The Next Generation of Pluggable Optical Module Solutions from the OSFP MSA

Sam Kocsis (Amphenol)

Brian Park (Arista)

With contributions from the OSFP MSA membership

September 18, 2022

Executive Summary

The OSFP MSA is proud to introduce OSFP1600 and OSFP-XD to the industry. This whitepaper highlights the key aspects and features of each solution with the expectation that both solutions will have a place in future data center applications.

The OSFP-XD solution has attracted significant interest in the market when it was publicly announced in June 2021. The opportunity to develop a pluggable IO solution that can address thermal challenges and meet electrical performance expectations of next-generation optical modules has engaged a large number of OSFP MSA members in the development of this specification and we wanted to take this opportunity to thank them for their many contributions.

This whitepaper summarizes the OSFP-XD mechanical, electrical, and thermal design capabilities. We expect the specification to be released early Q4 ’22 and the first 1.6 Tb/s OSFP-XD systems in the market in 2023.

OSFP

The OSFP has been broadly accepted for 400G (with 8x50 Gb/s host interface) and for 800G (8x100 Gb/s host interface) pluggable optics. The OSFP MSA has now completed the development of the OSFP1600 specification, supporting 8x200 Gb/s host interfaces. OSFP1600 carefully pushes the limits required to achieve the electrical performance while maintaining mechanical reliability and full backward compatibility with OSFP800.

With its integrated heatsink, the OSFP offers in excess of 30 Watt thermal performance which covers the full range of 1600G datacenter optics. On the electrical interface side, the OSFP1600 specification will also define cabled-connector footprints in addition to the traditional surface mount (SMT) connector.

The images below shows a 1U switch with 32 OSFP1600 ports supporting 51.2T throughput and a 2U switch with 64 OSFP1600 ports supporting 102.4T throughput, both using belly-to-belly configuration.

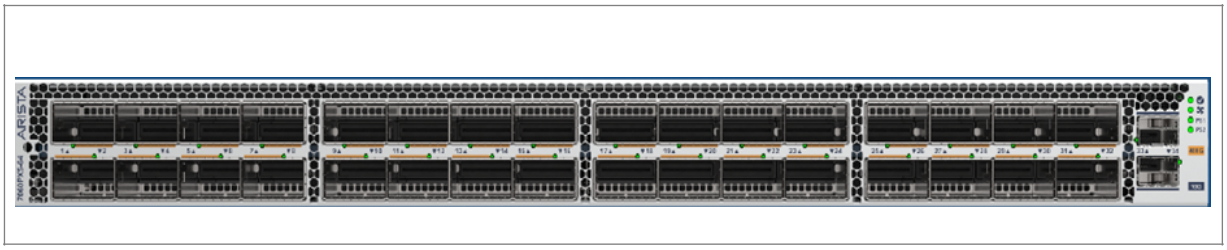

Example 1RU, 32-port Switch (OSFP)

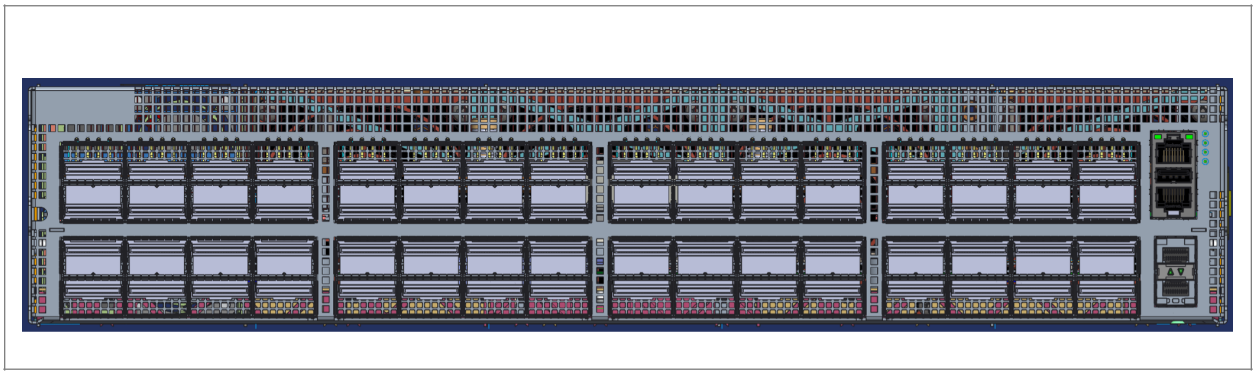

Example 2RU, 64-port Switch (OSFP)

OSFP-XD

While the OSFP1600 supports future switch silicon with 200 Gb/s electrical lanes, there is broad interest in 1.6 Tb/s optics modules with the 100 Gb/s electrical lane ecosystem.

The OSFP-XD (“eXtra Dense”) form factor was developed to meet this requirement. By doubling the number of electrical lanes from 8 to 16, the OSFP-XD offers 1.6T density with 16 lanes of 100 Gb/s and 3.2 Tb/s density with 16 lanes of 200 Gb/s in the future.

Recognizing that the success of any pluggable form factor depends on its ability to provide a robust mechanical, electrical, and thermal solution, OSFP-XD was designed around the following objectives:

1. Support up to 40W power for future 1600-ZR and 3200G optics modules
2. Support for passive copper cable solutions (DAC) per 100GBASE-CR1
3. Support 32-ports in 1RU and 64-ports in 2U chassis.

With these capabilities, OSFP-XD can support a wide range of system configurations, as illustrated in the chart below. In summary, OSFP-XD doubles the front panel density compared to the 8-lane OSFP or QSFP-DD form factor.

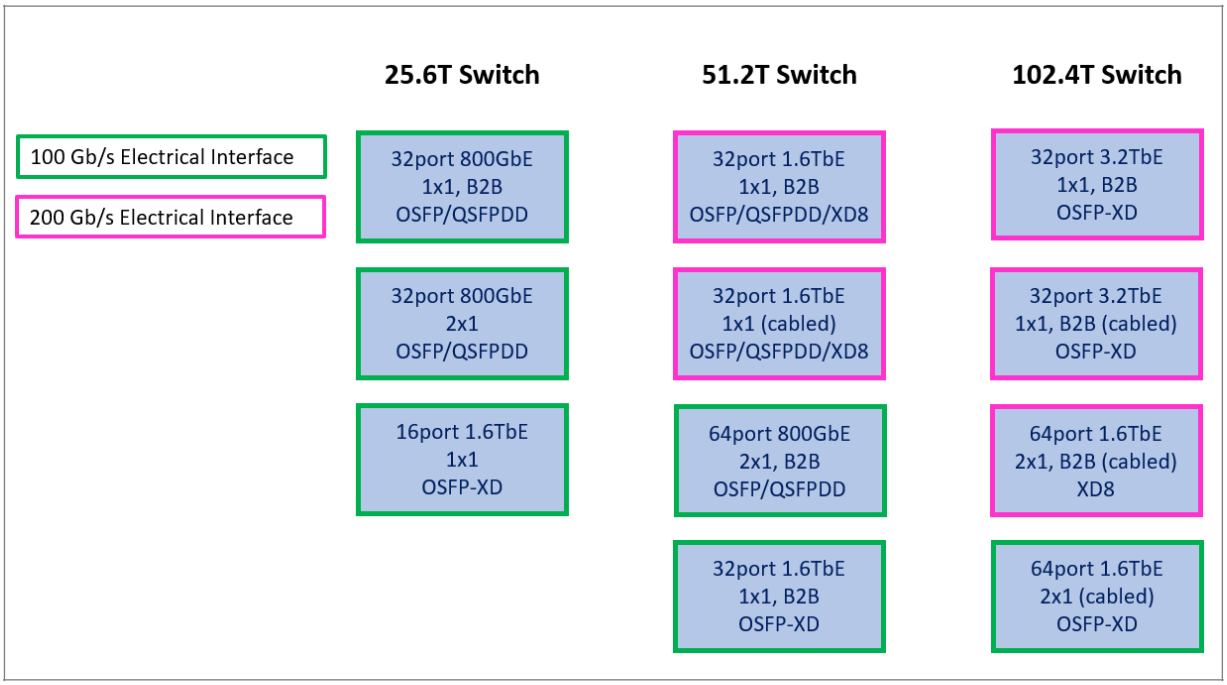

Anticipated system configurations for next-generation switches

Mechanical Details

The OSFP-XD solution doubles the number of high-speed electrical signals into the module by utilizing the well-known approach of adding a second row of contacts to the module's internal PCB or paddle card. This extends the length of the module slightly compared to the conventional OSFP module.

The latching mechanism for the OSFP-XD modules has been modified slightly from the OSFP module design to tighten tolerances to limit the maximum effective wipe, and thereby enabling a reduction of the parasitic effects on the electrical contacts. This work is considered critical to ensure readiness for 200 Gb/s electrical interfaces.

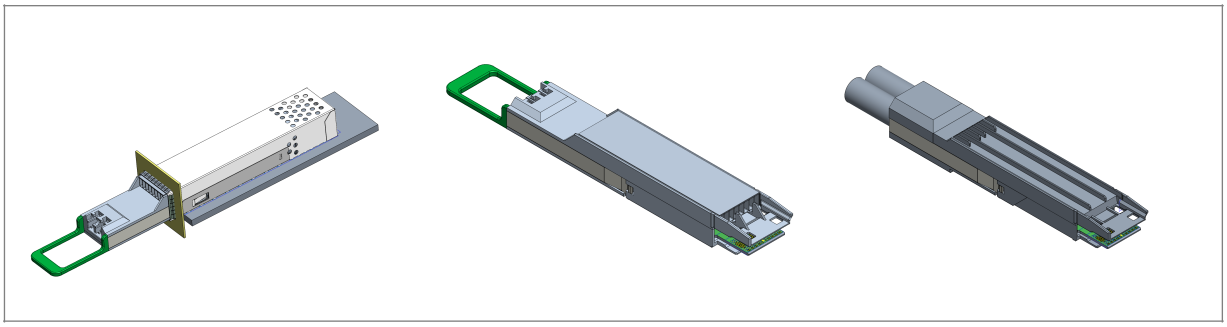

3D views of the OSFP-XD solutions

To accommodate both high-power optical and dense copper solutions, the specification will define separate but compatible heatsink specifications for both optical and copper modules, allowing module designs to optimize critical features for each implementation independently. Compared to the OSFP module, the OSFP-XD module is 20% taller by volume (15.5mm) and the internal volume is increased more than 12% for passive DAC applications. Designers will have the flexibility to trade-off internal volume for a taller integrated heat sink (IHS) depending on the intended application of the module.

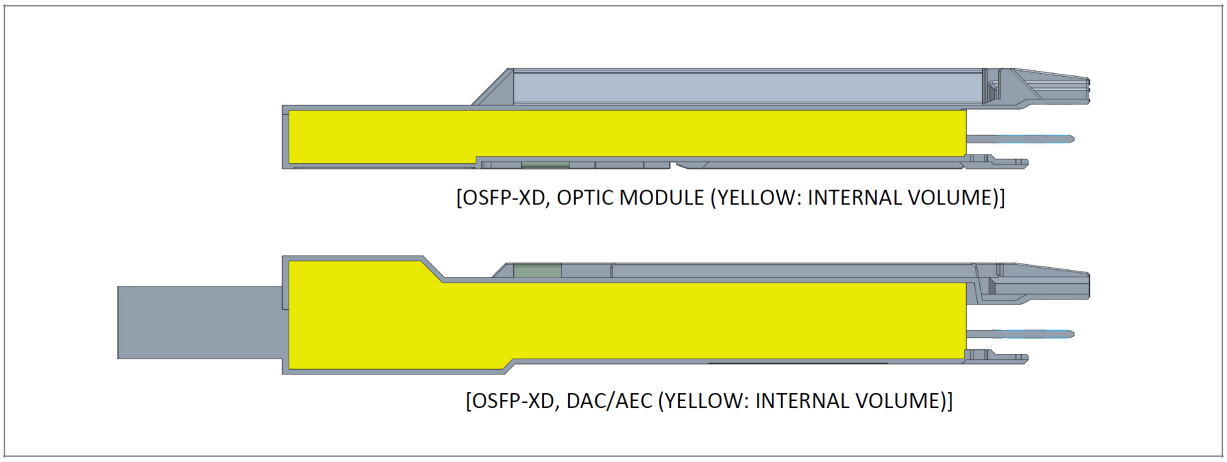

Illustration of trade-offs for Optical and Copper Module implementations

The OSFP-XD module will have a 1.2mm thick paddlecard in order to support 200 Gb/s electrical interfaces, to reduce routing complexities and for robust power distribution.

As a consequence of increasing the paddlecard thickness and the height of the module, OSFP-XD is not compatible with the existing OSFP form factor. To prevent any damage to an OSFP-XD port, keying features will be implemented in the OSFP-XD cages to keep OSFP modules from being inserted into an OSFP-XD port.

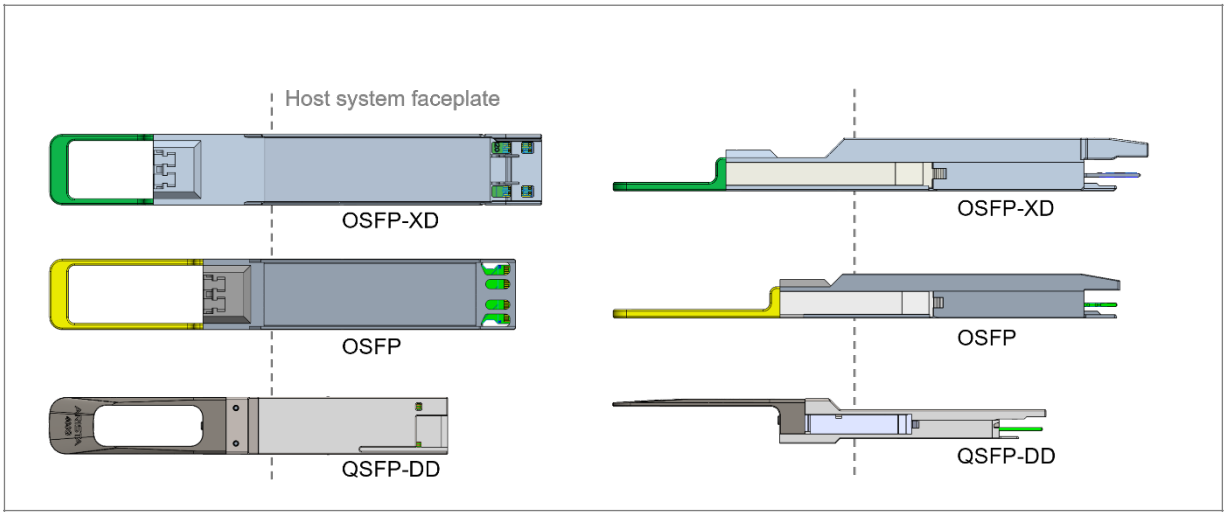

Relative size comparison of OSFP-XD, OSFP, and QSFP-DD modules

In addition to the integrated heatsink (IHS) modules, the OSFP-XD specification will also define a low-profile module that will use a riding heatsink (OSFP-XD RHS). The OSFP-XD RHS solution is not intended to support copper cable applications and is not expected to achieve the same thermal capabilities as the IHS solution unless alternative cooling techniques are employed.

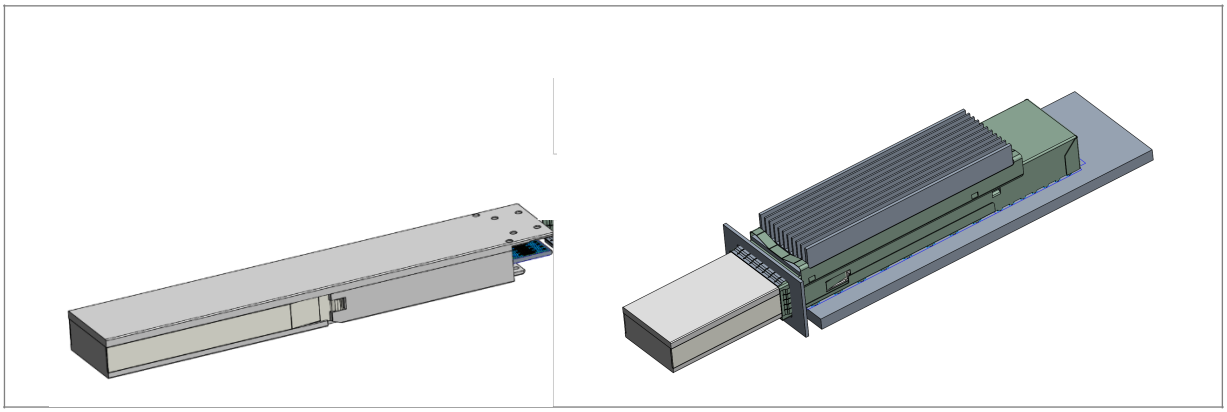

3D views of OSFP-XD RHS solution

On the host side, the OSFP-XD specification will offer options for both SMT and cabled-connector footprints. The SMT and cabled-connector footprints will be belly-to-belly compatible with each other and with the OSFP SMT and cabled-connector footprints.

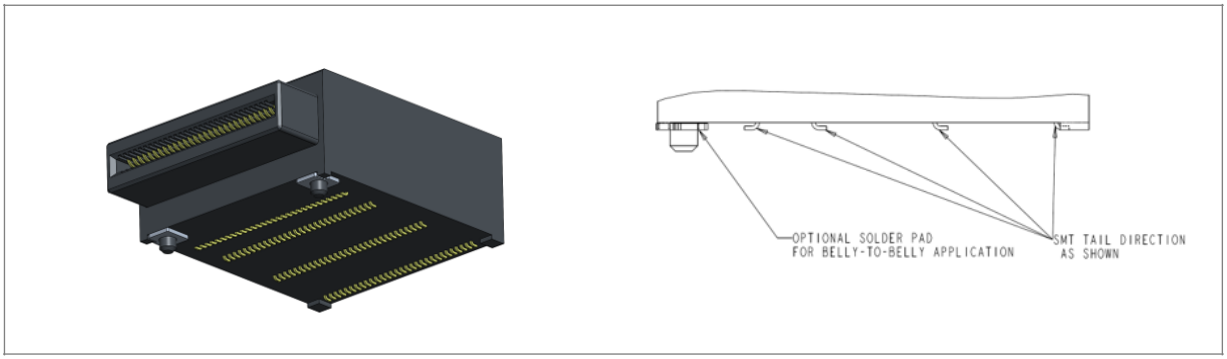

OSFP-XD Receptacle (SMT)

The OSFP-XD SMT footprint is identical to the OSFP 2x1 SMT footprint. The host footprint compatibility will simplify an upgrade path from 2x800G ports to a 1.6T port and allow for system implementers to leverage existing stack-up and routing techniques.

In addition to the traditional surface mount (SMT) connector, the OSFP-XD specification will define cabled-connector footprints. The cabled-connector uses twin-ax cable for the high-speed signals and press-fit attachment to the host PCB for the management and power signals. System integrators will have the flexibility to mix-and-match SMT and cabled-solutions in a variety of compatible belly-to-belly implementations to achieve their desired configuration. Footprint drawings for both the SMT and cabled-connector footprints are shown below.

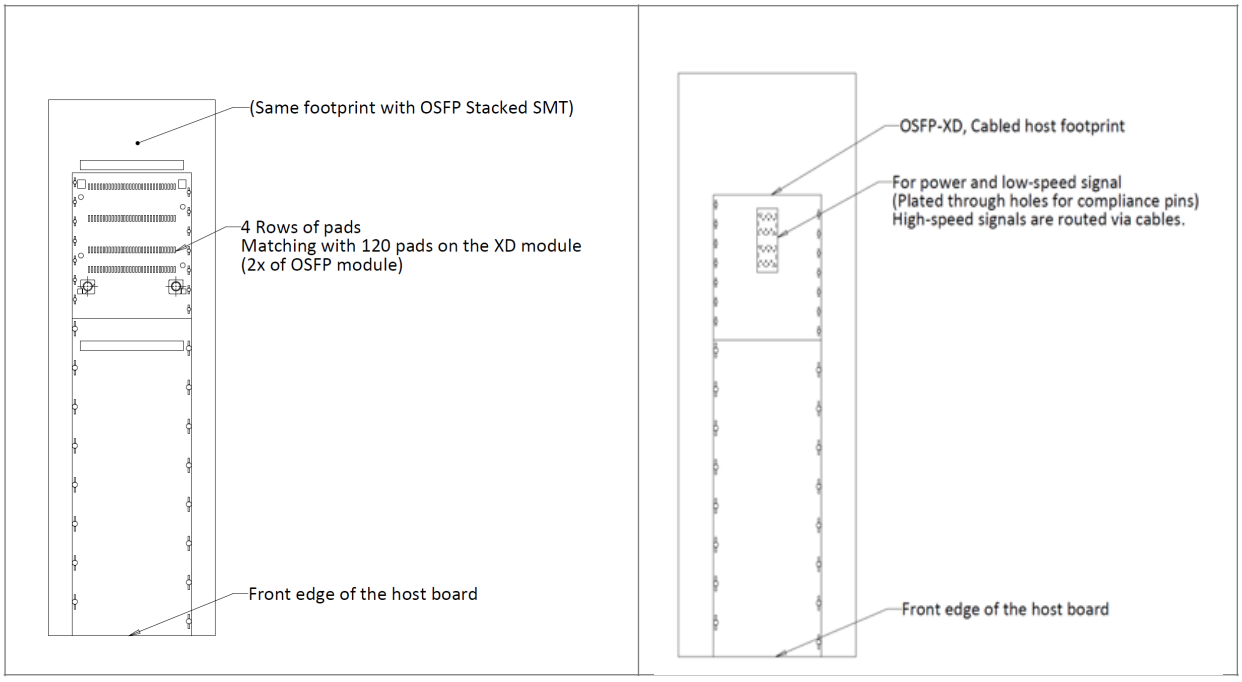

OSFP-XD Host Footprints (SMT and Cabled-Connector)

Thermal Management Capabilities

Power consumption for 1.6T and 3.2T modules is expected to reach up to 40W. It is well accepted that integrated heatsinks (IHS) provide better cooling performance compared to riding heatsinks (RHS) implementations, provided the fin height is optimized for the available area. The 1RU system provides a suitable boundary condition to support air-cooled switch applications as well as applications with alternative cooling techniques.

While 1.6T Datacenter optics are expected to consume less than 25W power, 1.6T-ZR coherent modules are expected to be in the 35-40W range and future 3.2T datacenter optics modules are also expected to be in the same range of power consumption.

A reference 1U system with 32 OSFP-XD ports was analyzed using airflow simulation tools for both 40W optical modules and 20W active electrical copper cables. As shown below, the OSFP-XD provides excellent cooling for optics modules up to 40W.

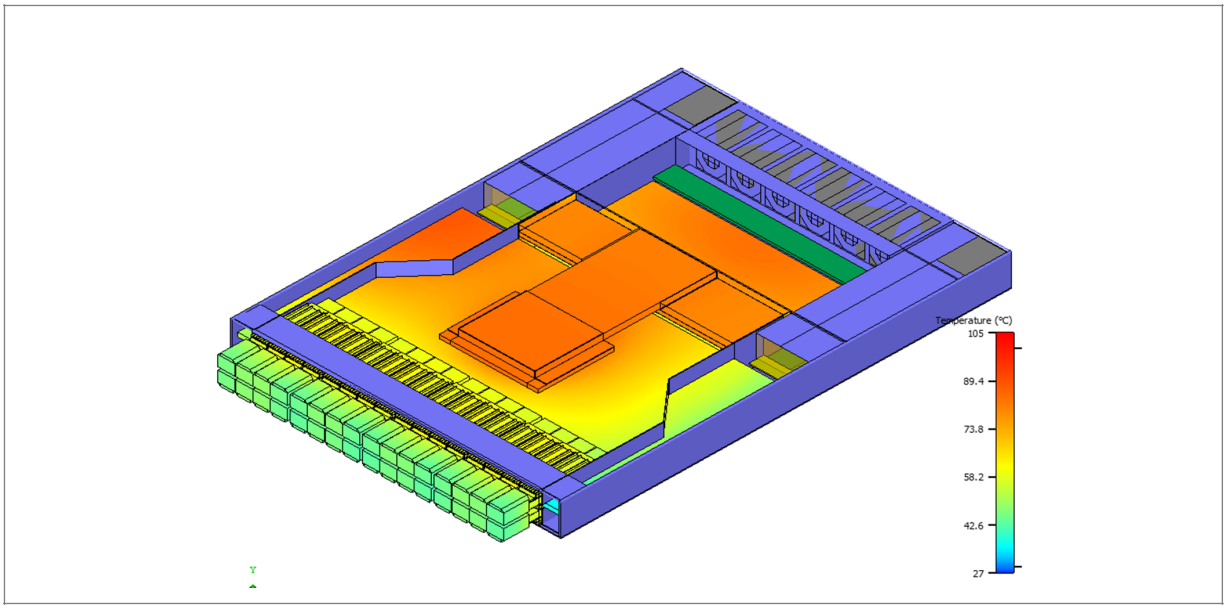

Snapshot from 1U System-level Thermal Simulation

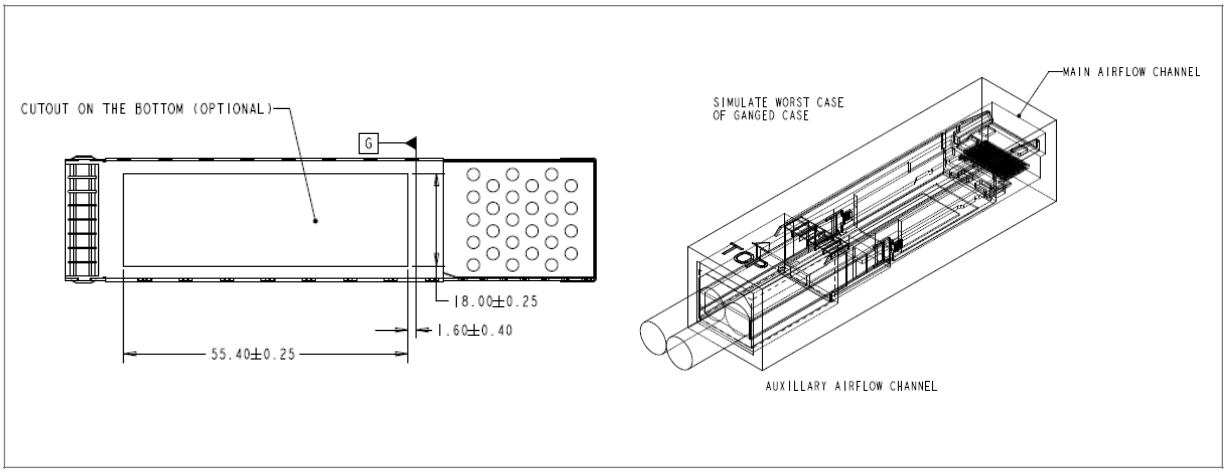

Illustration of example cage venting and normative Impedance Test Fixture

To create a reference for both suppliers and system implementors, the OSFP-XD specification will define a common test fixture. The resulting limits for the copper and optical modules are shown in the plot below.

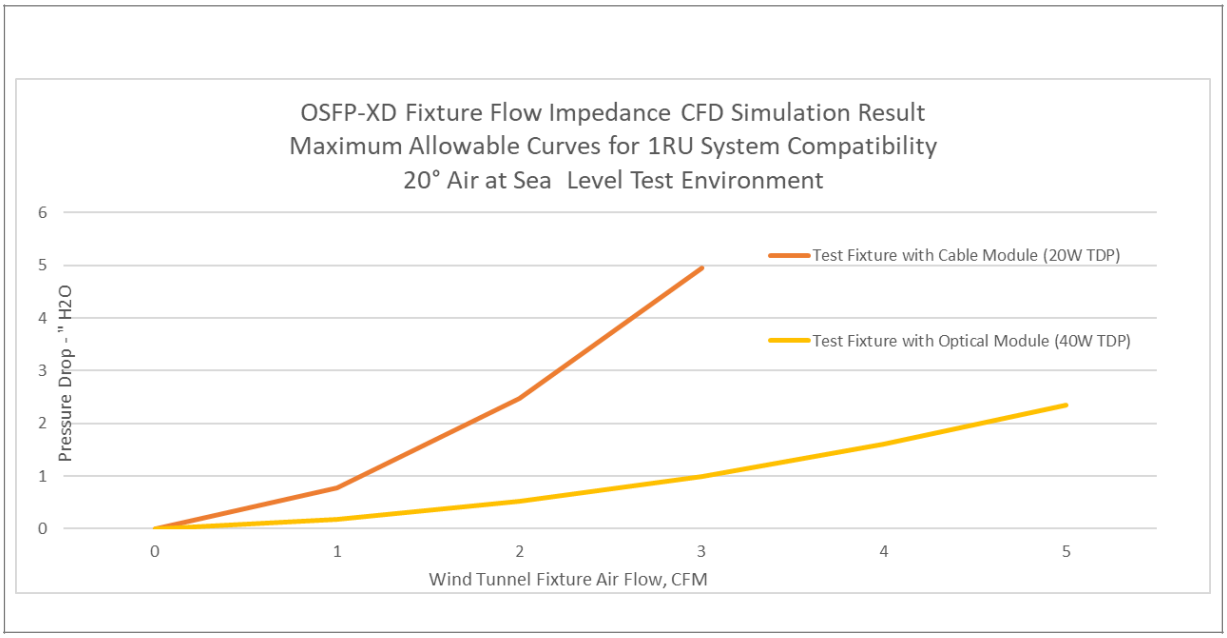

Target requirements for Copper and Optical Modules using MSA Test Fixture

While system designers have the final word regarding the capabilities of their solution, the key takeaway at this point should be that the thermal management capabilities of the OSFP-XD solution exceed those of any other pluggable module specification in the market today.

Electrical Interface

The electrical interface of an OSFP-XD module consists of a 120 contacts edge connector. It provides 32 contacts for 16 differential pairs of high-speed transmit signals, 32 contacts for 16 differential pairs of high-speed receive signals, 4 contacts for low-speed control signals, 8 contacts for power, 40 contacts for ground, and 4 contacts reserved for future use. There are several options for the future use pins under consideration of the MSA.

OSFP-XD Functional Pin Map (Module Perspective)

The host board connections to the OSFP-XD module will look identical to a stacked OSFP module. This provides a straightforward transition from host's using a 2x1 OSFP SMT solution to an OSFP-XD solution.

OSFP-XD modules will have the capability to support multiple Ethernet physical media configurations at both 112G-PAM4 and 224G-PAM4, ranging all the way up to 1.6TbE. The defined lane mapping will guide implementors and suppliers to help support the market transition from 8-lane to 16-lane solutions.

OSFP-XD can also support 8-lane optics modules that want to take advantage of thermal management capabilities and useable volume inside the module. An 8-lane OSFP-XD module (tentatively referred to as "XD-8") would utilize lanes 1 through 8 for

the high-speed signals and all the ground, power, and management signals available in the full OSFP-XD connector pinout.

The lane mapping for the 16-lane and 8-lane configurations is shown below. The 8-lane configuration will make use of the lanes 1 thru 8. The lane identifier corresponds to the TX/RX designation in the functional pin map.

(*L means Lane, L1 means Lane 1 in the port.)
(**P means Port, P1 means Port 1 in the PMD configuration)

PMD Configuration	Transmit and Receive Lane Assignments															
	L1*	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14	L15	L16
1x1.6T (112G-PAM4)	P1**															
2x1.6T (224G-PAM4) 2x800G (112G-PAM4)	P1								P2							
	L1	L2	L3	L4	L5	L6	L7	L8	L1	L2	L3	L4	L5	L6	L7	L8
4x800G (224G-PAM4) 4x400G (112G-PAM4)	P1				P2				P3				P4			
	L1	L2	L3	L4	L1	L2	L3	L4	L1	L2	L3	L4	L1	L2	L3	L4
8x400G (224G-PAM4) 8x200G (112G-PAM4)	P1		P2		P3		P4		P5		P6		P7		P8	
	L1	L2	L1	L2	L1	L2	L1	L2	L1	L2	L1	L2	L1	L2	L1	L2
16x200G (224G-PAM4) 16x100G (112G-PAM4)	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16

High-speed Lane Mapping (16-lane)

(*L means Lane, L1 means Lane 1 in the port.)
(**P means Port, P1 means Port 1 in the PMD configuration)

PMD Configuration	Transmit and Receive Lane Assignments															
	L1*	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14	L15	L16
1x1.6T (224G-PAM4) 1x800G (112G-PAM4)	P1								Not Used							
	L1	L2	L3	L4	L5	L6	L7	L8								
2x400G (224G-PAM4) 2x200G (112G-PAM4)	P1				P2				Not Used							
	L1	L2	L3	L4	L1	L2	L3	L4								
4x400G (224G-PAM4) 4x200G (112G-PAM4)	P1		P2		P3		P4		Not Used							
	L1	L2	L1	L2	L1	L2	L1	L2								
8x200G (224G-PAM4) 8x100G (112G-PAM4)	P1	P2	P3	P4	P5	P6	P7	P8	Not Used							

High-speed Lane Mapping Options (8-lane)

The designations for the reserved pins and management functionality could expand upon the functional capabilities of an OSFP module and would be developed in alignment with the Common Management Interface Specification (CMIS). The block diagram below provides functional view of the connections between a host and a module supporting OSFP-XD.

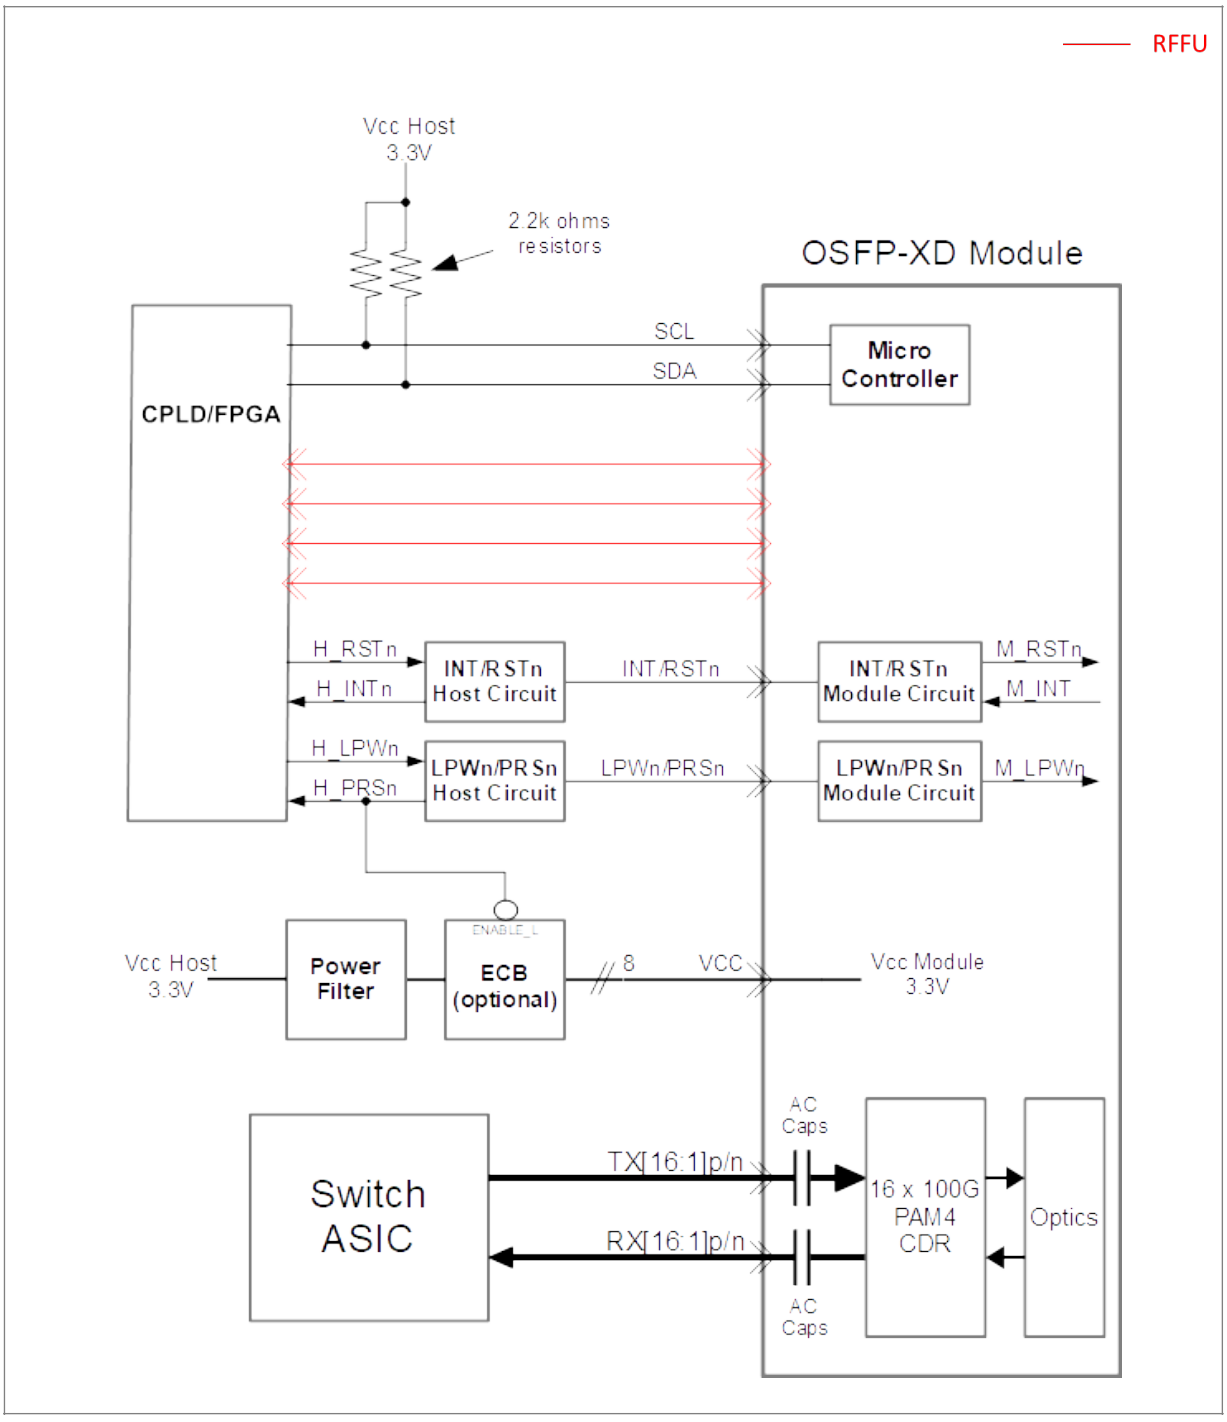

OSFP-XD Host and Module Block Diagram

Summary

The demand for optics modules is largely driven by the major cloud service providers. The next few years are expected to see rapid adoption in 800G modules as shown in the market forecast below for expected sales to the five largest cloud service providers. It is important to note that the majority of 800G optics modules initially will be used for break-out applications including dual 400G and octal 100G.

There are no market forecasts yet for 1600G modules, but one can expect a fast ramp once 1600G modules are being adopted by large cloud service providers. Similar to the 800G case, the majority of 1600G modules will be used for break-out applications.

The OSFP MSA roadmap provides an excellent mechanical and electrical solution for 800G, 1.6T, and 3.2T pluggable optics with best-in-class thermal performance and support for break-out applications, making these form factors a great choice to deliver the next-generations of pluggable optics modules.

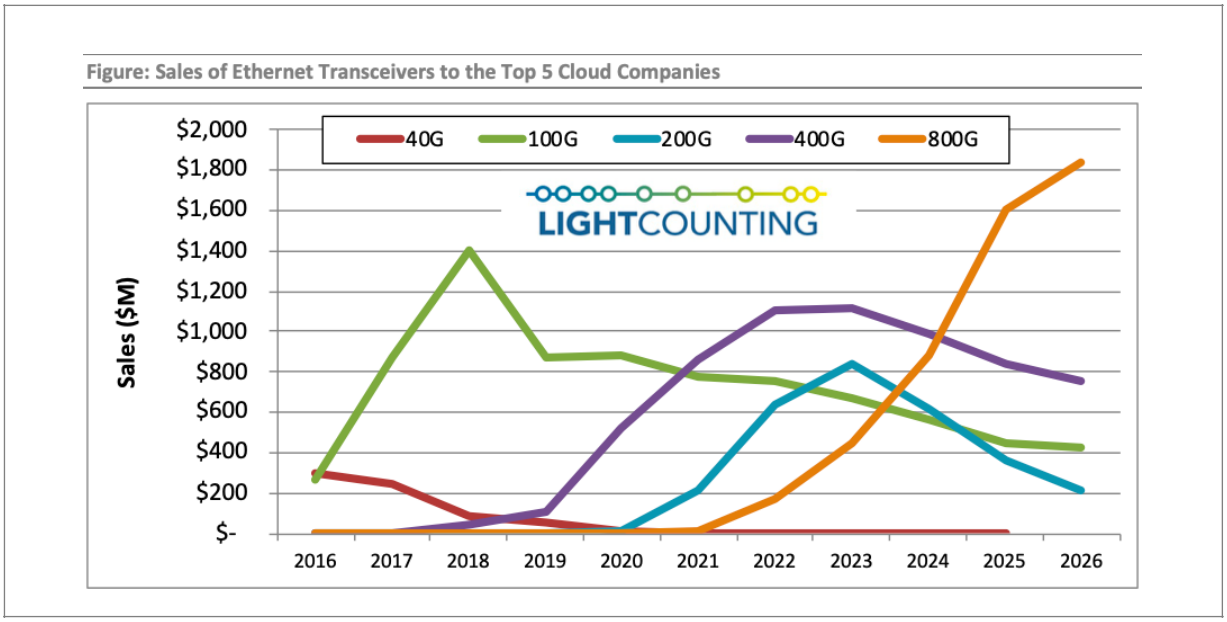

Source: Lightcounting, July 2021