




ADVANCED
PHOTONICS
COALITION

System Level Architecture Overview for Co-Packaged Optics

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.



Contributing Editors

Advanced Photonics Coalition,
Melissa Kallos, Working Group Facilitator and Editor

Alcyon Photonics, Dominik Rabus, CPO System Level Architecture Focus Group Chairperson

ColorChip, Bryan Hall, CPO Working Group Chairperson

ficonTEC, Rayhane Ghane, Editor

Furukawa Electric, Nasu Hideyuki, Editor

Genuine Optics, David Huff, Editor

Lightera, John George, Editor

Lightera, Yi Sun, Editor

Lightera, Satoshi Koyanagi, Editor

Mesh Semiconductor, Peter Troendle, Editor

MPI Corporation, Lawrence van der Vegt, Editor

Andrei Vankov, Senko, Editor

Hakusan, Tsuchiya Takuhiro, Editor

Hakusan, Ryu (Tatsuki) Kimbara, Editor

Hewlett Packard Enterprise, Mohammad Alavirad, Editor

Resolight, Yoel Polak, Editor

Resolight, Ofer Shapiro, Editor

Senko, Obie Roy, Editor

Senko, Max (Yohei) Omodaka, Editor

TE Connectivity Corporation, Sharon Wang, Editor

Tokyo Ohka Kogyo Co. (TOK) Ltd., Mitsuru Sato, Editor

Qnity Electronics, Hjalti Skulason, CPO High Density Optical Interfaces Focus Group Chairperson

Qnity Electronics, Jim Ryley, Editor

Wus Printed Circuit Co. Ltd., Joe Dickson, Editor

Wus Printed Circuit Co. Ltd., Dwight Mattix, Editor



CPO System Level Architecture and High-Density Optical Interfaces Focus 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.
Dexterials 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
2. System Overview	2
2.1 Current requirements	4
3. Optical Engine	7
3.1 Scope and Purpose	7
3.2 Baseline Optical Engine Reference Architecture (T0)	7
3.3 Three-Year Interface Scaling Roadmap	8
3.4 Phase Descriptions and Enabling Technologies	9
3.4.1 Phase 1 — Lane Densification (T0 to T0+12 months)	9
3.4.2 Phase 2 — WDM Introduction (T0+12 to T0+24 months)	9
3.4.3 Phase 3 — Full-Scale Target (T0+24 to T0+36 months)	10
3.5 Enabling Technology and Verification Summary	10
3.5.1 Wide-and-Slow Scaling Path	11
3.6 Out-of-band fabric control wavelength	11
3.7 Interoperability and Backward Compatibility Requirements	11
3.8 Risk Factors and Dependencies	12
3.9 Phase Gate Conformance Criteria	12
4. Optical Interposer	13
4.1 Purpose, Scope, and Relationship to Existing Standards	14
4.2 Glass Substrate Material Specifications	15
4.3 Through-Glass Via (TGV) Specifications	16
4.4 Optical Waveguide and Coupling Interface Specifications	19
4.5 Thermal, Mechanical, and Reliability Qualification	20
4.6 Qualification and Supplier Documentation Requirements	20
4.7 Three-Year Implementation Roadmap (2026–2029)	20
4.8 Year 1 (2026–2027): Prototype Convergence	21
4.9 Year 2 (2027–2028): Design Rule Lock and Panel-Scale Qualification	21
4.10 Year 3 (2028–2029): Production Readiness	22
4.11 Open Issues and Risks	23
4.12 References and Related Reference-Design Sections	23
5. Electrical/Optical Connection	24
6. Fiber-Array Units	26
6.1 Introduction	26
6.2 Mid-Board Connection	27
6.3 PMF FAU for SiPh-based OCS	28
7. Faceplate Connections	29
7.1 Optical Connectors for Data Ports	29



Table of Contents (cont.)

8. Laser Sources	31
8.1 External Laser Source Small Form Factor Pluggable (ELSFP)	32
8.2 Integrated Laser Source	33
9. Assembly Requirements	34
9.1 Overview and General Principles	34
9.2 Electronic and Semiconductor Assembly	34
9.2.1 Die Placement, Bonding, and Interconnects	35
9.2.2 Underfill, Encapsulation, and Material Control	35
9.2.3 Rework and Repair Control	35
9.3 Optical Assembly and Alignment	36
9.3.1 Optical Attachment and Alignment	36
9.3.2 Cure-Induced Alignment Stability	36
9.4 Assembly Critical-to-Quality Characteristics	36
9.5 Thermal and Mechanical Integration	37
9.5.1 Thermal and Mechanical Assembly	37
9.5.2 Downstream Thermal-Process Compatibility	37
9.6 Inline Verification and Assembly Quality Control	38
9.6.1 Inspection and Test Insertion Points	38
9.6.2 Contamination and Cleanliness Control	38
9.6.3 Assembly-Induced Performance Shift	38
9.7 Traceability, Data Management, SPC, and Change Control	39
10. CPO Test and Measurement	40
10.1 Wafer Level Testing	41
10.1.1 Electrical, Optical, RF, and Thermal Acceptance Criteria	42
10.1.2 Mechanical and Physical Acceptance Criteria	43
10.1.3 Wafer Yield Classification and Disposition	44
10.1.4 Known Good Wafer (KGW) Definition	45
10.1.5 Wafer Release Documentation Requirements	45
10.1.6 Statistical Process Control and Change Management	46
10.2 Die-Level Testing	47
10.2.1 Visual and Mechanical Inspection	49
10.2.2 Electrical Characterization	49
10.2.3 Optical and Electro-Optic Characterization	50
10.2.4 High-Speed RF Characterization	50
10.2.5 Thermal Characterization	50
10.2.6 Reliability Screening and Burn-In	51
10.2.7 Known Good Die (KGD) Qualification and Acceptance	51
10.2.8 Release Documentation and Traceability	51
10.2.9 Statistical Process Control and Change Management	52



Table of Contents (cont.)

10.3	Module-level Testing	53
10.3.1	Objectives and Scope	53
10.3.2	Module-Level Integration Verification	54
10.3.3	Visual, Mechanical, and Assembly Inspection	55
10.3.4	Electrical and Power Characterization	55
10.3.5	Optical and Fiber-Interface Characterization	55
10.3.6	Transmitter and Receiver Characterization	56
10.3.7	High-Speed and End-to-End Functional Testing	56
10.3.8	Thermal Characterization	57
10.3.9	Reliability Qualification and Screening	57
10.3.10	Module Qualification and Acceptance	57
10.3.11	Traceability, Data Management, SPC, and Change Control	58
10.4	System Level Testing	59
10.4.1	System Integration and Initialization	60
10.4.2	Electrical and Power-Delivery Network Verification	60
10.4.3	Optical-Link and End-to-End Data-Path Verification	60
10.4.4	Network and Mission-Mode Validation	61
10.4.5	Thermal Management and Power-Thermal Interaction	61
10.4.6	Robustness, Fault Recovery, and Protection	61
10.4.7	EMC/EMI and Environmental Qualification	62
10.4.8	Qualification, Acceptance, and Quality Control	62
10.4.9	Traceability, Data Integrity, SPC, and Change Control	63
11.0	System Level Architecture Future Requirements	64
11.1	Why the discrete-OE model has a horizon	64
11.2	Candidate future-requirement areas	66
11.2.1	Integration of optics into the chip/chiplet	66
11.2.2	Transition towards optical PCB	67
11.2.3	Internal laser sources reconsidered	67
11.2.4	Thermal management as a first-order architectural constraint	68
11.3	The APC working group roadmap	69
11.4	Outlook	70



List of Figures

Figure 1 - Example of a CPO architecture design.....	2
Figure 2 - The Advanced Photonics Coalition Open Ecosystem.....	3
Figure 3 - Example of a CPO Package coupled to an external laser source	4
Figure 4 - ASIC/FPGA stack	5
Figure 5 - Schematic showing one implementation of CPOs.....	6
Figure 6 - Typical losses in individual components making up a CPO.....	6
Figure 7 - Optical Interposer	13
Figure 8 - Schematic representation of the variables for optical waveguides	24
Figure 9 - Example fiber layout for co-packaged module design	27
Figure 10 - Examples of Mid Board Connectors.....	28
Figure 11 - Faceplate port requirements	29
Figure 12 - Schematic of an ELSFP coupled to a CPO Package.....	32
Figure 13 - Schematic of an integrated laser source within the CPO package.....	33
Figure 14 - Co-Packaged Optics (CPO) manufacturing and test flow	40
Figure 15 - Wafer-Level Testing Systems.....	46
Figure 16 - Die-Level Testing systems.....	52
Figure 17 - Module Level Testing Characteristics	58
Figure 18 - System-Level Testing.....	63
Figure 19 - Future CPO system architecture and requirement boundaries.....	65
Figure 20 - Future-state CPO package architecture	68



List of Tables

Table 1 - Optical Engine Interface Scaling Roadmap	8
Table 2 - Enabling Technologies and Verification Focus by Phase	10
Table 3 - Glass substrate material specification.....	15
Table 4 - Through-glass via (TGV) specification parameters and values	16
Table 5 - Optical waveguide coupling interface specification parameters and values	17
Table 6 - Signal integrity requirements.....	18
Table 7 - Thermal and mechanical reliability requirements of glass interposers.	19
Table 8 - Interposer production roadmap	22
Table 9 - List of variables defining the spatial arrangement waveguides - PIC surfaces	24
Table 10 - List of variables defining the spatial arrangement waveguides - interface of the OE	25
Table 11 - List of variables -- for span between two connections of the optical link.....	26
Table 12 - Faceplate Port Granularity	30
Table 13 - Electrical, Optical, RF, and Thermal Acceptance Criteria.....	42
Table 14 - Die Level Testing Characteristics	48
Table 15 - Module Level Testing Characteristics.....	54
Table 16 - System-Level Integration Verification	59
Table 17 - Proposed future system-level requirements set	66



1. Introduction

This document shall define hardware architectures for co-packaged optics (CPO) chassis and integrated systems, intended to support interoperability discussion, vendor alignment, and platform re-use analysis. It serves as a guide for organizations and developers to ensure their solutions meet the demands of current and emerging co-packaged optics applications. In addition, this document shall define the requirements that any optical link connectivity solution must satisfy, both now and in the foreseeable future. Optical connections are required to bring data into CPO based switches in data centers and AI servers, as well as to route data between CPOs, XPUs, and out of the server blade. Today, these connections are typically made using fiber array units (FAUs). As system requirements grow to hundreds or thousands of optical connections per package, current FAU based approaches will not scale. The physical dimensions of optical fibers limit connection density and even thinning or removing the cladding on FAUs only provides some pitch reduction but not enough to achieve the target dimensions. Consequently, they consume excessive beachfront and create significant routing challenges within the confined volume of a server chassis. Achieving higher input/output density will require smaller waveguides with reduced spacing and potentially 2D or 3D interconnect layouts. New technical solutions are needed to support the required high density optical connectivity.

This document is structured as follows:

A general system overview is provided in chapter 2, followed by the individual CPO building blocks, optical engine, optical interposer, electrical/optical connections, fiber array units, faceplate connections, external laser sources. The document closes with a system level architecture future requirements chapter.

2. System Overview

A CPO based system performs the same tasks as a pluggable-based system, routing data between the ASIC / FPGA and the faceplate, that connects to other systems. A CPO based system integrates, or co-packages, the ASIC / FPGA and the optical engine onto a single substrate. Solutions need to co-package the ASIC / FPGA 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 switch. Capabilities for electrical, optical, and thermal cooling, and packaging designs will depend on the design choices for various solutions. An example of a generic CPO switch architecture is shown in the figure below.

- A** CPO module with ASIC and fiber-optic connectors
- B** ELSFP laser modules
- C** Passive fiber-optic connectors
- D** Faceplate
- E** Fiber-optic cables




Figure 1 – Example of a CPO architecture design based on external laser small form factor (ELSFP)

The individual building blocks with their respective components are described in the following chapters including their interoperability and common interdependent variables such as for example the internal loss budget and why this matters.

The Advanced Photonics Coalition open ecosystem supplies these building blocks and components to enable different CPO designs for current and future demand.




Figure 2 – The Advanced Photonics Coalition Open Ecosystem as of Sept. 2026



2.1 Current Requirements

The figure below shows the optical path of a single optical link from an external laser source via the fiber array unit to the photonic integrated circuit / optical engine to the electronic integrated circuit to the ASIC / FPGA.




Figure 3 – Example of a CPO Package coupled to an external laser source.




Figure 4 – ASIC / FPGA stack.

A simplified optical link shall be defined as an ELS-OE connected to an ELS-OE. The major elements of a CPO link ASIC / FPGA-to-switch (switch-to-ASIC / FPGA is the same in reverse) consist of an ASIC / FPGA tray, holding the following:

1. External laser source form-factor pluggable (ELSFP);
2. ELSFP connector
3. Polarizing maintaining fiber (PMF);
4. Detachable connector or non-detachable connector;
5. Optical engine transmittance;
6. Detachable connector or not-detachable connector (same as #4)
7. Singlemode fiber (SMF);
8. Faceplate optical port connector
9. Custom fiber routing and layout may potentially add to the power budget (for example via mid board connectors)




Figure 5 – Schematic showing one implementation of CPOs and individual components.

The end-to-end link budget is a major challenge where trading off power, performance, and cost is necessary and vendor specific. An example of relative losses in individual components and how these losses stack up to yield an overall loss in a system are shown in the figure below.




Figure 6 – Typical losses in individual components making up a CPO. Total loss from laser-die inside an ELS and to the faceplate after modulation can reach upward of -20 dB

3. Optical Engine

One of the primary challenges of co-packaged optics is the large amount of fiber that connects directly to the CPO substrate. For example, CPO fiber counts for 51.2 Tbps bandwidth would require hundreds to over 1000 fibers. Fibers may be different lengths and have optical connectors attached to the ends, including MT-type multi-fiber connectors. End users will not have easy access to this fiber since these connectors are on-board and behind the front panel. Although there are many optical factories trained in handling on-board fiber, connecting and routing the required number of fibers to an ASIC is a new challenge.

Design safeguards to the optical engine (OE) interface are recommended to allow for replacement of components, thus impacting the yield and reliability of the entire CPO switch. It is critical that the components are designed such that they can be physically separated from the CPO. As shown in the following paragraphs, this separation can be achieved electrically, optically, or both depending on design needs.

3.1 Scope and Purpose

This section specifies the architecture of the optical engine within a co-packaged optics (CPO) module and defines a three-year technical roadmap for scaling the total number of optical interfaces the engine presents at the package shoreline. It establishes the baseline configuration, the phased architectural changes required to achieve the scaling target, the enabling technologies each phase depends on, and the verification activities that member companies shall use to demonstrate conformance at each milestone. This roadmap is normative for planning purposes and informative with respect to specific implementation choices; alternative implementations that meet the stated interface-count and interoperability targets at each milestone date are permitted.

The roadmap begins at the publication baseline, designated T0 (2026), and extends through T0+36 months (2029), organized into three annual phases plus the baseline.

3.2 Baseline Optical Engine Reference

At baseline, the optical engine is defined as the smallest replaceable optical-electrical conversion unit within the CPO module, comprising a laser source (integrated or co-located), modulator or receiver array, and a fiber-attach interface presenting discrete optical lanes to the package shoreline. The baseline engine implements 8 optical lanes, each carrying a single wavelength in the O-band at a per-lane rate consistent with existing lane electrical interface specifications, with a 1:1 mapping between optical lane and physical fiber. A baseline package could for example integrate 8 such engines, yielding 64 total optical interfaces and an equal fiber-shoreline count of 64 fibers.

For the purposes of this section, an “optical interface” is defined as one logical optical I/O channel — the combination of one fiber lane and one wavelength carried on it — measured at the package shoreline. This metric is distinct from the physical fiber count, which may scale independently once wavelength-division multiplexing (WDM) is introduced in Phase 2.

3.3 Three-Year Interface Scaling Roadmap

The table below summarizes the phased scaling path from baseline to the Phase 3 target. The roadmap achieves a 16x increase in total optical interfaces per package over 36 months through two independent scaling levers applied in sequence: first, lane densification within the existing single-wavelength architecture (Phase 1); second, introduction of wavelength multiplexing to decouple interface growth from fiber-shoreline growth (Phases 2 and 3). This sequencing was selected because lane densification can be implemented with incremental changes to existing driver, TIA, and fiber-attach designs, while WDM introduction requires new mux/demux and laser-source qualification work that benefits from an additional development cycle.

Phase	Milestone Date	Engines / Package	Lanes / Engine	λ / Lane	Optical Interfaces / Engine	Optical Interfaces / Package	Fiber Shoreline / Package
Phase 0 — Baseline	2026 (T0)	8	8	1	8	64	64
Phase 1 — Lane Densification	2027 (T0+-12mo)	8	16	1	16	128	128
Phase 2 — WDM Introduction	2028 (T0+-24mo)	12	16	2	32	384	192
Phase 3 — Full Scale Target	2029 (T0+-36mo)	16	16	4	64	1024	256

Table 1 – Optical Engine Interface Scaling Roadmap

Optical Interfaces / Package = Engines/Package $\times$ Lanes/Engine $\times$ λ /Lane. Fiber Shoreline / Package = Engines/Package $\times$ Lanes/Engine $\times$ (1 / fiber-sharing factor), where the fiber-sharing factor equals λ /Lane once WDM is active. Note that fiber shoreline count grows more slowly than interface count from Phase 2 onward, and the ratio of interfaces to fibers reaches 4:1 at Phase 3.



3.4 Phase Descriptions and Enabling Technologies

3.4.1 Phase 1 — Lane Densification (T0 to T0+12 months)

Phase 1 doubles the lane count per optical engine from 8 to 16 while preserving the single-wavelength, 1:1 lane-to-fiber mapping. This phase requires no new photonic multiplexing technology and is intended to be achievable using an evolution of baseline driver, transimpedance amplifier (TIA), and fiber-attach designs. Package-level engine count remains unchanged at 8, yielding 128 total optical interfaces per package.

Engine redesign: driver and TIA arrays extended to 16 channels within the existing engine footprint, requiring channel pitch reduction and updated thermal shunt design.

Fiber attach: fiber-array unit pitch reduced to accommodate 16 lanes per engine without increasing engine footprint; pull-strength and epoxy process requalified at the new pitch.

Backward compatibility: shoreline connector footprint and keying remain compatible with baseline modules; Phase 1 packages shall be electrically and mechanically interoperable with baseline host boards.

3.4.2 Phase 2 — WDM Introduction (T0+12 to T0+24 months)

Phase 2 introduces two-wavelength (dual- λ) multiplexing per fiber lane and increases engine count per package from 8 to 12. Each of the 16 lanes per engine now carries 2 wavelengths, so the optical-interface count per engine rises from 16 to 32, and total package interfaces rise to 384, while the fiber-shoreline count increases only from 128 to 192 fibers because two logical interfaces now share each fiber.

Multiplexing: integration of an on-engine 2-channel multiplexer/demultiplexer (ring-resonator or arrayed-waveguide-grating based) between the laser/modulator array and the fiber-attach interface.

Light source: dual-wavelength laser source qualified for the required channel spacing and thermal-drift tolerance; wavelength locking or thermal control added at the engine level.

Connector evolution: shoreline fiber connector updated to a WDM-capable variant; the coalition shall publish an interoperability addendum defining the shared-fiber connector keying before this phase is implemented industry-wide.

Package integration: engine count per package increased from 8 to 12, requiring a revised package floorplan and an updated thermal budget to accommodate the additional engines.

3.4.3 Phase 3 — Full-Scale Target (T0+24 to T0+36 months)

Phase 3 completes the transition to four-wavelength (WDM4-class) multiplexing per fiber lane and increases engine count per package from 12 to 16. This yields 64 optical interfaces per engine and 1,024 total optical interfaces per package, while the fiber-shoreline count grows only to 256 fibers — a fourfold decoupling of interface density from physical fiber count relative to Phase 1.

Full WDM integration: on-engine mux/demux extended to four channels per fiber lane, with insertion-loss budget re-verified across the full channel plan.

Disaggregated laser source: adoption of an external laser source (ELS) architecture to remove laser heat load from the engine and enable hot-swappable light supply, improving package-level thermal headroom at the higher engine count.

Shoreline standardization: the WDM shoreline connector introduced in Phase 2 is finalized as the standard CPO optical shoreline interface for all Phase 3 and later conformant modules.

Serviceability: ELS hot-swap procedures and failover behavior are specified so that light-source replacement does not require package removal or interrupt in-flight traffic on unaffected channels.

Phase	Milestone Date	Engines / Package	Lanes / Engine	λ / Lane	Optical Interfaces / Engine	Optical Interfaces / Package	Fiber Shoreline / Package
Phase 0 — Baseline	2026 (T0)	8	8	1	8	64	64
Phase 1 — Lane Densification	2027 (T0+12mo)	8	16	1	16	128	128
Phase 2 — WDM Introduction	2028 (T0+24mo)	12	16	2	32	384	192
Phase 3 — Full Scale Target	2029 (T0+36mo)	16	16	4	64	1024	256

Table 2 – Enabling Technologies and Verification Focus by Phase

3.5.1 Wide-and-Slow Scaling Path

Alongside the narrow-fast progression described above, each nominal high-rate (e.g., 100G/200G) lane may alternatively be implemented as a wide-and-slow aggregate of N lower-rate sub-channels (e.g., CWDM, DWDM etc.) whose aggregate throughput equals that of the original lane. The per-sub-channel rate and modulation format are implementation choices and are not limited herein.

Reducing the per-channel rate in this manner relaxes the associated driver, modulator, and equalization requirements, at the cost of a proportionally larger channel count. In the wide-and-slow approach the N sub-channels are carried on distinct wavelengths multiplexed onto the same fiber, so that the per-lane wavelength count of any given phase scales by the factor N while the physical fiber-shoreline count is preserved. For instance, with $N = 4$ the two-wavelength scheme of Phase 2 becomes an eight-wavelength scheme, and the four-wavelength scheme of Phase 3 becomes a sixteen-wavelength scheme, with a corresponding increase in multiplexer/demultiplexer channel count and a re-verification of the per-package insertion-loss budget.

3.6 Out-of-band fabric control wavelength

When incorporated in OCS-based network, a CPO interface may reserve at least a single dedicated control-wavelength which is WDM-separable from the data channels. The dedicated control channel is used to schedule transmissions through the OCS and manage the operation of the CPO being part of an OCS-based system. The data channels and control channels may be distinct in transport, timescale, and processor path.

3.7 Interoperability and Backward Compatibility Requirements

Each phase shall preserve mechanical and electrical shoreline compatibility with the immediately preceding phase's host-board interface, except for the fiber connector change introduced at Phase 2, which shall be published as a coalition interoperability addendum no later than three months before Phase 2's milestone date.

Modules conformant to an earlier phase shall remain deployable in host systems designed for that phase for a minimum of 5 years following the introduction of the next phase, to support mixed-generation deployments during transition windows.

Any deviation from the lane, wavelength, or engine counts specified in Table 1 shall be documented against the equivalent total optical-interface and fiber-shoreline count for the target phase, so that conformance can be assessed on an interface-count basis rather than a specific implementation basis.

3.8 Risk Factors and Dependencies

The following dependencies and risk factors have been identified as material to achieving the roadmap on schedule and shall be tracked by the working group through each phase gate.

Laser and mux/demux qualification lead time: dual- and quad-wavelength source and multiplexer qualification (Phases 2 and 3) typically require longer reliability qualification cycles than electrical channel-count increases; schedule contingency should be planned into the Phase 2 milestone.

Thermal density: increasing engine count per package from 8 to 16 while adding wavelength-locking control raises per-package thermal load; the Phase 3 external laser source architecture is a direct mitigation for this risk and should not be deferred.

Test time scaling: per-wavelength test requirements at Phases 2 and 3 increase production test time per package; parallel wavelength-domain test methods should be qualified ahead of the Phase 2 milestone to avoid becoming a manufacturing bottleneck.

Connector ecosystem readiness: the Phase 2 WDM shoreline connector depends on multi-source agreement across connector suppliers; the coalition's interoperability addendum is the pacing dependency for this phase.

Assembly yield at increased lane density: Phase 1's reduced fiber-array pitch may initially reduce fiber-attach yield; process requalification data should be reviewed at the Phase 1 gate before committing to Phase 2 package floorplans.

3.9 Phase Gate Conformance Criteria

A module or package is considered conformant to a given phase only when it meets all of the following at that phase's milestone date: (a) the total optical-interface count specified in Table 1 for that phase; (b) the fiber-shoreline count at or below the value specified in Table 1; (c) the verification activities listed in Table 2 for that phase and all prior phases; and (d) the interoperability requirements of Section 3.4. Conformance shall be demonstrated through the coalition's standard test report format prior to a module being designated as Phase 1, Phase 2, or Phase 3 conformant.

4. Optical Interposer

A Three-Year Implementation Roadmap (2026–2029) for the CPO Reference Design

This section defines the technical specifications, qualification methods, and standards-alignment path for glass interposers used as the electro-optical packaging substrate in the co-packaged optics (CPO) reference design. It is scoped to guide working-group members, module vendors, and substrate suppliers through a coordinated three-year path from prototype-grade glass interposers to a qualifiable, second-sourced production specification suitable for volume CPO deployment in switch and accelerator packages.




Figure 7 – Optical Interposer



4.1 Purpose, Scope, and Relationship to Existing Standards

Glass interposers are addressed here as a distinct packaging substrate class, separate from organic and silicon interposer specifications covered elsewhere in this reference design. Glass offers a combination of properties that make it attractive for CPO: a coefficient of thermal expansion (CTE) that can be tuned close to silicon, low and highly uniform dielectric loss at frequencies relevant to 100+ Gbps electrical signaling, high dimensional stability across large panel formats, and — uniquely among interposer materials — the ability to host optical waveguides directly in the substrate alongside through-substrate electrical vias. This last property allows the interposer to serve simultaneously as electrical redistribution layer and optical routing layer, which is the core motivation for its inclusion in the CPO reference design.

This section does not attempt to define a new standard. Instead, it consolidates applicable requirements from existing and emerging bodies and identifies the gaps this working group must close over the next three years. Primary standards bodies and specification efforts referenced throughout this section include:

- **Optical Internetworking Forum (OIF)** — CPO framework documents and the 3.2T/6.4T co-packaged module implementation agreements, which define optical and electrical interface requirements at the module boundary but do not yet specify interposer material properties.
- **JEDEC JC-11/JC-15** — mechanical, thermal, and reliability test methods for advanced packaging substrates, applicable to glass core qualification by reference.
- **IEC/SEMI panel-level packaging (PLP) working groups** — panel size, warpage, and handling standards relevant to glass panel processing.
- **UCle Consortium** — die-to-die electrical interface specifications that constrain interposer routing pitch and channel loss budgets when UCle is used across the interposer.
- **Open Co-Packaged Optics Collaboration and OCI-MSA** — multi-source agreements defining optical engine footprints and fiber-attach interfaces that the glass interposer must accommodate at its edge or surface.
- **IPC** — laminate and via reliability test methods adapted for through-glass vias (TGVs) pending a dedicated IPC glass-substrate specification.

As of mid-2026, no single body has published a complete, ratified glass-interposer material and process specification. This section therefore presents target specifications derived from published supplier and consortium data, flagged where values remain provisional pending working-group ratification.

4.2 Glass Substrate Material Specifications

The following material targets apply to the glass core used as the CPO interposer. These values represent the working group's consolidated target window, informed by current supplier roadmaps (borosilicate and alkali-free alumino-silicate glass formulations) and reliability data published through 2026. Values are given as a target range to accommodate multiple qualified suppliers rather than a single-source specification.

Property	Target Specification	Test Method	Priority
Panel format	510 mm × 515 mm (Gen 3 / rectangular panel); 300 mm round wafer format retained for early prototypes	Panel metrology per SEMI PLP guidelines	Must-have Y1
Glass thickness	0.1–0.3 mm core, application-dependent; 0.2 mm nominal for switch ASIC packages	Contact/non-contact thickness gauge	Must-have Y1
CTE (25–260°C)	3.0–5.5 ppm/°C, tunable to match adjacent silicon/organic layers within ±1 ppm/°C	TMA per ASTM E228 (adapted)	Must-have Y1
Dielectric constant (Dk) @ 20 GHz	≤ 5.0, target 4.6–4.8	Split-post dielectric resonator or SPDR	Must-have Y1
Dissipation factor (Df) @ 20 GHz	≤ 0.002	SPDR / ring resonator	Must-have Y1
Surface roughness (Ra)	≤ 0.5 nm on optical routing regions; ≤ 2 nm elsewhere	AFM	Must-have Y2
Total thickness variation (TTV)	≤ 5 μm across panel	Capacitive/optical panel scanner	Must-have Y2
Warpage (post-anneal)	≤ 60 μm bow across 510 mm panel	Shadow moiré	Should-have Y2
Optical transmission loss (waveguide core, O-band)	≤ 0.1 dB/cm target; ≤ 0.3 dB/cm acceptable for Y1 prototypes	Cutback method	Must-have Y2
Young's modulus	70–75 GPa	Nanoindentation / resonance method	Should-have Y1

Table 3 – Glass substrate material specification.

Priority tags follow reference-design convention: "Must-have" items gate qualification sign-off for the design phase indicated (Y1 = Year 1, Y2 = Year 2); "Should-have" items are recommended but not blocking.

4.3 Through-Glass Via (TGV) Specifications

TGVs provide the electrical redistribution path through the glass core and are the single largest driver of yield and cost in glass interposer processing. Two TGV formation routes are in active use across the supply base — laser-induced deep etching (LIDE) and laser-assisted electrochemical etching — and both are treated as qualifiable within this specification provided they meet the geometric and electrical targets below.

Parameter	Year-1 Target (Prototype)	Year-3 Target (Production)
Via diameter (top)	60–80 μm	20–35 μm
Via pitch (min.)	150 μm	40–60 μm
Aspect ratio	$\geq 3:1$	$\geq 8:1$
Via taper / sidewall angle	$\leq 5^\circ$ from vertical	$\leq 2^\circ$ from vertical
Via fill metallization	Cu paste / partial fill acceptable	Full Cu electroplate fill, void-free
Via resistance (per via)	$\leq 50 \text{ m}\Omega$	$\leq 15 \text{ m}\Omega$
Insertion loss @ 56 GBd Nyquist	$\leq 1.0 \text{ dB}$ per via transition	$\leq 0.5 \text{ dB}$ per via transition
Via-to-via electrical isolation	$\geq 30 \text{ dB}$ crosstalk suppression	$\geq 40 \text{ dB}$ crosstalk suppression

Table 4 – Through-glass via (TGV) specification parameters and values

Electrical characterization of TGV structures should be reported using de-embedded S-parameters from 10 MHz to at least 67 GHz, consistent with the measurement bandwidth conventions used in OIF electrical implementation agreements for 100G-per-lane and 200G-per-lane signaling. Vendors are expected to supply Touchstone (.s2p/.s4p) files for each qualified via geometry as part of the production part approval process (PPAP) package described in Section 7.7.



4.4 Optical Waveguide and Coupling Interface Specifications

In-Glass Waveguide Requirements

Two integration architectures are supported under this reference design, reflecting the divergent approaches already demonstrated across the supplier base: (a) direct laser-written or ion-exchange waveguides formed within the glass core, and (b) discrete planar lightwave circuit (PLC) or silicon-photonics chiplets surface-mounted onto a glass redistribution layer that carries only passive optical routing between chiplet and fiber-attach interface. Both architectures share the coupling and reliability requirements below; only the waveguide-loss and fabrication tolerances differ.

Parameter	Direct In-Glass Waveguide	Surface-Mounted PIC + Passive Routing
Propagation loss	≤ 0.1 dB/cm (target), 0.3 dB/cm (Y1 floor)	≤ 0.05 dB/cm (glass routing only)
Waveguide pitch at edge coupling	127 μ m / 250 μ m (fiber-array standard pitches)	127 μ m / 250 μ m
Mode field diameter matching	± 0.5 μ m to standard SMF-28 at O-band	± 0.5 μ m
Polarization dependent loss	≤ 0.2 dB	≤ 0.2 dB
Waveguide crossing loss (per crossing)	≤ 0.05 dB	N/A (routing layer only)
Coupling loss, waveguide-to-fiber	≤ 1.0 dB per interface (Y1); ≤ 0.5 dB (Y3)	≤ 0.5 dB per interface

Table 5 – Optical waveguide coupling interface specification parameters and values.

Fiber Attach and Alignment

Fiber-attach interfaces must be compatible with detachable and permanently bonded connector schemes under evaluation by the OCI-MSA and Open CPO collaboration groups. The reference design requires that any glass interposer submitted for qualification support both of the following attach classes, since the working group has not converged on a single mandatory approach as of this writing:

- Fixed (epoxy or solder-bonded) fiber-array unit (FAU) attach, with passive alignment tolerance of ± 1 μ m and active alignment fallback for first-article qualification.
- Detachable expanded-beam or lensed connector attach, targeting field-replaceable optical connections at the package edge for serviceability in switch line cards.



Alignment fiducials for both classes must be defined in the glass at the panel level (not added post-dicing) so that photonic and electrical assembly can share a common coordinate reference frame across the panel.

Electrical Design Rules and Signal Integrity Requirements

Where the glass interposer also carries high-speed electrical signaling (e.g., SerDes lanes between the switch ASIC and co-packaged optical engines, or UCle die-to-die links), the following channel budget applies through the interposer segment only — end-to-end link budgets are defined in the SerDes and optical-link sections of this reference design.

Signaling Class	Max Insertion Loss (interposer segment)	Target Reach
112G PAM4 (56 GBd)	≤ 6 dB at Nyquist (28 GHz)	≤ 25 mm
224G PAM4 (112 GBd)	≤ 4 dB at Nyquist (56 GHz)	≤ 15 mm
≤ 50 G NRZ (25/32/50 GBd)	≤ 6 dB at Nyquist (28GHz)	≤ 25 mm
UCle-Advanced (die-to-die)	Per UCle channel compliance mask, interposer contribution $\leq 30\%$ of budget	≤ 10 mm

Table 6 – Signal integrity requirements.

Return loss at all interposer-level connectors and via transitions should be maintained at ≥ 12 dB across the Nyquist bandwidth of the fastest supported signaling class. Crosstalk budgets between adjacent electrical and optical routing channels must account for the electro-optic nature of the substrate — specifically, verifying that RF signal coupling does not induce measurable phase or amplitude modulation in adjacent optical waveguides (a failure mode unique to glass electro-optical substrates that has no organic-interposer analog).

4.5 Thermal, Mechanical, and Reliability Qualification

Glass is mechanically stiffer and more brittle than organic laminate, which changes the dominant failure modes the working group must screen for. Fracture and crack-propagation behavior, rather than laminate delamination, is the primary reliability concern, particularly at TGV stress concentrations and at the interposer edge where dicing introduces micro-cracks.

Test	Condition / Method	Pass Criterion
Temperature cycling	-40°C to 125°C, 1000 cycles (JEDEC JESD22-A104 adapted)	No TGV opens/shorts; waveguide loss delta $\leq$ 0.05 dB/cm
Highly accelerated stress test (HAST)	130°C / 85% RH, 96 h	No delamination; insertion loss delta $\leq$ 10%
Edge/corner chip resistance	Four-point bend per ASTM C1499 (adapted for panel coupons)	Fracture strength $\geq$ 100 MPa
Panel warpage post-reflow	260°C peak reflow profile	$\leq$ 80 μ m bow
Drop / shock (module level)	JEDEC JESD22-B111	No functional failure
Moisture sensitivity level (MSL)	J-STD-020 classification	MSL 3 or better

Table 7 – Thermal and mechanical reliability requirements of glass interposers.

Fracture toughness and crack-propagation test methods for glass-core interposers are not yet standardized under a dedicated JEDEC or IPC document. The working group's Year-1 deliverable includes a proposal, developed jointly with glass suppliers, to adapt ASTM C1499 and JEDEC mechanical shock methods into a combined test flow specific to electro-optical glass substrates (see Section 7.8, Workstream C).

4.6 Year 1 (2026–2027): Prototype Convergence

Any glass interposer submitted into the CPO reference design flow must be accompanied by a production part approval process (PPAP)-equivalent package. At minimum, this includes:

- Full material data sheet covering all properties in Section 7.2, with lot-to-lot variation data across at least three panel lots.
- TGV process capability report ($C_{pk} \geq 1.33$ on diameter, position, and resistance) across a statistically representative sample per panel.
- De-embedded S-parameter files for all qualified via and routing geometries, from DC to at least 67 GHz.
- Optical loss and coupling reports for all supported waveguide/coupling classes, including wavelength-dependent loss curves across the relevant O-band or C-band window.
- Reliability test reports for all items in Section 7.6, or a documented equivalency claim referencing prior qualification on a materially identical process.
- Panel-level defect density map and yield model, to support the working group's cost-reduction tracking described in Section 7.8.

4.7 Three-Year Implementation Roadmap (2026–2029)

The roadmap below sequences the working group's activities across three overlapping workstreams: (A) specification and standards alignment, (B) supplier qualification and second-sourcing, and (C) reliability methodology development. Dates are targets set at the time of this reference design's publication and are reviewed at each quarterly working-group meeting.



4.8 Year 1 (2026–2027): Prototype Convergence

- Workstream A: Publish this section's provisional specification for working-group ballot; align terminology and test-method references with OIF CPO framework updates and JEDEC JC-11 glass-substrate task group formation.
- Workstream A: Submit joint proposal to JEDEC and IPC to open a dedicated glass-core interposer specification project, using this document as a starting contribution.
- Workstream B: Qualify at least two glass-core suppliers against the Year-1 TGV and material targets (Section 7.2–7.3) on 300 mm and/or sub-panel formats.
- Workstream B: Complete first-article electro-optical test vehicles combining TGV electrical channels and edge-coupled optical waveguides on a single panel.
- Workstream C: Draft combined mechanical/optical reliability test flow for glass electro-optical substrates; run pilot HAST and temperature-cycling data on test vehicles.

4.9 Year 2 (2027–2028): Design Rule Lock and Panel-Scale Qualification

- Workstream A: Ratify electrical design rules (Section 7.5) and optical coupling interface specification (Section 7.4) as stable reference-design requirements; freeze fiber-attach fiducial placement convention.
- Workstream A: Publish crosstalk/electro-optic interference test method (RF-to-optical coupling) as a working-group technical report, feeding into JEDEC/IPC ballot.
- Workstream B: Move qualification to full 510×515 mm panel format; demonstrate TTV, warpage, and TGV yield targets at panel scale across at least two independent suppliers.
- Workstream B: Complete second-source qualification for at least one supplier per major sub-component (glass core, TGV fill process, waveguide formation).
- Workstream C: Finalize fracture/crack-propagation qualification method; incorporate into formal reliability qualification suite; begin field-return data collection from Year-1 prototype deployments.

4.10 Year 3 (2028–2029): Production Readiness

- Workstream A: Converge Year-3 TGV and electrical targets (Section 7.3, right-hand column) into the shipping reference-design specification; retire provisional-status flags pending JEDEC/IPC ratification.
- Workstream B: Demonstrate production-representative yield (target Cpk ≥ 1.33 across all must-have parameters) at panel scale with at least two qualified, second-sourced suppliers per critical process step.
- Workstream B: Complete system-level qualification of at least one full CPO module using a glass interposer meeting the complete Year-3 specification set.
- Workstream C: Publish finalized reliability and test-method specification through JEDEC/IPC; retire interim working-group test methods in favor of the ratified standard.

Phase	Primary Enabling Technologies	Verification Focus
Phase 1	Engine array redesign to 16-lane driver/TIA channel count; denser fiber-array attach (2x pitch reduction); doubled thermal shunt capacity per engine footprint.	Per-lane BER and eye-margin regression at doubled channel density; fiber-attach pull-strength and insertion-loss uniformity across the expanded array.
Phase 2	On-engine integrated multiplexer/demultiplexer (ring-resonator or AWG based, 2-channel); dual-wavelength laser source qualification; shared-fiber WDM connector interface.	Per-wavelength channel isolation and crosstalk; mux/demux insertion-loss budget; thermal-drift wavelength stability across engine operating range.
Phase 3	Full WDM mux/demux integration; disaggregated external laser source (ELS) architecture with hot-swappable light supply; standardized WDM CPO shoreline connector.	Full WDM channel-plan conformance; ELS failover and hot-swap continuity; end-to-end interface count and shoreline density conformance to Section target.

Table 8 – Interposer production roadmap.



4.11 Open Issues and Risks

- Standards gap: no ratified JEDEC or IPC document currently covers glass-core TGV reliability or fracture qualification; the working group's proposed test methods (Section 7.6) remain provisional until formal ballot completes, which may not align with the Year-2 design-freeze target.
- Panel-scale yield: TGV aspect ratios above 6:1 at full 510×515 mm panel scale have limited public qualification data as of this writing; Year-2 and Year-3 targets carry schedule risk if yield does not scale from coupon-level results.
- Fiber-attach convergence: the industry has not converged on a single mandatory fiber-attach class (fixed vs. detachable); carrying both through Year-2 design-rule freeze increases interposer edge complexity and cost.
- Second sourcing: as of mid-2026, glass electro-optical substrate production remains concentrated among a small number of suppliers; the Year-2/Year-3 second-source targets in Workstream B are a primary schedule and cost risk for the overall reference design.
- RF-to-optical crosstalk methodology: this is a reliability and signal-integrity failure mode specific to electro-optical glass substrates with no established measurement standard; Workstream C's Year-1/Year-2 deliverables are on the critical path for electrical design-rule ratification.

4.12 References and Related Reference-Design Sections

- OIF Co-Packaged Optics framework and 3.2T/6.4T module implementation agreements (electrical and optical interface boundary conditions).
- JEDEC JC-11 (mechanical) and JC-15 (test methods) working documents, as extended for glass-core substrates.
- UCle Consortium specification, for die-to-die channel budgets where UCle traverses the glass interposer.
- OCI-MSA and Open Co-Packaged Optics Collaboration fiber-attach and optical-engine footprint specifications.
- Section 4 (Optical Engine and PIC Integration) of this reference design, for end-to-end optical link budget context.
- Section 6 (Organic and Silicon Interposer Specifications) of this reference design, for comparative electrical design rules.
- Section 9 (Thermal and Mechanical Co-Design) of this reference design, for package-level thermal budget interactions with glass core properties.

This section reflects working-group consensus as of the publication date and is expected to be revised at each quarterly review as JEDEC, IPC, and OIF workstreams progress. Values marked in Section 7.2–7.3 as provisional should not be treated as final until the corresponding standards ballot referenced in Section 7.8 closes.

5. Electrical/Optical Connection

Transporting optical signals from the OE to an external fiber or other waveguide remains a key challenge for co-packaged optics. Effective optical coupling must address waveguide size and refractive index mismatches, pitch differences, and the stringent alignment requirements of separable connectors. Pitch adaptation is commonly achieved using intermediate fan-out/fan-in waveguides fabricated from glass or polymer materials. Optical insertion loss between the OE and the external fiber or waveguide should be less than 3 dB, with a target of less than 1.5 dB.

The spatial arrangement of the optical waveguides carrying signals into and out of the optical engine (OE) is a key consideration in the design of the optical interface, as it determines the coupling approach and connector technology that can be used. The figure below illustrates the parameters that define the OE optical interface geometry, with corresponding definitions provided in Table 1. Table 2 defines the optical properties of the photonic integrated circuit (PIC) interface.




Figure 8 – Schematic representation of the variables defining the spatial arrangement of optical waveguides in the PIC that the connector technology must adapt to. A) Waveguide cross-sectional dimensions. B) 1D or 2D array definitions. C) Waveguide length in case of adiabatic coupling. Variables are listed in Table 9.

Variable	Description	Future Recommended Feature size	Future Tolerances
wgh	Waveguide height	TBD	TBD
wgw	Waveguide width	TBD	TBD
wgp	Waveguide pitch	TBD	TBD
wgl	Waveguide length	TBD	TBD
wrh	Waveguide row height	TBD	TBD
wro	Waveguide row offset	TBD	TBD
wgn1, wgn2, etc	Number of waveguides in row	TBD	TBD

Table 9 – List of variables defining the spatial arrangement of optical waveguides on the PIC surface.

Variable	Description	Future Recommended Feature size	Future Tolerances
wlr	Wavelength range	TBD	TBD
wri	Waveguide refractive index	TBD	TBD
spx	Optical signal power, maximum	TBD	TBD
spn	Optical signal power, minimum	TBD	TBD
lpx	Laser signal power, maximum	TBD	TBD
lpn	Laser signal power, minimum	TBD	TBD

Table 10 – List of variables defining the optical properties of the waveguide interface of the OE.

CPO detachable connectors are new technology. Scale-out optics use pluggable transceivers that are too bulky for CPO. CPO detachable connectors need to be very small to fit in the GPU/Switch packages. Detachable connectors are required to ease handling of the multi-chip package (MCP) for compute/switch. Without detachable connectors, dozens of fibers would have to be connected to the switch/compute MCP, which could break during assembly of the compute/switch tray.

6. Fiber-Array Units

6.1 Introduction

The primary purpose of the fiber-array units (FAUs) is to route optical signals between the connectors on both the OE and the faceplate sides, while minimizing both physical volume and optical loss. Today, this connection is typically implemented using fiber array units (FAUs), which consist of optical fiber bundles terminated with appropriate connectors at each end. Single-mode fibers (SMFs) carry data signals between the OE and the faceplate connector and onward outside the chassis. In systems that use external lasers, the FAU may also include polarization-maintaining fibers (PMFs) to deliver optical power from the faceplate to the OE. PMFs preserve the laser polarization state, which is critical for polarization-sensitive OE functions.

The number of fibers in an FAU depends on the OE design, connector technology, and chassis architecture. Common configurations include 16 or 32 fibers, while 64- and 96-fiber FAUs are used in high-density implementations. In systems that employ external lasers, PMFs typically account for one fiber for every eight SMFs, though architectures using a 1:16 PMF-to-SMF ratio are also emerging.

As the number of OE-to-faceplate optical links increases, efficient use of chassis space becomes increasingly important. Key metrics include the cross-sectional area required per optical connection and the average connection length. These parameters are defined in the table below.

Variable	Description	Future Recommended Feature size	Future Tolerances
avl	Average length	TBD	TBD
csl	Cross-sectional area per optical link	TBD	TBD

Table 11 – List of variables defining the requirements for the span between two connections of the optical link.

6.2 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 can 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 in Figure 9.




Figure 9: Example fiber layout for co-packaged module design

For an electrically separable socketed optical engine, fibers in the Figure 9 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 there may be 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. Pigtail length variation brings concerns not only about design but also the risk of fiber damage during installation. The fibers pigtails will need to be routed in tight spaces surrounded by other components. This can 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 can contribute negatively to the overall costs of the switch production. In the case where pigtails would be connectorized, connector positions and size could constrain options for heating or fiber density. Pigtail lengths, if varied, may documentation detailing fiber layouts to prevent unwanted breakage.

Implementing a mid-board optical interconnect solution can 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 can 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 can be easily replaced with another jumper. Yield-affecting risks are shifted to jumper cables from the more expensive OE and ASIC components. The change 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 in the figure below.




Figure 10: Examples of Mid Board Connectors

6.3 PMF FAU for SiPh-based OCS

A CPO connected to a SiPh-based OCS may extend its switch-facing outputs to polarization-maintaining fiber (beyond the CW laser-distribution harness). For these segments the following should apply:

1. Keyed PMF connectors at every mated switch-facing interface, with a defined maximum per-interface angular misalignment ($-/+3^\circ$), and axis-alignment verification as part of assembly.
2. PMF axis preservation through the fiber shuffle: any shuffle carrying fabric-facing fiber should be a polarization-maintaining, keyed assembly with axis continuity from ingress to egress.
3. Source-side launch-SOP alignment onto the PMF axis at the edge device output, and FAU/PIC input-coupling alignment to the switch's supported polarization at each OCS stage.

7. Faceplate Connections

The faceplate can 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 >1 kW systems. For example, as shown in Figure 11, the faceplate of a 51.2 Tbps switch box may be configured as 32 ports with 1.6 Tbps throughput per port, or 128 ports x 400 Gbps per port, with impact on the required panel height (rack real estate) and faceplate surface available for airflow. Basic guidance is provided below on the possible configurations and their influence on thermal management.




Figure 11: Faceplate port requirements

7.1 Optical Connectors for Data Ports

Table 12: Faceplate Port Granularity lists the fiber channel density for various optical connectors, and the corresponding panel height required DR-based and FR-based switch systems, assuming 100Gb/s/lane communication. For ease of discussion, DR-based 51.2 Tbps bandwidth switches were assumed to require 1024 fibers, while FR-based 51.2 Tbps bandwidth switches were assumed to require 256 fibers. Due to the reduced fiber requirements for FR-based front panels, additional connector options could be available for a 1RU 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 1RU switch design.



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	N/A
MMC-16	16	216	3456	N/A
SN-MT16	16	216	3456	N/A
SN-MT32	32	216	6912	N/A
MMC32	32	216	6912	N/A
SN-MT12	12	216	2592	N/A

Table 12: Faceplate Port Granularity

Connector selection is generally driven by density, accessibility and thermal considerations rather than optical loss performance. Most modern single-mode and multimode connector families (LC, SN, CS, MDC, MPO, SN-MT, MMC and related variants) can support insertion losses below 0.35 dB per mated pair, provided that appropriate manufacturing controls and cleanliness procedures are maintained.

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.

The implementation of the optical source, as described in 8.1, 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.



Typical Insertion Loss Across Faceplate Connectors

Optical attenuation through the faceplate connection is one of a key contributor to the overall link loss budget. The optical path typically consists of a mated connector pair between the internal fiber assembly (FAU) and the external cabling infrastructure. Industry standards such as ANSI/TIA-568.3-E specify a maximum insertion loss of 0.75 dB per mated connector pair for multimode and single-mode structured cabling applications, although modern data center connectors generally perform significantly better than this limit.

Typical insertion loss values for high-performance data center connectors are in the range of 0.1 to 0.35 dB per mated pair, with premium single-mode connectors often achieving ≤ 0.2 dB under controlled manufacturing conditions. MPO and SN-MT style multi-fiber connectors are typically specified at 0.35 dB maximum insertion loss for low-loss grades and 0.75 dB maximum insertion loss for standard grades. Duplex connectors such as LC, SN, CS and MDC generally exhibit similar optical performance when manufactured to equivalent quality levels, with differences primarily related to port density rather than optical attenuation.

For system-level optical budget calculations, a design value of 0.35 dB per faceplate connection is commonly assumed, while 0.2 dB may be appropriate for optimized low-loss implementations. This insertion loss should be included in the end-to-end channel budget together with losses associated with fiber attenuation, internal routing, splices and any additional connector interfaces.

8. Laser Sources

The continuous-wave laser light supplied to the optical modulators can originate from either an integrated laser source (ILS) within the OE or an external laser source (ELS) connected through the faceplate. Historically, ELS architecture has been preferred because III-V lasers have exhibited relatively high failure rates, and a pluggable external module can be replaced more easily than an integrated laser. In contrast, failure of an ILS can degrade optical link capacity or result in complete unit failure, requiring replacement of the entire ASIC/OE package to restore functionality.

The laser operates in the O-band wavelength range of 1260 to 1360 nm. Optical fibers and waveguides can simultaneously carry multiple discrete wavelengths, a capability that is widely used to increase transmission capacity. Wavelength-division multiplexing (WDM) enables multiple optical signals to be transmitted over a single fiber by assigning each signal a unique wavelength. Dense Wavelength Division Multiplexing (DWDM) specifies frequency differences as small as 100 GHz which seems to be the target of many OEMs. Various WDM standards define the supported wavelength channels and channel spacing. Each wavelength channel requires a dedicated laser source operating at the specified wavelength.

Note that for OCS-based networks, the external laser source may be sized to cover the summed switch-fabric insertion loss in addition to the conventional link losses. The excess loss should allow the utilization of standardized ELSFP or Integrated LS.

8.1 External Laser Source Small Form Factor Pluggable (ELSFP)

The ELS delivers laser light to each OE through optical fibers, increasing fiber count and routing density within the system.




Figure 12: Schematic of an ELSFP coupled to a CPO Package

To simplify deployment and maintenance, ELS modules typically use standard pluggable form factors and fiber connectors that can be integrated into the switch faceplate. Each ELS module generally contains multiple lasers, with each laser coupled to a dedicated fiber supplying continuous-wave light to an OE. The most common configuration today supports eight lasers per ELS module.

Thermal considerations also influence the choice between ELS and ILS. O-band III-V semiconductor lasers are typically less than 30% efficient, with the majority of their electrical power dissipated as heat. Because OEs often incorporate temperature-sensitive micro-ring resonators, locating the lasers outside the OE can help reduce thermal loading and improve overall system performance in some designs.

The optical path inside the ELSFP consists of the following components:

- Laser die
- Collimating lens
- Isolator
- Focusing lens
- Mechanical transfer (MT) connector: a multiple-fiber optical connector

The laser light from the ELSFP is focused into a polarizing maintaining fiber (PMF) that carries it to the connector on the OE.

8.2 Integrated Laser Source




Figure 13: Schematic of an integrated laser source within the CPO package.

Integrating laser sources directly within the OE eliminates the need for external laser distribution networks, reducing system complexity and improving integration density. ILS can also operate at lower power levels than ELS, as optical routing losses between the laser and OE are eliminated.

However, incorporating ILS into the OE requires heterogeneous integration of III-V materials (e.g., InP, GaAs) onto silicon wafers, increasing manufacturing complexity and potential yield risk. In addition, lasers are highly temperature-sensitive, and their proximity to a high-power, heat-generating ASIC / FPGA can create thermal management challenges. This is particularly critical in DWDM systems, where precise wavelength control is required.

Despite these challenges, significant advancements in integrated laser technology have been made in recent years, and several foundries now offer ILS-enabled PIC platforms as part of their process portfolios.



9. Assembly Requirements

9.1 Overview and General Principles

The Co-Packaged Optics (CPO) assembly process shall establish the mechanical, electrical, optical, and thermal interfaces required to integrate applicable Photonic Integrated Circuits (PICs), Electronic Integrated Circuits (EICs), substrates or interposers, optical-source and coupling interfaces, and thermal-management elements into the intended optical engine, module, or CPO assembly.

Assembly processes shall operate within defined and qualified process windows and shall not introduce unacceptable degradation of the electrical, optical, thermal, mechanical, or reliability performance established for the constituent components. The assembly sequence shall preserve applicable alignment references, optical surfaces, electrical interfaces, and test access required for subsequent inspection, test, qualification, and system integration.

Assembly requirements shall be architecture-dependent. The applicable assembly sequence, integration technology, materials, process controls, inspection methods, and acceptance criteria shall be established according to the product architecture, identified failure mechanisms, manufacturing capability, and downstream performance and reliability requirements.

9.2 Electronic and Semiconductor Assembly

9.2.1 Die Placement, Bonding, and Interconnects

Die placement, attachment, and interconnection processes shall maintain the positional accuracy, mechanical integrity, and electrical performance required by the applicable interface design. Qualified integration technologies may include flip-chip bonding, thermo-compression bonding, micro-bump interconnection, direct or hybrid bonding, or other applicable die-to-die or die-to-substrate processes.

Placement and bonding capability shall be established relative to applicable interconnect pitch, alignment features, coplanarity requirements, and assembly tolerance budgets. Critical process parameters such as placement position, bonding force, temperature, time, atmosphere, and applicable surface conditions shall be controlled where they influence interconnect integrity or downstream performance.

Electrical interconnections shall satisfy applicable requirements for continuity, isolation, contact resistance, power delivery, and high-speed signal integrity. Assembly-induced parasitics, mechanical stress, voiding, interconnect defects, or other process effects shall remain within qualified acceptance limits.



9.2.2 Underfill, Encapsulation, and Material Control

Underfills, encapsulants, adhesives, cleaning materials, Thermal Interface Materials (TIMs), and other assembly materials shall remain within the qualified material and process baseline.

Material selection and application shall control voiding, contamination, cure-induced displacement, mechanical stress, moisture sensitivity, outgassing, and other effects capable of degrading electrical, optical, thermal, mechanical, or reliability performance.

Materials adjacent to optical interfaces shall not introduce unacceptable contamination, residue, outgassing, obstruction, or optical degradation within active optical regions. Material compatibility shall be evaluated with respect to applicable manufacturing processes, operating environments, and expected product lifetime.

9.2.3 Rework and Repair Control

Rework or repair shall be permitted only where a qualified process has demonstrated that the operation does not compromise the performance or reliability of the reworked component, adjacent components, electrical interconnects, optical interfaces, substrates, dielectric structures, or thermal-management materials.

Where rework introduces additional thermal, mechanical, chemical, or optical exposure, the allowable number and severity of rework cycles shall be defined and controlled according to qualified process limits. Cumulative process exposure shall remain within the demonstrated capability of affected materials and interfaces.

Following rework, all Critical-to-Quality (CTQ) characteristics potentially affected by the operation shall be re-inspected or re-tested using qualified methods. Reworked units shall remain identifiable within the product genealogy, including the affected component or interface, rework operation, applicable process conditions, and resulting acceptance status.



9.3 Optical Assembly and Alignment

9.3.1 Optical Attachment and Alignment

Optical assembly shall establish and maintain the alignment and interface integrity required for the completed optical path. Depending on the CPO architecture, the assembly may include fiber arrays, waveguides, lenses, optical connectors, optical-source interfaces, micro-optics, or other coupling structures.

Alignment may be passive, active, or a qualified combination of both. The alignment process shall provide sufficient positional and angular capability to satisfy applicable optical coupling, channel-uniformity, and interface-performance requirements. Alignment tolerances shall be derived from the optical design and applicable coupling-loss budget rather than from a universal positional requirement. Optical attachment processes shall maintain alignment and mechanical stability throughout subsequent manufacturing operations and qualified operating conditions.

9.3.2 Cure-Induced Alignment Stability

For optical interfaces using adhesives or other cure-dependent attachment processes, the effect of curing on optical alignment shall be characterized during process qualification. Applicable translational and angular displacement, coupling-performance change, and channel-to-channel variation shall be evaluated before and after cure.

Where cure-induced displacement is significant relative to the optical alignment tolerance budget, the assembly process shall incorporate appropriate compensation, process control, or alignment methodology to maintain the completed interface within qualified acceptance limits. Qualification shall consider applicable ultraviolet (UV), thermal, or combined cure processes and subsequent stabilization effects where they can influence final optical alignment, mechanical integrity, or optical performance.

9.4 Assembly Critical-to-Quality Characteristics

Assembly CTQ characteristics shall be established according to the product architecture, integration technology, downstream requirements, and identified failure mechanisms. Applicable CTQs may include die position and orientation, coplanarity, interconnect integrity, bond or joint quality, bond-line thickness, under-fill or adhesive condition, substrate or interposer warpage, optical alignment, coupling performance, channel uniformity, optical-interface cleanliness, and thermal-interface integrity.

Acceptance criteria shall be derived from applicable electrical, optical, mechanical, thermal, manufacturing, and reliability requirements. Measurement and inspection methods shall demonstrate sufficient capability, repeatability, and reproducibility relative to the applicable acceptance limits.

CTQ characteristics requiring 100% inspection or test shall be identified by the qualified production control plan. Statistically justified sampling may be applied to appropriate characteristics where process stability, measurement capability, defect-detection risk, failure criticality, and downstream impact support reduced inspection.



9.5 Thermal and Mechanical Integration

9.5.1 Thermal and Mechanical Assembly

Thermal and mechanical assembly shall maintain the structural integrity and thermal performance of the integrated CPO assembly without introducing unacceptable stress, deformation, or displacement of electrical or optical interfaces. TIMs, heat spreaders, cold plates, stiffeners, lids, or other applicable thermal and mechanical elements shall be installed according to qualified process conditions. Their integration shall not introduce unacceptable die tilt, substrate warpage, optical misalignment, interconnect stress, or other assembly-induced degradation.

Substrate, interposer, and package deformation during bonding, curing, reflow, or other applicable thermal processes shall remain within limits established by the mechanical, electrical, optical, and thermal tolerance budgets. Optical fibers and other mechanically sensitive optical structures shall be routed and retained according to their qualified bend-radius, strain, attachment, and environmental requirements throughout assembly and enclosure integration.

9.5.2 Downstream Thermal-Process Compatibility

The assembly sequence shall account for all subsequent thermal processes to which the CPO assembly or its constituent interfaces will be exposed, including applicable solder reflow, bonding, curing, bake, and other manufacturing thermal cycles. Where optical, electrical, or thermal interfaces are assembled before a downstream reflow or other elevated-temperature process, those interfaces shall be qualified for the applicable temperature profile and cumulative thermal exposure. Post-process verification shall demonstrate that critical alignment, attachment integrity, electrical interconnection, optical performance, and material properties remain within qualified limits.

Where an optical interface, attachment, or material is not qualified for the applicable downstream thermal exposure, the manufacturing sequence shall place the corresponding attachment operation after that thermal process or implement an alternative qualified assembly sequence.

9.6 Inline Verification and Assembly Quality Control

9.6.1 Inspection and Test Insertion Points

Inspection and test insertion points shall be established at appropriate stages of the assembly flow to identify defects before subsequent high-value or non-reworkable integration operations. Constituent dies or chiplets shall have the required qualification status prior to assembly. Pre-assembly verification may include identity confirmation, interface inspection, dimensional or alignment-feature verification, and cleanliness assessment.

Following applicable die-attachment and interconnection processes, qualified inspection or test techniques may include machine vision, electrical continuity testing, X-ray imaging, acoustic microscopy, optical metrology, or other methods appropriate to the targeted failure mechanism. Following optical attachment, applicable optical measurements shall verify the integrity and performance of the completed optical interface at the appropriate assembly test insertion point.

9.6.2 Contamination and Cleanliness Control

Assembly and handling of exposed optical interfaces shall be performed within a controlled environment appropriate to the contamination sensitivity of the optical design and assembly process. Processes shall control particles, residues, fibers, oils, moisture, and other contaminants capable of degrading optical coupling, adhesion, bonding, or long-term reliability. Cleaning methods shall be compatible with the applicable optical surfaces, coatings, adhesives, and assembly materials.

Where a defined cleanroom classification is required, the applicable ISO 14644 cleanliness class and monitoring requirements shall be specified by the qualified manufacturing process. Localized environmental control may be used where demonstrated to provide the required contamination control at the critical assembly interface.

9.6.3 Assembly-Induced Performance Shift

Where a CTQ is measurable both before and after an assembly operation, the change in that characteristic should be monitored to identify assembly-induced degradation independently of absolute pass/fail status. Applicable parameters may include electrical interconnect characteristics, optical coupling performance, alignment-sensitive characteristics, thermal behavior, or other architecture-dependent CTQs.

Expected pre-assembly-to-post-assembly shifts shall be characterized during process qualification. Statistically significant or systematic deviations from the qualified correlation baseline shall trigger investigation according to the applicable process-control plan.

9.7 Traceability, Data Management, SPC, and Change Control

Each assembled CPO unit shall maintain traceability to applicable constituent dies or chiplets, originating wafer and lot information, assembly and material lots, substrates or interposers, optical components, and other critical constituent elements.

Manufacturing records shall include, as applicable, equipment identification, process and recipe revision, material lots, relevant process parameters, inspection results, measured CTQs, rework history, and final assembly disposition. First-pass, retest, rework, and final-disposition records shall remain distinguishable within the Manufacturing Execution System (MES), Quality Management System (QMS), or equivalent controlled data system. Original records shall not be overwritten by subsequent processing or testing.

Selected assembly CTQs shall be monitored using Statistical Process Control (SPC) or equivalent statistical methods where appropriate. Monitoring should support identification of equipment-, material-, process-, position-, and assembly-site-dependent variation, systematic parameter shifts, and recurring failure signatures.

Changes to critical materials, die-attachment or bonding processes, optical alignment or attachment methods, substrates or interposers, thermal interfaces, assembly equipment, process recipes, inspection methods, or other assembly-critical elements shall be evaluated through formal change control. Correlation, regression testing, or requalification shall be performed according to the potential impact of the change on assembly quality, downstream CPO performance, manufacturability, and reliability.

10. CPO Test and Measurement

Co-Packaged Optics (CPO) test and measurement require a coordinated verification strategy across multiple manufacturing and integration stages because electrical, photonic, optical, mechanical, and thermal characteristics become observable at different points in the product flow. Testing is therefore organized across four primary test boundaries:

1. **Wafer-Level Characterization:** Provides early parametric, functional, and process information prior to singulation.
2. **Die-Level Qualification:** Establishes Known Good Die (KGD) status and detects defects or performance shifts associated with singulation, handling, and applicable post-wafer processing.
3. **Module-Level Integration Verification:** Evaluates effects introduced by heterogeneous integration, electrical interconnection, optical attachment and alignment, packaging, and thermal integration.
4. **System-Level Validation:** Verifies the behaviour and interaction of the complete CPO platform under representative operating conditions.

Each stage shall establish the Critical-to-Quality (CTQ) characteristics appropriate to its test boundary while maintaining correlation and traceability to relevant upstream results.

The overall test methodology shall be architecture-aware and supported by demonstrated measurement capability. Test coverage shall evolve according to the failure mechanisms and interaction risks introduced at each successive level rather than unnecessarily repeating characteristics already demonstrated upstream. Results shall be correlated across wafer, die, module, and system levels to identify process shifts, post-singulation changes, assembly-induced degradation, and system-integration effects while maintaining complete traceability and genealogy from the originating wafer through final system configuration. Production test coverage, sampling strategies, validated surrogate measurements, Statistical Process Control (SPC), and qualification activities shall be established according to measurement capability, process stability, defect-detection risk, failure criticality, and downstream impact, providing a continuous technical basis for qualification and production release throughout the CPO manufacturing flow.




Figure 14: Co-Packaged Optics (CPO) manufacturing and test flow from wafer-level characterization through die, assembly, optical-engine, and module-level verification to final system-level validation, with traceability, correlation, and process feedback maintained across all stages.



10.1 Wafer Level Testing

Wafer acceptance criteria establish the minimum electrical, optical, RF, mechanical, thermal, and manufacturing requirements that shall be satisfied before a silicon photonics (SiPh) wafer is released for die singulation and heterogeneous integration into a co-packaged optics (CPO) assembly. The objective is to ensure that only wafers capable of producing a sufficient population of Known Good Die (KGD) proceed to downstream assembly, maximizing cumulative packaging yield and minimizing high-cost packaging risk.

Wafer acceptance shall be based on a combination of parametric, functional, structural, thermal, metrology, and statistical process-control metrics. Silicon photonic wafers require electro-optic qualification of passive waveguides, active modulators, integrated or heterogeneously attached lasers, photodetectors, optical coupling interfaces, and other applicable photonic structures in addition to conventional semiconductor parametric testing.

Wafer-level test methodology shall account for the intended CPO integration architecture:

- Single-sided configuration: Applicable optical and electrical test/integration interfaces are accessible from one primary device surface.
- Double-sided configuration: Functional, optical, electrical, or assembly access is required on both device surfaces.

For double-sided configurations, the wafer-level test setup shall provide appropriate access, alignment, probing, handling, and measurement capability for both surfaces without introducing mechanical damage, contamination, or thermal/mechanical stress that could affect device performance or downstream packaging.

All mandatory tests defined in applicable wafer-level test sections shall be completed before wafer release. Wafers failing a mandatory acceptance criterion shall be rejected, quarantined for engineering review, or classified as non-production development material. Requirements associated with lasers, photodetectors, modulators, heaters, or other architecture-specific functions shall apply only where those functions are present on the wafer.



10.1.1 Electrical, Optical, RF, and Thermal Acceptance Criteria

Wafer acceptance shall evaluate the following critical-to-quality (CTQ) categories across all applicable wafer-level test and integration interfaces:

CTQ Category	Minimum Acceptance Requirement
Electrical	Parametric and functional yield shall satisfy qualified product/process limits.
Optical	Optical device and optical-path yield shall satisfy qualified product/process limits.
RF / High-Speed	High-speed functional performance shall satisfy applicable bandwidth, signal-integrity, or functional requirements.
Thermal	Integrated heater and thermal-control functions shall satisfy applicable product limits.
Laser	Integrated or heterogeneously attached lasers shall satisfy applicable functional requirements.
Photodetector	Integrated photodetectors shall satisfy applicable functional requirements.
Waveguide	Continuity and optical-loss requirements shall be satisfied.
Critical Defects	No wafer-level critical defect that compromises downstream processing, reliability, or CPO integration shall be permitted.

Table 13: Electrical, Optical, RF, and Thermal Acceptance Criteria

Direct high-speed functional testing, such as Bit Error Rate (BER) or eye-quality measurement, should be used where practical. A validated surrogate methodology, including electro-optic S-parameters or other correlated measurements, may be used where correlation to the required high-speed performance has been demonstrated and documented.

Laser functional acceptance should include applicable parameters such as threshold current, optical output power, operating current/voltage, center wavelength, slope efficiency, and Side-Mode Suppression Ratio (SMSR). Photodetector acceptance should include applicable parameters such as responsivity and dark current.

Waveguide continuity and optical-loss performance should be evaluated separately, as a continuous optical path may still exhibit unacceptable propagation or insertion loss.

10.1.2 Mechanical and Physical Acceptance Criteria

Mechanical and physical wafer acceptance shall evaluate, as applicable to the device architecture and process flow, total thickness variation (TTV), wafer bow/warpage, surface roughness, particle/defect density, alignment-fiducial accuracy, facet or trench geometry, and optical coupling-surface quality.

Acceptance limits shall be defined by qualified process and downstream assembly requirements rather than by a universal numerical limit where the required value is product- or process-dependent.

Metrology methods shall provide sufficient resolution, repeatability, and measurement uncertainty relative to applicable acceptance limits. Suitable methods include optical profilometry, high-precision optical metrology, Atomic Force Microscopy (AFM), dark-field inspection, Optical Coherence Tomography (OCT), confocal metrology, or an equivalent validated method.

Where a measurement technique evaluates localized wafer regions, sampling quantity and locations shall be defined by the process-control plan and shall represent areas critical to optical coupling, bonding, singulation, or heterogeneous integration.

For double-sided CPO configurations:

- Features on both surfaces used for optical coupling, electrical access, thermal interfaces, bonding, or alignment shall be included in the applicable inspection and metrology plan.
- Front-to-back alignment accuracy shall be verified using a qualified metrology method where features on opposite surfaces require registration for downstream integration.
- Wafer handling, chucking, probing, and optical alignment shall not damage or contaminate qualified structures on either surface.

Critical defects shall be classified according to their impact. A localized defect may result in exclusion of the affected die without requiring rejection of the entire wafer unless the defect is systematic, repeating, located within a designated wafer-level critical area, or indicates loss of process control.

10.1.3 Wafer Yield Classification and Disposition

Composite Functional Yield shall represent the percentage of testable dies or chiplets on a wafer that satisfy all applicable mandatory Known Good Die (KGD) functional acceptance criteria within the defined wafer-level test boundary. Applicable criteria shall be established according to the device architecture and may include electrical, optical, electro-optical, RF/high-speed, and thermal characteristics. Physical, mechanical, and metrology requirements necessary for downstream assembly shall be evaluated through the applicable wafer and die acceptance criteria.

Composite Functional Yield shall be determined directly from the pass/fail status of individual testable dies or chiplets and shall not be calculated by mathematically multiplying separate domain-specific yields. Composite Functional Yield and relevant domain-, parameter-, or failure-mode-specific yield indicators shall be monitored statistically by wafer, lot, product, process revision, and other applicable manufacturing factors to establish qualified process baselines and identify abnormal variation. Release limits, warning limits, and engineering-review criteria shall be defined by the applicable product or manufacturing control plan based on sufficient process, measurement, and production data.

Wafer disposition shall consider Composite Functional Yield together with individual Critical-to-Quality (CTQ) results, spatial yield distribution, defect signatures, process-control data, and measurement validity. No minimum Composite Functional Yield value shall, by itself, establish wafer acceptability or override failure of a mandatory safety-, reliability-, performance-, or process-critical requirement.

Systematic spatial signatures, repeating reticle-related patterns, abnormal centre-to-edge behaviour, localized failure clusters, or other statistically significant departures from the qualified process baseline shall be identified through wafer-map and statistical analysis. Wafer maps and associated test data shall retain sufficient spatial and test-channel traceability to support identification of process-, equipment-, reticle-, test-site-, channel-, and location-dependent failure signatures.

Abnormal yield, CTQ performance, or spatial behaviour shall result in appropriate material control and disposition through the applicable engineering or quality review process. Disposition may include production release, additional testing, hold for investigation or failure analysis, controlled engineering use, or rejection based on the identified failure mechanism, affected CTQs, process history, measurement confidence, and downstream CPO integration risk.

10.1.4 Known Good Wafer (KGW) Definition

A Known Good Die (KGD) is an individual die that has satisfied all applicable mandatory die-level electrical, optical, RF, thermal, and functional criteria required for downstream assembly, together with applicable physical and metrology release criteria. For double-sided configurations, KGD status shall include all applicable mandatory criteria associated with both device surfaces.

A Known Good Wafer (KGW) is a wafer that has successfully completed all applicable wafer-level screening, metrology, and release requirements and contains a spatially identified population of KGD suitable for downstream singulation and CPO integration.

KGW qualification shall include 100% die-level testing of designated KGD CTQ parameters. Parameters demonstrated through validation to have sufficient process capability, measurement-system capability, and correlation to other qualified production controls may be monitored through statistically justified sampling or SPC where appropriate.

Process capability indices such as C_p and C_{pk} should be applied to designated production CTQs only after the associated process has demonstrated statistical stability. C_{pk} alone shall not be used as justification for removing a critical parameter from 100% screening.

Complete spatial traceability shall be maintained between wafer ID, die coordinates, surface/interface orientation where applicable, test results, KGD status, and final wafer disposition.

10.1.5 Wafer Release Documentation Requirements

Wafer release documentation shall establish conformance, traceability, and correlation with downstream die-, package-, and CPO-level test results. The release record shall include:

- Electrical and optical wafer maps with die-coordinate traceability;
- RF/high-speed test results or validated surrogate-test results;
- Thermal characterization results;
- Wafer metrology and defect/inspection results;
- SPC and applicable process-capability results;
- Equipment calibration and measurement-system status;
- Wafer, substrate, material, and process-lot traceability;
- Test station and applicable probe/fixture identification;
- Test-program and measurement-recipe revision; and
- Mask, process-flow, and applicable manufacturing revision history.

For double-sided configurations, wafer maps and test records shall identify surface-specific measurement data where applicable while maintaining a unified wafer and die-coordinate reference throughout wafer acceptance, singulation, and downstream packaging/test insertions.

Electronic wafer-map data shall use a recognized substrate-mapping format, such as SEMI E142 or an agreed equivalent.

First-pass test results shall be retained separately from retest and final-disposition results. Retesting shall not overwrite the original measurement record.

10.1.6 Statistical Process Control and Change Management

For qualified production, selected electrical, optical, RF, thermal, and physical CTQs shall be monitored using Statistical Process Control (SPC). SPC shall monitor wafer-to-wafer and lot-to-lot variation, within-wafer spatial distributions, first-pass yield, retest rate, KGD population, process capability, and recurring failure signatures.

For double-sided configurations, SPC shall include surface- or interface-specific monitoring where fabrication, handling, optical coupling, or testing can produce asymmetric behaviour between device surfaces.

SPC warning and action limits shall be established independently from absolute product specification limits where sufficient process data are available. An out-of-control condition shall require containment or engineering review even when individual devices remain within specification limits.

Significant changes to wafer processing, mask revision, materials, optical coupling structures, heterogeneous integration methods, test equipment, probe/fixture configurations, test software, or measurement recipes shall be evaluated through formal change control. The required correlation or requalification shall be determined according to potential impacts on wafer acceptance, KGD classification, measurement correlation, and downstream CPO performance.



Figure 15: Wafer-Level Testing Systems



10.2 Die-level testing

Die-level testing shall be performed following wafer singulation to verify that individual silicon photonic (SiPh) and applicable electronic/optoelectronic dies satisfy the requirements necessary for subsequent heterogeneous integration into a Co-Packaged Optics (CPO) assembly.

The primary objective is to establish Known Good Die (KGD) status before high-value assembly operations and to evaluate performance shifts or defects introduced or revealed by wafer singulation, die handling, backside processing, facet formation, or other post-wafer operations.

Die-level testing shall complement, rather than duplicate, wafer-level testing. Wafer-level results shall be correlated with die-level measurements to identify post-singulation changes. Parameters demonstrated to remain stable and strongly correlated between wafer- and die-level measurements may use reduced test coverage or statistically justified sampling at die level, provided KGD confidence and defect-detection capability are not compromised.

Die-level test methodology shall account for the intended CPO integration architecture:

- **Single-sided configuration:** Applicable electrical, optical, thermal, and assembly interfaces are accessed primarily from one die surface.
- **Double-sided configuration:** Functional, electrical, optical, thermal, or assembly interfaces require access to both die surfaces.

For double-sided configurations, the die-handling and test setup shall provide controlled access to the required interfaces on both surfaces without damaging, contaminating, obstructing, or mechanically loading functional structures. Sequential or simultaneous surface access may be used according to the device architecture and qualified test methodology.

The die-level qualification program shall provide sufficient coverage to establish KGD status according to device functionality, accessible interfaces, downstream integration requirements, and demonstrated wafer-to-die correlation.



Test Category	Primary Die-Level Objective	Typical Characteristics
Visual / Mechanical	Detect defects introduced or revealed by singulation and handling that may affect downstream integration.	Cracks, chipping, contamination, edge/corner integrity, die dimensions, flatness/warpage, bonding surfaces, alignment features
Electrical	Verify DC functionality and electrical-interface integrity.	Continuity, leakage current, supply/bias current, resistance, interface integrity
Optical / Electro-Optic	Verify passive and active photonic functionality and applicable optical interfaces.	Optical loss, output power, wavelength response, modulator behavior, photodetector response, laser characteristics, coupling-interface performance
High-Speed RF	Verify applicable high-speed electrical and electro-optic performance.	S-parameters, impedance, insertion/return loss, crosstalk, bandwidth, eye quality or BER where applicable
Thermal	Verify temperature-dependent die functionality under representative die-level boundary conditions.	Parameter drift, self-heating, thermal tuning, wavelength behavior, temperature-dependent electrical/optical performance
Reliability Screening	Screen applicable latent or early-life failure mechanisms before high-value integration.	Pre-/post-stress functional and parametric verification; burn-in where required
KGD / Traceability	Establish production release status and preserve correlation from wafer through downstream assembly.	KGD disposition, wafer-to-die correlation, die identification, wafer coordinates, surface/interface identification, test genealogy

Table 14: Die Level Testing Characteristics

Not every characteristic listed above is required for every die architecture. The production test plan shall define applicable CTQs and screening coverage based on device functionality, integration architecture, qualification results, process capability, measurement capability, and failure risk.



10.2.1 Visual and Mechanical Inspection

Visual inspection shall be performed after singulation to identify defects affecting subsequent assembly, optical coupling, electrical connection, thermal integration, reliability, or device performance. Inspection shall evaluate, as applicable, die edges and corners, cracks, chipping, delamination, contamination, metallization and bond-pad condition, optical coupling interfaces, facets, alignment features, and other assembly-critical structures.

Inspection and metrology methods may include Automated Optical Inspection (AOI), machine vision, optical microscopy, profilometry, interferometry, or equivalent qualified techniques, as appropriate to the characteristic being evaluated.

Mechanical and dimensional inspection shall verify applicable die dimensions, thickness, flatness/warpage, edge integrity, bonding surfaces, and alignment features.

For double-sided configurations, critical functional and assembly features on both surfaces shall be included in the applicable inspection plan. Front-to-back alignment accuracy shall be verified using a qualified metrology method where registration between opposite surfaces is required for downstream integration.

Critical defects introduced or revealed by singulation shall result in die rejection or formal engineering disposition.

10.2.2 Electrical Characterization

Electrical characterization shall verify applicable DC and electrical-interface functionality before release for CPO integration.

Testing shall evaluate, as applicable, continuity, leakage current, supply/bias current, resistance, and electrical interface integrity. High-speed electrical measurements may be included where required to establish KGD status or where die-level probing provides a representative measurement of the intended interface.

Electrical acceptance limits shall be defined by qualified product and process specifications.

10.2.3 Optical and Electro-Optic Characterization

Optical characterization shall verify the applicable passive and active photonic functions required for downstream CPO integration. Measurements may include, as applicable, optical insertion loss, propagation loss, optical output power, wavelength-dependent response, polarization-dependent behavior, modulator performance, photodetector responsivity and dark current, laser characteristics, and optical coupling-interface performance.

Electro-optic characterization shall be performed where combined electrical excitation and optical response, or optical excitation and electrical response, are required to verify device functionality.

The optical test configuration shall provide controlled alignment and coupling to the die-level optical interface while minimizing measurement variation caused by positioning, polarization, coupling condition, or fixture loading. Optical interfaces that become accessible only after singulation, such as edge or facet coupling structures, shall be evaluated where required for KGD qualification.

10.2.4 High-Speed RF Characterization

High-speed RF characterization shall be performed where RF performance is required to establish KGD status. Characterization may include S-parameters, impedance, insertion loss, return loss, crosstalk, bandwidth, and other signal-integrity parameters appropriate to the device architecture.

Direct high-speed functional testing, including eye-quality or Bit Error Rate (BER) measurements, should be used where practical and technically representative. A validated surrogate methodology, such as electro-optic S-parameter characterization, may be used where correlation to the required high-speed functional performance has been demonstrated and documented.

10.2.5 Thermal Characterization

Die-level thermal characterization shall verify that applicable electrical, optical, and electro-optic functions remain within qualified limits over the required die-level temperature conditions. Characterization may evaluate temperature-dependent parameter variation, self-heating, thermal tuning behavior, wavelength drift, laser behavior, photodetector response, and other temperature-sensitive CTQs.

The die test fixture shall provide controlled and repeatable thermal contact appropriate to the device architecture without obstructing or damaging required interfaces.

Full product-level thermal performance shall not be inferred solely from bare-die measurements where the final package, heat spreader, thermal interface material, or CPO assembly materially changes the thermal boundary conditions.



10.2.6 Reliability Screening and Burn-In

Reliability screening or burn-in may be performed at die level where required by device technology, qualification strategy, product reliability requirements, or identified failure mechanisms.

Stress conditions shall be defined by a qualified reliability methodology and shall not introduce failure mechanisms that are not representative of the intended device application. Applicable CTQs shall be evaluated before and after stress to identify functional failure or unacceptable parameter drift.

Burn-in shall not be considered a universal requirement for all SiPh or CPO dies. Its application, sampling level, stress conditions, and duration shall be established by the applicable qualification and reliability plan.

10.2.7 Known Good Die (KGD) Qualification and Acceptance

A Known Good Die (KGD) is a singulated die that has satisfied all applicable mandatory die-level electrical, optical, RF, thermal, functional, inspection, and reliability-screening requirements necessary for downstream CPO integration.

For double-sided configurations, KGD status shall include all applicable mandatory requirements associated with both die surfaces. Correlation or requalification shall be performed according to the potential impact on KGD classification and downstream CPO integration.

Production test plans shall identify CTQs requiring 100% screening and those permitted to use statistically justified sampling or other qualified process controls based on validation results, process capability, measurement capability, wafer-to-die correlation, defect-detection risk, and failure criticality.

A die failing a mandatory KGD criterion shall not be released for production integration unless formally dispositioned through an approved Material Review Board (MRB) or equivalent engineering/quality process.

10.2.8 Release Documentation and Traceability

Die-level test records shall provide objective evidence of KGD qualification and maintain correlation to the originating wafer and downstream CPO assembly.

Complete traceability shall be maintained between the unique die identifier, originating wafer ID and die coordinates, applicable surface/interface orientation, process and manufacturing revision, test station and fixture/probe configuration, calibration status, test-program/recipe revision, test results, and final KGD disposition.



Electronic wafer-map data used to maintain wafer-to-die genealogy shall use a recognized substrate-mapping format, such as SEMI E142 or an agreed equivalent.

For double-sided configurations, surface-specific measurement data shall remain associated with the same unified die record.

First-pass, retest, and final-disposition results shall remain distinguishable within the Manufacturing Execution System (MES) or equivalent data-management system. Retesting shall not overwrite the original measurement record.

10.2.9 Statistical Process Control and Change Management

For qualified production, selected die-level CTQs shall be monitored using Statistical Process Control (SPC). SPC shall monitor, as applicable, die-to-die and lot-to-lot variation, wafer-to-die parameter shifts, first-pass yield, retest rate, KGD yield, recurring failure signatures, and relevant spatial relationships to the originating wafer.

For double-sided configurations, SPC shall include surface- or interface-specific monitoring where fabrication, singulation, handling, probing, optical coupling, or testing can produce asymmetric behaviour between die surfaces.




Figure 16: Die-Level Testing systems

Significant changes to singulation, die handling, backside processing, facet preparation, test equipment, chuck/fixture configuration, probe configuration, optical coupling methods, test software, or measurement recipes shall be evaluated through formal change control.

10.3 Module-level testing

10.3.1 Objectives and Scope

Module-level testing shall be performed after completion of the applicable heterogeneous integration and assembly processes to verify that the assembled Co-Packaged Optics (CPO) module or optical engine satisfies the requirements for downstream system integration and production release.

The primary objective is to verify integrated module functionality and identify performance shifts or defects introduced or revealed by die attach, bonding, electrical interconnection, optical attachment and alignment, thermal integration, or package assembly.

Module-level testing shall complement, rather than duplicate, wafer- and die-level testing. Results shall be correlated with available upstream data to identify integration-induced changes and support process control. Parameters demonstrated to remain stable and strongly correlated across manufacturing stages may use reduced test coverage or statistically justified sampling where product risk and defect-detection capability permit.

Test configurations shall represent the intended module interfaces and operating conditions sufficiently to evaluate the interaction of integrated photonic, electronic, optical, mechanical, and thermal elements. Defined electrical and optical measurement reference planes shall be established for applicable interface measurements to support consistent interpretation and correlation of results.



10.3.2 Module-Level Integration Verification

Module-level testing shall emphasize characteristics that depend on the interaction of constituent dies, interconnects, optical interfaces, package structures, and thermal-management elements.

Integration Area	Module-Level Risk	Verification Focus
Die-to-Die / Die-to-Substrate Integration	Electrical discontinuity, parasitics, mechanical stress, or performance shift introduced during assembly	Interconnect integrity, interface functionality, assembled signal paths, upstream-to-module parameter shifts
Optical Attachment and Coupling	Coupling degradation, misalignment, contamination, or instability introduced during optical attachment	Completed optical-path performance, alignment stability, channel uniformity, polarization integrity where applicable
Electrical-to-Optical Signal Path	Performance degradation resulting from interaction of EIC, PIC, interconnect, and optical interfaces	End-to-end Tx/Rx functionality, modulation quality, receiver performance, high-speed link behavior
High-Speed Interconnect	Added loss, reflections, crosstalk, jitter, or bandwidth limitation from assembled interconnects	Signal integrity at defined reference planes and high-speed functional performance
Power and Control Integration	Excessive power, bias instability, control failure, or abnormal interaction among integrated components	Module power states, bias/control functions, telemetry and monitoring functions
Thermal Integration	Local heating, temperature gradients, thermal crosstalk, or performance drift	Temperature-dependent performance, tuning behavior, thermal stability, interaction among heat-generating components
Mechanical / Package Integration	Bond, substrate, fiber, connector, or package defects affecting performance or long-term integrity	Assembly integrity, package condition, optical attachment, alignment, and mechanical interfaces
End-to-End Module Operation	Failure of the assembled optical engine despite compliant constituent dies	Module initialization and operation, Tx/Rx path functionality, link performance, fault and monitoring behavior

Table 15: Module Level Testing Characteristics

The production test plan shall define module-level Critical-to-Quality (CTQ) characteristics necessary to detect integration-induced defects and verify complete module functionality. Measurements established at upstream manufacturing stages should not be repeated unless required to detect assembly-induced change, verify an integrated performance characteristic, support correlation, or control downstream manufacturing risk.



10.3.3 Visual, Mechanical, and Assembly Inspection

Visual and mechanical inspection shall verify the integrity of the completed module and identify assembly defects that may affect electrical, optical, thermal, mechanical, or long-term functional performance. Inspection shall evaluate, as applicable, package and substrate condition, die and bond integrity, contamination, mechanical damage, electrical interconnects, optical interfaces, fiber or fiber-array attachment, connectors, alignment features, and other assembly-critical structures.

Inspection and metrology methods may include Automated Optical Inspection (AOI), machine vision, microscopy, X-ray imaging, acoustic microscopy, interferometry, coordinate metrology, or equivalent qualified techniques according to the characteristic being evaluated.

Defects exceeding qualified acceptance criteria shall result in rejection, rework through an approved process, or formal engineering disposition.

10.3.4 Electrical and Power Characterization

Electrical and power characterization shall verify power delivery, applicable electrical interfaces, control and monitoring functions, and assembled electrical signal paths.

Measurements may include supply voltage and current, module or subsystem power consumption, bias conditions, interface continuity, and energy-efficiency metrics where required by the product architecture or specification.

Measurements shall be performed under defined reference conditions to support correlation across modules, test stations, and manufacturing stages.

10.3.5 Optical and Fiber-Interface Characterization

Optical characterization shall verify module performance at defined optical reference planes following applicable optical attachment and alignment processes.

Measurements may include optical output power, insertion or coupling loss, wavelength-dependent behaviour, spectral characteristics, channel uniformity, polarization-dependent characteristics, and other applicable optical-interface CTQs.

Where fibers, fiber arrays, lenses, connectors, or other coupling structures form part of the module, testing shall verify the optical performance and integrity of the completed interface.

The measurement configuration shall control or characterize relevant sources of measurement variation and uncertainty, including alignment, polarization, optical-interface repeatability, reference-plane definition, and instrumentation calibration.



10.3.6 Transmitter and Receiver Characterization

Transmitter and receiver characterization shall verify applicable communication functions of the assembled module at defined electrical and optical reference planes.

Transmitter characterization may include optical output power, modulation amplitude, wavelength characteristics, eye quality, timing characteristics, and transmitter-quality metrics required by the applicable interface specification.

Receiver characterization may include sensitivity, overload behaviour, dynamic range, electrical response, noise-related performance, and other receiver characteristics required by the applicable interface specification.

For architectures using an External Light Source (ELS), module testing shall verify applicable characteristics at the module optical-source interface rather than requiring characterization of laser generation not contained within the module. Relevant characteristics may include input optical power, wavelength, polarization, noise-related performance, and interface integrity.

10.3.7 High-Speed and End-to-End Functional Testing

High-speed testing shall verify communication performance of the integrated module at the applicable operating data rates and interface conditions.

Characterization may include S-parameters, impedance, insertion loss, return loss, crosstalk, bandwidth, jitter, and eye-quality parameters as applicable.

End-to-end functional testing shall evaluate the combined electrical, photonic, optical, and control functions using representative operating conditions and test patterns. Bit Error Rate (BER) and, where applicable, Forward Error Correction (FEC)-related performance shall be evaluated according to the applicable interface specification.

Direct functional measurements should be used where practical and technically representative. Validated surrogate measurements may be used for selected production CTQs where correlation to the required functional performance has been demonstrated and documented.



10.3.8 Thermal Characterization

Thermal characterization shall verify that the assembled module maintains required electrical, optical, and functional performance under applicable operating temperature and power conditions.

Testing may evaluate temperature-dependent parameter variation, self-heating, temperature gradients, thermal interaction between integrated components, thermal tuning behavior, wavelength drift, and other temperature-sensitive CTQs.

The test configuration shall provide controlled and repeatable thermal boundary conditions appropriate to the qualification objective. Where module performance depends on external heat spreaders, Thermal Interface Materials (TIMs), system-level liquid cooling, or airflow, the applicable thermal boundary conditions shall be defined and documented.

Module-level thermal results shall not be extrapolated to system-level performance where the final system materially changes the thermal boundary conditions.

10.3.9 Reliability Qualification and Screening

Module-level reliability activities shall address applicable failure mechanisms associated with heterogeneous integration, optical attachment, electrical interconnects, package materials, and thermal/mechanical interfaces.

Qualification stresses such as temperature cycling, damp heat, mechanical shock, vibration, or operating-life stress shall be applied according to the applicable reliability plan and identified failure mechanisms. Production burn-in or other reliability screening shall be applied only where required by the qualified manufacturing and reliability strategy.

Applicable CTQs shall be evaluated before and after reliability stress to identify failures or unacceptable parameter shifts. Qualification and ongoing reliability results may be used to establish or modify production screening requirements where supported by documented failure-mechanism analysis, process capability, and product risk.

10.3.10 Module Qualification and Acceptance

A module shall be eligible for production release when all applicable mandatory module-level requirements have been successfully completed and the required manufacturing and test records are complete.

Production test plans shall identify CTQs requiring 100% screening and those permitted to use statistically justified sampling, validated surrogate measurements, or other qualified process controls. Test coverage shall be established based on validation results, process capability, measurement capability, defect-detection risk, failure criticality, and downstream system impact.

Acceptance limits shall be defined by the applicable product, interface, manufacturing, and system requirements.

A module failing a mandatory acceptance criterion shall not be released for production use unless successfully reworked through a qualified process or formally dispositioned through an approved Material Review Board (MRB) or equivalent engineering/quality process.

10.3.11 Traceability, Data Management, SPC, and Change Control

Complete manufacturing genealogy shall be maintained between the module identifier and the applicable originating wafers, constituent dies, assembly lot, process revisions, and upstream qualification records.

Module-level records shall include, as applicable, test-station and fixture configuration, measurement reference planes, equipment and calibration status, test-program or recipe revision, relevant test conditions, measured CTQs, and final module disposition.

First-pass, retest, rework, and final-disposition results shall remain distinguishable within the Manufacturing Execution System (MES) or equivalent data-management system. Original measurement records shall not be overwritten by subsequent testing.

For qualified production, selected module-level CTQs shall be monitored using Statistical Process Control (SPC). Monitoring should include applicable module-to-module and lot-to-lot variation, first-pass yield, retest and rework rates, upstream-to-module parameter shifts, channel uniformity, recurring failure signatures, and assembly-related process trends.

Significant changes to die attach, bonding, electrical interconnection, optical attachment or alignment, package or thermal materials, test hardware, measurement reference planes, test software, or test recipes shall be evaluated through formal change control. Correlation or requalification shall be performed according to the potential impact on module performance, reliability, and downstream system integration.




Figure 17: Module Level Testing Characteristics

10.4 System Level Testing

System-level testing shall verify that the fully integrated Co-Packaged Optics (CPO) platform satisfies applicable functional, electrical, optical, network, power, thermal, management, and reliability requirements necessary for system qualification and production release. The system test boundary shall be defined by the product architecture and may include the host ASIC or processor, optical engines, system board, Power Delivery Network (PDN), thermal-management system, firmware, management interfaces, and associated networking functions.

Testing shall evaluate functions and interactions observable only after complete platform integration, including system initialization, host-to-optical-engine communication, end-to-end data transmission, multi-port operation, power and thermal interactions, management functions, and recovery from defined fault conditions.

Parameters qualified at wafer, die, or module level should not be repeated unless required to evaluate a system-level interaction, operating margin, correlation, or integration risk.

System Interaction	Primary Integration Risk	System Verification Focus
Host–Optical Engine Integration	Initialization, configuration, synchronization, or interface failure	Engine discovery, initialization, clock/reset behavior, host-to-engine communication
Optical Link and Data Path	Degradation across the complete electrical-to-optical communication path	Link establishment, stability, recovery, data integrity, BER/FEC-related performance
Multi-Port / Multi-Engine Operation	Shared-resource contention, crosstalk, common-mode effects, or performance degradation	Concurrent operation, aggregate traffic, interaction effects, workload scaling
Power–Workload Interaction	Voltage instability, transient response, or sequencing failure under dynamic load	PDN behavior, power states, load transitions, workload-dependent stability
Thermal–Workload Interaction	Hot spots, thermal coupling, throttling, or performance drift	Component temperatures, thermal gradients, cooling effectiveness, performance stability
Firmware and Management	Configuration, monitoring, telemetry, or control failure	Management/control interfaces, telemetry, alarms, configuration and status reporting
Fault, Recovery, and Protection	Failure to detect, report, or recover from defined abnormal conditions	Fault response, reset/restart, link reacquisition, diagnostics, applicable protection functions

Table 16: System-Level Integration Verification

System-level Critical-to-Quality (CTQ) characteristics shall be selected to demonstrate complete platform functionality and detect integration or interaction failures not adequately observable at earlier manufacturing stages.

10.4.1 System Integration and Initialization

System integration testing shall verify the sequence required to bring the fully assembled platform to its intended operational state. The system shall successfully execute applicable power-up, reset, clocking, firmware loading, configuration, optical-engine discovery, and management sequences.

Applicable control and management interfaces shall be verified for communication, configuration, status reporting, telemetry, and fault indication. Initialization and configuration failures shall provide sufficient diagnostic information to support failure isolation and manufacturing disposition.

10.4.2 Electrical and Power-Delivery Network Verification

System-level electrical verification shall evaluate electrical interfaces and power-delivery characteristics whose performance depends on complete platform integration. High-speed electrical paths shall be evaluated where required to verify host-to-optical-engine operation or identify system-level signal-integrity limitations. Measurements may include applicable timing, jitter, signal-integrity, and interface characteristics at defined system reference planes.

The complete Power Delivery Network (PDN) shall be verified under relevant operating states and workload transitions. Testing may evaluate supply regulation, ripple/noise, current transients, load response, sequencing, PDN stability, and other power-integrity characteristics required by the system architecture. Test conditions shall be sufficient to identify significant dynamic interactions among the host device, optical engines, control electronics, and other system loads.

10.4.3 Optical-Link and End-to-End Data-Path Verification

Optical-link testing shall verify that applicable system links can establish, maintain, and recover communication under representative operating conditions. Testing shall evaluate complete link behaviour rather than repeat intrinsic optical characteristics established at module level. Applicable verification may include link initialization, lane or port operation, link stability, recovery behaviour, and operating margin.

End-to-end data-path testing shall verify the combined operation of electrical, photonic, optical, firmware, and interface functions using representative data patterns or traffic. BER, FEC-related performance, error statistics, or other data-integrity metrics shall be evaluated as required by the applicable interface and system specifications.



10.4.4 Network and Mission-Mode Validation

Network and mission-mode testing shall verify system performance under traffic patterns and workload conditions representative of the intended deployment application. Testing may evaluate aggregate throughput, data integrity, latency, multi-port operation, traffic distribution, congestion behaviour, Quality of Service (QoS), scalability, and other functions required by the system specification.

Concurrent operation of multiple ports or optical engines shall be evaluated where required to identify shared-resource contention, crosstalk, power or thermal interaction, or performance limitations not observable during isolated link operation. Sustained or dynamic workload testing may be used to verify operational stability and expose intermittent or workload-dependent failures. Test duration, traffic profile, and workload severity shall be defined by the applicable validation, qualification, or production test plan.

10.4.5 Thermal Management and Power-Thermal Interaction

System thermal validation shall verify stable operation of the integrated CPO platform under representative power, traffic, environmental, and cooling conditions. Testing may evaluate component temperatures or qualified temperature indicators, thermal gradients, localized hot spots, workload-dependent heating, thermal coupling between system elements, cooling effectiveness, and temperature-dependent electrical, optical, or functional performance.

Where thermal throttling, active cooling control, alarms, or protective shutdown functions are implemented, their required operation and recovery behaviour shall be verified using qualified test conditions. Thermal boundary conditions and workload states shall be defined and documented sufficiently to support measurement repeatability and correlation with the intended deployment environment.

10.4.6 Robustness, Fault Recovery, and Protection

System robustness testing shall verify the ability of the integrated platform to detect, report, and respond appropriately to defined operational faults and abnormal conditions. Testing may include loss and restoration of optical links or optical-source power, reset or restart events, power-state transitions, communication interruptions, recoverable interface faults, and other scenarios identified through system risk analysis.

Where safety, protection, redundancy, telemetry, alarm, or diagnostic functions are implemented or required, their specified response and recovery behaviour shall be verified. Recovery requirements shall be defined by the applicable system specification. Fault testing shall remain within qualified product operating limits unless performed as part of a specifically defined robustness or qualification activity.



10.4.7 EMC/EMI and Environmental Qualification

Electromagnetic Compatibility (EMC), Electromagnetic Interference (EMI), and environmental qualification shall be performed where required by the applicable product specification, deployment environment, regulatory requirements, or qualification plan. EMC/EMI evaluation may include applicable conducted and radiated emissions, immunity, Electrostatic Discharge (ESD), and other requirements defined by the relevant regulatory or product qualification framework.

Representative system functionality shall be monitored during applicable immunity testing to identify communication errors, link interruptions, resets, or other unacceptable operational effects. Environmental and long-duration qualification shall address applicable system-level failure mechanisms and deployment conditions, including thermal, humidity, mechanical, or operational stresses where required. Qualification stresses shall not be interpreted as universal 100% production screening requirements.

Production environmental or stress screening may be applied where justified by the qualified manufacturing strategy, demonstrated failure mechanisms, and product risk.

10.4.8 Qualification, Acceptance, and Quality Control

System qualification shall establish that the defined CPO platform design and manufacturing baseline satisfy applicable functional, interface, reliability, environmental, and regulatory requirements. Production system acceptance shall require successful completion of applicable mandatory production CTQs and availability of the required manufacturing and test records. Qualification evidence applicable to the released design and manufacturing baseline need not be regenerated for each production unit.

Production test plans shall identify CTQs requiring 100% verification and those permitted to use statistically justified sampling, validated diagnostics, surrogate measurements or telemetry, or other qualified process controls. Test coverage shall consider validation results, process and measurement capability, failure criticality, defect-detection risk, and downstream deployment impact.

Acceptance criteria shall be defined by the applicable product, interface, system, and regulatory requirements. A system failing a mandatory production acceptance criterion shall not be released unless successfully reworked through a qualified process or formally dispositioned through an approved Material Review Board (MRB) or equivalent engineering/quality process.

10.4.9 Traceability, Data Integrity, SPC, and Change Control

Controlled and traceable system genealogy shall be maintained between the system identifier and applicable constituent optical engines or modules, host devices, board assemblies, hardware revisions, firmware/software versions, and system build configurations.

System-level records shall include, as applicable, test-station and fixture configuration, optical and electrical connectivity, test-program revision, traffic or workload configuration, equipment and calibration status, thermal and power conditions, measured CTQs, and final system disposition.

First-pass, retest, rework, and final-disposition records shall remain distinguishable within the Manufacturing Execution System (MES), Quality Management System (QMS), or equivalent controlled data system. Original measurement and diagnostic records shall not be overwritten by subsequent testing.

For qualified production, selected system-level CTQs shall be monitored using Statistical Process Control (SPC) or equivalent statistical methods where appropriate. Monitoring should include applicable first-pass yield, retest and rework rates, recurring link or initialization failures, parameter shifts, and manufacturing or test-station trends.

Significant changes to optical engines or modules, host silicon, board design, optical connectivity, power delivery, cooling architecture, firmware/software, system configuration, test hardware, or test methodology shall be evaluated through formal change control. Correlation, regression testing, or requalification shall be performed according to the potential impact on system functionality, performance, reliability, manufacturability, and deployment requirements.



Figure 18: Example of System-Level Testing Equipment

11. System level architecture future requirements

Chapters 3 through 8 of this specification specify a scaling path for CPO built on an implicit architectural premise: Optical Engines remain discrete, separable components mounted around the switch ASIC / FPGA, connected to the faceplate through fiber attach units and mid-board jumpers. That premise holds through the Phase 3 target defined as 1,024 optical interfaces per package on a 256-fiber shoreline. It does not hold indefinitely.

This chapter addresses the architectural transition the working group must plan for once the current discrete-OE, fiber-routed model reaches its practical limits, and defines what a next-generation reference architecture needs to specify.

11.1 Why the discrete-OE model has a horizon


Three independent scaling pressures identified elsewhere in this document converge on the same conclusion:

- **Interface density outpaces shoreline area.** The Phase 3 target achieves a 4:1 ratio of optical interfaces to physical fibers through WDM, but this only decouples interface count from fiber count, and it does not eliminate the physical fiber-attach problem at the OE. Every fiber in a 256-fiber shoreline still terminates at a discrete FAU connector with its own alignment tolerance, pull-strength requirement, and assembly yield exposure.
- **Assembly yield degrades with lane pitch.** Section 3.4's own risk register flags that Phase 1's reduced fiber-array pitch "may initially reduce fiber-attach yield," and each subsequent phase compounds this by adding engines and channels within the same package footprint.
- **Faceplate real estate is a hard physical ceiling.** Figure 9 (Faceplate Port Granularity) shows that even the highest-density connector families top out in the low thousands of fibers per 1RU panel height, and that density trades directly against the airflow surface a >1 kW air-cooled switch needs (Section 7, faceplate discussion).

Each of these is manageable individually within the current architecture through T0+36 months.

Collectively, they describe a model whose per-generation scaling factor is bounded by mechanical fiber handling (a constraint that does not track the electrical/compute scaling curve the rest of the switch or accelerator package is on).

Chapter 9 exists because the working group needs a plan for the architecture after mechanical fiber handling stops scaling, not merely a plan for making it scale one more generation. Future requirements shall not assume that optical-interface count, physical fiber count, aggregate bandwidth, and package-edge density scale at the same rate. These quantities shall be specified separately because wavelength count, lane rate, duplexing method, redundancy, and laser-distribution architecture can change their relationship



Requirement boundaries:

- | | |
|--|---|
| A Die-to-die electrical interface and protocol | D Detachable optical connection and external channel |
| B Optical-engine functional and thermal interface | E Management, telemetry, fault-containment, and safety interface |
| C Continuous-wave laser-delivery interface | |

Figure 19: Future CPO system architecture and requirement boundaries



11.2 Candidate future-requirement areas

11.2.1 Integration of optics into the chip/chiplet

The most direct answer to the fiber-handling ceiling is to remove the discrete OE as a separable, fiber-terminated component and integrate the optical I/O directly into the compute chiplet. This is analogous to how UCIe moved die-to-die electrical I/O from board-level SerDes into the package. This has direct precedent in the disaggregated compute trend the industry has already normalized electrically; extending it optically means the electrical/optical boundary Chapter 5 currently places at the OE shoreline would instead sit at the chiplet edge.

The requirements this creates are non-trivial and should be scoped explicitly in the next specification-design cycle:

Requirement area	What changes vs. current (Ch. 3–10) architecture	Open question for the working group
Serviceability	Chapter 3 OE is explicitly designed to be a replaceable unit, separable “electrically, optically, or both.” Chip-integrated optics removes this separation.	Does field-replaceability move up a level (whole chiplet/package swap) or is it accepted as a yield/reliability tradeoff at manufacturing time?
Test and qualification	OE-level optical test (Section 3.8’s per-lane BER/eye-margin regression) currently happens on a known-good discrete part before package assembly.	Does known-good-die optical test become mandatory pre-integration, and at what cost to test time relative to Section 3.8’s Phase 2/3 concerns?
Standards boundary	OIF CPO framework and OCI-MSA currently define footprints at the module/OE boundary (Section 4.1). A chiplet-integrated optical I/O already has an emerging multi-source basis: the ratified OCI-MSA specification (adopted in shipping CPO products as of 2026) and UCIe’s own roadmap toward optical die-to-die options in its next major release. A standard-package optical UCIe chiplet has already been publicly demonstrated (OFC 2025).	Advanced-package-level, multi-vendor interoperability (comparable to what UCIe’s electrical standard-package ecosystem already achieves) has not yet converged industry-wide. Should the working group build on OCI-MSA and UCIe’s optical roadmap as the foundation, or does chiplet-level optical integration still need its own dedicated MSA at the advanced-package tier?
Yield economics	Section 6.2 already documents that a non-replaceable, soldered OE concentrates fiber-damage risk into the highest-value component.	Chip-integrated optics inherit this risk at chiplet scale unless a mid-board or on-package decoupling point (Section 6.2’s jumper-cord model) is preserved somewhere in the stack.

Table 17: Proposed future system-level requirements set



11.2.2 Transition towards optical PCB

As discussed in previous chapters, optical connections between a package containing an ASIC and PIC and either the chassis backplane or a second package are currently implemented using FAUs. Due to the stringent alignment requirements between the PIC waveguides and detachable optical connectors, these connectors are inherently bulky. As optical I/O requirements scale to hundreds or thousands of connections per package, conventional FAU architecture will demand larger package footprints and substantial edge real estate. In addition, routing large numbers of multi-fiber FAUs within the limited volume of a chassis introduces significant system-level design challenges.

To support future optical connectivity requirements, new interconnect architecture is needed. The most natural approach is to extend the PCB paradigm used for electrical interconnects by integrating optical waveguides on or within the main PCB. The leading approaches are based on glass or polymer waveguides, each offering distinct advantages and tradeoffs. Similar to the evolution of electrical PCBs, two key challenges must be addressed for optical PCBs (OPCBs):

- PICs within the CPO package must be aligned to OPCB waveguides with sub-micron precision to maintain low optical loss over the server lifetime.
- A scalable, cost-effective manufacturing process is required to fabricate low-loss, three-dimensionally routed waveguides on or within PCBs.

Addressing these challenges will require significant advancements in both optical packaging and PCB manufacturing technologies.

11.2.3 Internal laser sources reconsidered

Section 8's ILS/ELS discussion frames the choice primarily around thermal load and field-replaceability, and the current Phase 3 roadmap (Section 3.4) resolves it in favor of ELS specifically to remove laser heat from the engine at higher engine counts. A chip-integrated optical architecture reopens this question rather than settling it: if the OE no longer exists as a separable unit, ELS's principal advantage (hot-swappable light supply without touching the compute package) must be re-derived at the new integration boundary. The future requirement is not "should ILS return" but "what does laser serviceability mean when there is no OE boundary to swap at." This needs to be scoped alongside heterogeneous III-V-on-silicon integration yield data, since chip-level integration raises the cost of a laser failure from "replace one OE" to "replace one compute chiplet."

Also, the supplier shall specify laser junction-temperature limits, thermal resistance to the cooling solution, wavelength drift, aging behavior, failure containment, and the capacity impact of a laser failure. Internal-laser qualification shall be performed at the package-level thermal corner, not only under isolated photonic-die conditions.

11.2.4 Thermal management as a first-order architectural constraint




Figure 20: Future-state CPO package architecture - from discrete OE to chip-integrated optics

Every phase in Section 3.5's risk register treats thermal density because of engine count and wavelength-locking overhead, something managed after the electrical/optical architecture is fixed. Chip-integrated optics inverts this: laser heat, modulator temperature sensitivity, and ASIC / FPGA thermal load now share the same silicon, at a point in the design flow where thermal floor planning must happen concurrently with the optical and electrical architecture, not after it. This is consistent with where the industry's own thermal ceiling already sits (referring to Section 6 faceplate discussion notes that connector density trades directly against airflow surface at >1 kW per switch box), and that pressure only intensifies as more optical function moves onto the compute die itself. The next specification cycle should treat thermal co-design as a Chapter 3 level input requirement, sized against real wavelength-drift and crosstalk budgets, rather than a Chapter 9 afterthought.

11.3 The APC Working Group Roadmap

None of the above displaces the Phase 1–3 roadmap in Chapter 3.3, which remains the near-term, normative scaling path through T0+36 months. Chapter 7's role is to flag that the discrete-OE/FAU/faceplate architecture, this entire specification specifies, is a bounded solution, and that the Advanced Photonics Coalition should open a parallel workstream (which is analogous to the glass-interposer workstreams already running in Chapter 4) scoped specifically to chip-integrated optical I/O, laser-serviceability models under integration, and concurrent thermal-optical co-design methodology. The deliverable of that workstream shall be the requirements baseline for whatever follows Phase 3, not a Phase 4 addendum to the current roadmap. It shall demonstrate, under representative environmental and traffic conditions:

- Compliance with the end-to-end optical budget on every supported wavelength.
- Stable operation across the specified thermal envelope.
- Containment and reporting of engine, lane, wavelength, and laser-source failures.
- Recovery of unaffected capacity following a single replaceable-unit failure.
- Manufacturing test coverage sufficient to protect final package yield.
- Multi-vendor interoperability at each standardized boundary.
- Safe and repeatable field-service procedures.
- Telemetry sufficient to identify loss of operating margin before traffic interruption.
- A credible second-source path for system-critical optical and mechanical interfaces.
- Documented power, cooling, reliability, and cost scaling for the next two bandwidth generations.

These acceptance criteria shift the evaluation of CPO from isolated component performance to deployable system behavior. That transition is essential because future CPO success will depend not only on achieving higher optical density, but on delivering predictable availability, manufacturability, interoperability, and lifecycle cost at data-center scale.

11.4 Outlook

Scale-up optical links require careful planning across a complex combination of elements to deliver high performance at low cost per bandwidth. New technologies like lasers and connectors will be initially very costly, but both are more like DRAM than ASIC / FPGA, so costs should drop after a few years of rapid-ramp demand outstripping supply. Silicon photonics is more than 30 years old, but only recently have production volumes ramped, and they will ramp much more. There will be rapid learning and technical improvements/inventions to reduce losses at all stages, which will further boost performance and lower cost per bandwidth. Major changes to the CPO system level architecture, such as moving away from FAUs (because of number of optical connections required) and moving to waveguides inside the PCBs and the challenges associated with these changes could represent a major shift in architectural CPO designs in the coming years as optical and electrical power densities increase.

Future requirements would focus on:

- Integration of optics into the chips/chiplets
- Move to waveguides inside PCB to accommodate number of optical connections required and alleviate the crowded beachfront
- Internal lasers
- Thermal management and novel materials to handle the high-power densities optically and electrically

Future compute platforms may incorporate ASICs, FPGAs, chiplets, accelerators, and mesh-based compute architectures. The working group should therefore avoid defining future optical interfaces around a single compute implementation model. Future optical and electrical interconnect standards should support heterogeneous compute fabrics, including architectures that distribute processing, data movement, scheduling, and memory functions across interconnected compute elements. This approach preserves interoperability while allowing innovation in both conventional and emerging compute architectures.



12. Acronyms

ASIC	Application Specific Integrated Circuit.
BER	Bit Error Rate.
CDR	Clock and Data Recovery.
CLTE	Continuous-time Linear Equalizer.
COBO	Consortium for On-Board Optics.
CPO	Co-Packaged Optics.
DFE	Decision-feedback Equalizer.
DR4	Module output with 4 optical channels using 8 single-mode fibers, 500 m reach.
ELS	External Laser Source
FFE	Feed-forward Equalizer
FR4	Module output with 4 optical channels using 2 single-mode fibers, 2 km reach
IC	Integrated Circuit
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
MDI	Medium Dependent Interface
MSA	Multi-Supplier Agreement
NPO	Near Packaged Optics
OBO	On Board Optics
PIC	Photonic Integrated Circuit
PUE	Power Usage Effectiveness
QSFP-DD	Quad Small Form-factor Pluggable Double Density.
SiPh	Silicon Photonics
SN(R)	Signal-to-Noise (Ratio) – the ratio of signal power to the noise power, often expressed in decibel (dB) units
SR4	Module output with 4 optical channels using 8 multi-mode fibers, 100 m reach
SR8	Module output with 8 optical channels using 16 multi-mode fibers, 100 m reach
TP2	Optical Test Point 2, as described in IEEE 802.3 Ethernet standards, at which optical power is measured from the source PMD.
TP3	Optical Test Point 3, as described in IEEE 802.3 Ethernet standards, at which optical receive signal is measured going to the destination PMD.
TIA	Telecommunications Industry Association
VSR	Very Short Reach
XSR, XSR+	Extra Short Reach



Thank you



ADVANCED
PHOTONICS
COALITION

www.advanced-photonics.org