SSC	Spot-Size Converter
SUDM	Scale-Up Domain Manager
TDEC	Transmitter and Dispersion Eye Closure
TDP	Thermal Design Power
TEC	Thermoelectric Cooler
TFLN	Thin-Film Lithium Niobate
TIA	Transimpedance Amplifier
TIE	Trans Interval Error
TRX	Transceiver
TWD	Technical Working Group
UCIe	Universal Chiplet Interconnect Express
VR	Voltage Regulator
WDM	Wavelength Division Multiplexing
WG	Working Group
WPE	Wall-Plug Efficiency
WUE	Water Usage Effectiveness
YANG	Yet Another Next Generation
ZTP	Zero-Touch Provisioning

14.0 External Standards Alignment

Normative and Informative References

This document is intended to be consistent with and contribute toward the following industry working group activities. Conflicts between this document and final published versions of these standards should be resolved in favor of the external standard; this document will be updated at the next revision cycle. The following standards are referenced normatively. Each referenced standard should apply unless superseded by a subsequent edition:

ASHRAE A2	Thermal guidelines for data processing
ASHRAE TC 9.9	Data Center Cooling Guidelines A1–A4
ASTM C1239	Standard Practice for Reporting Uniaxial Strength Data
CMIS 5.0	Common management interface spec
DMTF DSP0236	MCTP Base Specification
DMTF DSP0266	Redfish Scalable Platforms Management API Specification
DMTF DSP0267	PLDM for Firmware Update
DMTF SPDM v1.1	Security Protocol and Data Model
Telcordia GR-468	Generic reliability assurance requirements for optoelectronic devices
GRPC Protocol	GRPC Remote Procedure Call Framework
GR-63-CORE	NEBS Level 3 Environmental Stress
IEC 60068-2-1/2/6/14	Environmental testing (cold, dry heat, vibration, thermal shock)
IEC 60529	Degrees of protection (IP rating)
IEC 60825-1 Ed.3	Safety of laser products
IEC 61000-4 Series	Electromagnetic compatibility
IEC 61753-1	Fiber optic interconnecting devices - Performance standard
IEC 61754-7	Fiber optic connector interfaces, MPO connector family
IEC 61754-20	LC fiber optic connectors
IEC 62343	Dynamic modules – operations, administration, and maintenance
IEC 62368-1 Ed.3 / UL 62368-1:	Audio/video, IT and communications equipment safety
IEEE 1588-2019 -	Precision time protocol for optical frequency comb timing reference

Normative and Informative References (cont')

IEEE 802.1AR-2018	Secure Device Identity
IEEE 802.3bs	400G PHY & MAC layers
IEEE 802.3ck-2022	100G, 200G, and 400G Ethernet over Single-Mode Fiber
IEEE 802.3df-2024	800G and 1.6T Ethernet
IETF RFC 2119	Key words for use in RFCs to Indicate Requirement Levels
IETF RFC 6241	NETCONF Configuration Protocol
IETF RFC 8040	RESTCONF Protocol IETF
IETF RFC 9334	Remote Attestation procedureS (RATS) Architecture
IETF RFC 9393	Concise Software Identification Tags (CoSWID)
ISO 9022-2	Optics and photonics environmental test methods
ITU-T G.652D, G.654E	Single-mode optical fiber and cable characteristics
ITU-T G.657A2	Characteristics of a single-mode optical fiber cable
ITU-T G.694.1	Spectral grids for WDM applications - DWDM frequency grid
ITU-T G.698.4	Multi-channel seeded DWDM applications
JEDEC HBM3	High Bandwidth Memory 3 Standard
JEDEC JESD22-A104	Temperature Cycling
JEDEC JESD22-A110	Highly Accelerated Stress Test
JEDEC JESD79	DRAM interface standards (for memory-near-optical extensions)
MIL-STD-461G	EMI/EMC requirements
MIL-STD-810H:	Environmental Engineering Considerations and Laboratory Tests
MIL-STD-883 Method 2019	Die shear strength
OCP Baseline Hardware Management	Specification v2.0
OCP Optical Mezzanine Card (OMC)	Open Compute Project CPO form factor specification
OIF CEI-112G-LR-PAM4	Optical Internetworking Forum Electrical Interface Standards
OIF CEI-112G-VSR	Common Electrical I/O 112 Gbps Very Short Reach. Optical
OIF CEI-112G-XSR	Common Electrical I/O 112G XSR specification
OIF CEI-224G	Common Electrical I/O 224G specification
OIF CMIS 5.2	Common Management Interface Specification, Rev 5.2.
OIF-CPO-01.0	Co-Packaged Optics Electrical Interface Implementation Agreement

Normative and Informative References (cont')

OIF-CPO-04.0 (draft)	OIF CPO Thermal and Mechanical Interface Requirements
OpenConfig	Network Configuration Data Models
PICMG ATCA 3.0	(referenced for form factor guidance only)
PMBus Specification Rev 1.3.1	Power Management Bus
Protocol Buffers	Google Protocol Buffer serialization format
REACH Reg. 1907/2006	Chemical registration, Materials compliance
RoHS Dir. 2011/65/EU	Hazardous substances restriction
SEMI E187	Specification for Cybersecurity of Fab Equipment
SFF-8679	OSFP Form Factor MSA
SNIA SFF-8024, SFF-TA-1007 (QSFP-DD)	Pluggable baseline comparison
Telcordia SR-332 Issue 4	Reliability Prediction for Electronic Equipment
TIA-568.3-D	Optical fiber cabling, Fiber plant interface

Thank you



**ADVANCED
PHOTONICS
COALITION**

www.advanced-photonics.org