

OpenZR+ Advances Coherent Interoperability

November 2019

Abstract

Coherent optical technology, traditionally a closed-system technology with proprietary implementations, is evolving to establish implementable standardization and interoperability. This migration can help network operators by ensuring multiple suppliers for their critical components, which reduces the risk of adopting new technology. This white paper looks at several of these industry standardization efforts, as well as recent industry efforts to leverage existing standards and introduces a new and widely supported 400G interoperability solution called OpenZR+.

Introduction

Standardization and interoperability of optical networking technology provide benefits to both network operators as well as suppliers. History has shown that adoption of optical transceiver/transponder solutions accelerates when implementable standardization and interoperability are established. Evidence of this can be seen in the widespread use of client optics based on optical interfaces defined in IEEE and pluggable form factors defined by Multi-Source Agreements (MSAs). Coherent optical technology, traditionally a closed-system technology with proprietary implementations, is evolving towards a similar path. This migration can help carriers and hyperscale data center network operators by ensuring multiple suppliers for their critical components, which reduces the risk of adopting new technology.

Initially adopted in long-haul applications, coherent optical interfaces have traditionally been deployed in bookended configurations, in which equipment from the same supplier are used on both sides of the link. In this model, vendors were differentiated from one another based on small performance differences that could result in cost savings by reducing the need for costly regeneration nodes.

In recent years, coherent technology has been utilized for shorter distance interfaces to support increasing bandwidth demands on the network edge. Some hyperscale data center networks require large numbers of regional interconnections, which has been driving requirements for coherent technology optimized for cost and power. Several hyperscale network operators have also expressed a need for coherent module solutions in the same form factors that support client optics in the emerging 400 Gigabit per second (Gbps) I/O generation.

Carriers and multiple-system operators (MSOs) have also been driving requirements for coherent optical transport interoperability in adjacent applications, such as access aggregation and metro core networks. These applications may require higher performance, additional data rates, OTN functionality and client multiplexing onto line-side transport channels for optimizing network efficiency.

This white paper looks at several of these industry standardization efforts, as well as recent industry efforts to leverage existing standards and introduces a new 400G interoperability solution called OpenZR+. Already publicly supported by numerous industry leaders, OpenZR+ provides an open, flexible and interoperable coherent solution in a small form factor pluggable module that addresses hyperscale data center applications for higher-performance edge and regional interconnects, as well as potential future carrier applications.

Historical Background

Coherent Interoperability Still Nascent

Historically, optical transport solutions were comprised of an end-to-end solution that provided differentiated technology to achieve high-capacity, long-reach links. Although the data being transmitted followed standardized protocols (e.g., SONET/SDH), the underlying optical layer consisted of proprietary solutions typically resulted in bookended transport links. This approach continued as vendors began introducing coherent transport solutions. Coherent interconnect technology required digital signal processors, utilizing

unique framing, proprietary signal processing, and Forward Error Correction (FEC) algorithms, increasing the justification for proprietary solutions.

When coherent technology moved into pluggable form-factor solutions (initially in the CFP followed by the CFP2 form factor), there were early carrier efforts, particularly by AT&T and Deutsche Telekom, to demonstrate interoperable 100G coherent transmission to promote multi-supplier interoperability. Through industry organizations such as OpenROADM MSA as well as the ITU, service providers drove standards which could be supported by multiple vendors, while leaving opportunity for higher performance differentiated proprietary implementations. In particular, a hard-decision FEC and differential encoding were defined in the standard. Hard-decision FEC offered less coding gain than widely used proprietary soft-decision FEC algorithms. Differential encoding simplified interoperability, but introduced a performance penalty compared to non-differential encoding.

Though multiple equipment vendors supported this 100G standard, network operator adoption to date has been limited, primarily due to the fact that higher performance offered by proprietary solutions outweighed the benefit of interoperability. Also, efforts to define a standard took several years to materialize despite 100G solutions being already available. During this time, a large performance gap between the standard and proprietary implementations limited the adoption of the standardized solution.

Lessons Learned: OIF 400ZR for DCI Edge 400G and the Power Consumption Target

Some network operators chose to take a different approach at 400G upon recognizing the lessons learned from the efforts at 100G. In late 2016, these network operators and a few vendors identified 400G as an intersection point for the industry to support coherent optics in the same form factors as emerging high-volume client optics, such as QSFP-DD and OSFP. At that time, it was anticipated that these form factors would support approximately 15W power dissipation. Hyperscale network operators proposed to the Optical Internetworking Forum (OIF) a new project targeting DCI edge applications up to 120km (See Figure 1), with an objective that anticipated module power dissipations not exceed 15W. The resulting OIF implementation agreement was known as 400ZR¹. The OIF hoped that by targeting a very specific application and starting their effort before vendors were well down the development path, the OIF standardization activity could be driven expeditiously.

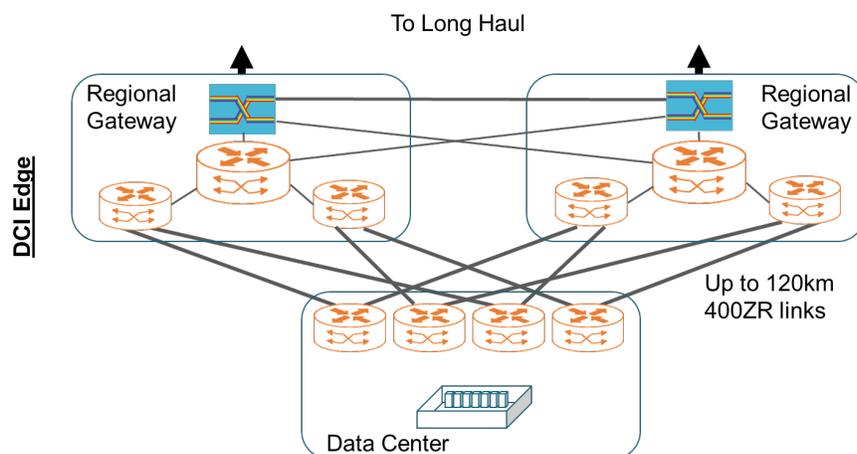


Figure 1. Target application space of DCI edge for 400ZR.

OIF was quite successful in expediting this effort, with power optimization taking the highest priority in each decision. 16QAM and ~60Gbaud were adopted for the modulation format and baud rate, respectively, after a trade-off comparison of these parameters was performed. In addition, a concatenated (hard-decision + soft-decision) FEC was selected after several FEC proposals were debated. Low-power support and sufficient coding gain for edge applications were driving factors towards this decision. In less than one year, the OIF defined most of the 400ZR interface technical details, which helped motivate increased industry investment in pluggable, interoperable coherent interfaces.

The Carrier/MSO Perspective: OpenROADM & CableLabs

In a separate effort, the OpenROADM MSA group embarked on a coherent standardization effort focused on carrier applications. This effort included carrier-centric features that were not within the scope of 400ZR (e.g., the 400ZR host-board electrical interface scope was limited to 400GbE). Carrier applications needed the flexibility to support multiple rates, multiplexing functionality and additional protocols such as OTN. Additionally, compared to 400ZR, these carrier applications had more stringent performance requirements and were not restricted to client optic form factors. This would permit form factors such as CFP2 to be developed.

Based on different priorities compared to OIF, the OpenROADM MSA chose a high performance soft-decision FEC, called openFEC (a.k.a, oFEC), that was contributed by Acacia Communications. OpenFEC offered higher coding gain when compared to 400ZR, which was comparable to proprietary bookended implementations typical in long-haul applications. OpenROADM MSA also defined a framing that supported 200G, 300G or 400G line rates all utilizing oFEC, with greater host-board interface flexibility.

In the same timeframe, CableLabs, a standardization body representing the cable/MSO industry, developed a coherent interconnect standard to address the cable access aggregation requirements, targeting data rates at 100G and 200G. Although this application had its own unique set of requirements, CableLabs recognized the benefit of aligning with other standards bodies. CableLabs decided to align with the 100G and 200G standards in OpenROADM because it offered the best combination of benefits for this application, including the high-performance oFEC for 200G. The CableLabs standard utilized a common protocol with OpenROADM, while keeping with optical specifications that addressed cable access aggregation requirements.

OpenZR+: Taking the Ambiguity Out of “ZR+”

While the OIF focused on specifications, leading switch equipment manufacturers diligently worked on increasing the QSFP-DD/OSFP power consumption ratings that their next generation equipment could accommodate. As a result, the per-port power consumption tolerance was increased beyond the 15W target set by OIF 400ZR. Thanks to this achievement, support for additional features and higher performance modes that required higher than 15W became possible. “ZR+” became a term loosely used by the industry to specify or capture potential broader modes of operation beyond 400ZR. Unfortunately, due to lack of a uniform definition, the use of the term ZR+ created confusion within the industry.

While the term ZR+ created confusion, the need behind the desire for higher performance beyond 400ZR remained strong. A key requirement was to accommodate hyperscale DCI links beyond 120km, while

maintaining the same QSFP-DD/OSFP form-factors. A survey of 400G standardization efforts pointed to elements of OpenROADM that could provide a standard-based high-performance addition to the 400ZR standard. Thus, the industry began looking at the next logical step, which would be to combine these vetted specifications and achieve the goal of addressing 400G Ethernet-centric solutions beyond 120km. This would enable the extension of hyperscale DCIs beyond the edge to regional distances, and expand the addressable market for module suppliers, providing greater economies of scale that benefits the entire distribution chain. This combination of the 400ZR standard with elements of OpenROADM became known as OpenZR+.

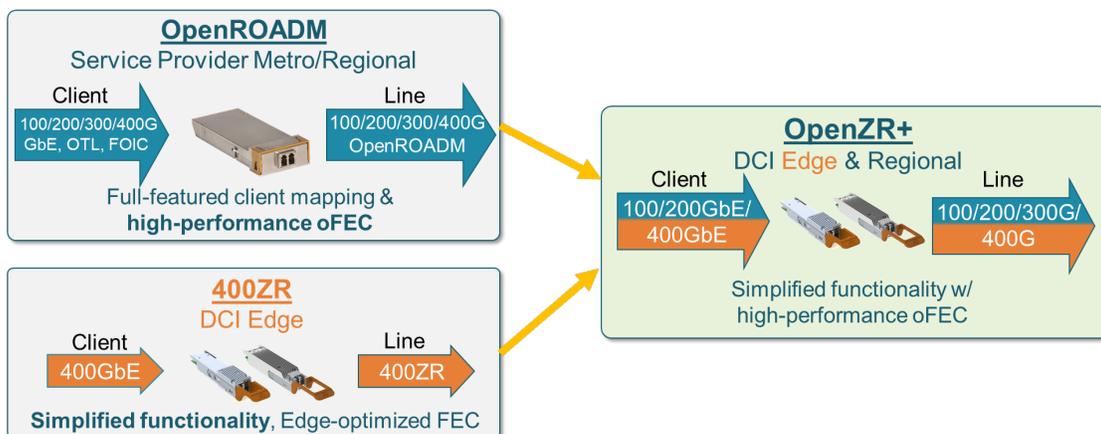


Figure 2. OpenZR+ is the logical combination of two industry standardization efforts that enables high performance DCI pluggable modules supporting multi-vendor interoperability.

As illustrated in Figure 2, OpenZR+ is a combination of two industry standardization efforts created to maintain the simple Ethernet-only host interface of 400ZR while adding support for features such as: (1) higher coding gain using oFEC from the OpenROADM standard, which extends the reach capability; (2) multi-rate Ethernet, which enables the multiplexing of 100GbE and 200GbE clients over the line-side link, providing optimization options for the switch/router equipment to channelize the traffic over the transport link; (3) adjustable 100G, 200G, 300G or 400G line-side transport links (using QPSK, 8QAM, or 16QAM modulation), which enables reach/capacity optimization over various fiber links; and (4) higher dispersion tolerance. All of these enhanced capabilities would exist in a QSFP-DD or OSFP module designed to utilize OpenZR+, supporting reaches well beyond the 120km supported by 400ZR. Table 1 summarizes the benefits of modules designed to use OpenZR+, compared to 400ZR and OpenROADM.

	400ZR	OpenROADM	OpenZR+
Target Application	Edge DCI	Carrier Metro/LH	Regional DCI
Client Traffic	400GbE only	100-400GbE & OTN	100-400GbE Multirate
Reach	<120km	>120km*	>120km*
Form Factor	QSFP-DD/OSFP	CFP2 or other	QSFP-DD/OSFP
SD-FEC	CFEC	oFEC	oFEC
Max Pwr Consumption	15W	Not Defined	20W*
Standards	OIF	OpenROADM	Leverages other standards work

*Reach and power consumption are dependent on modulation format and data rate.

Table 1. Next generation pluggable coherent module types

OpenZR+ Applications and Benefits

The higher performance of OpenZR+ can benefit hyperscale network operators by enabling a wider geographical extension of the physical connections to cover not only data center edge interconnects, but also regional interconnects (see Figure 3). Because modules designed using OpenZR+ are housed in QSFP-DD/OSFP form-factors, hyperscale operators can also reduce operational complexity by having one module type to cover 400ZR and OpenZR+ modes.

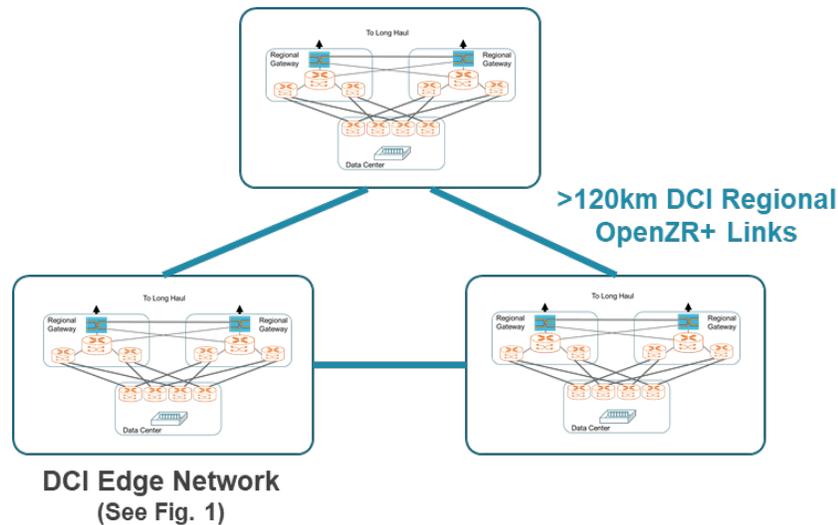


Figure 3. OpenZR+ enables the expansion of a DCI network's geographical footprint over regional distances beyond 120km in the same coherent pluggable form-factor as 400ZR.

Hyperscale networks with 75GHz DWDM grid channel spacing line systems can also benefit from modules utilizing OpenZR+. Because 400ZR was initially targeted to operate over 100GHz channels, a 400ZR transmission might suffer from degradation due to the roll-off frequency response of optical filters when deployed with 75GHz channels spacing. The additional gain from oFEC, specified in OpenZR+, can be used to compensate for these degradation effects over the 75GHz channel.

Efforts by carriers to transform or create a data center topology, such as CORD within the Open Networking Foundation (ONF), may also benefit from OpenZR+ utilizing this Ethernet-centric coherent optical solution in a similar fashion to a hyperscale edge/regional interconnect as an overlay to an existing carrier infrastructure.

OpenZR+ can also provide network operators with operational advantages. Because OpenZR+ specifications are derived from existing standards, network operators can benefit from interoperability among multiple suppliers that adopt OpenZR+. Module suppliers can also benefit from OpenZR+ by supporting applications beyond 400ZR, reaching a wider range of customer applications resulting in a larger addressable market.

Conclusion

Module vendors have quickly recognized the benefits of the OpenZR+ mode. For example, Acacia and NEL announced support for the OpenZR+ mode, and have exchanged DSP test vectors to confirm interoperability of their implementations. This is a milestone in the industry as it paves the way for the development of OpenZR+ modules in a QSFP-DD or OSFP form factor that support performance and functionality beyond 400 ZR.

OpenZR+ utilizes a well-defined common sense implementation of functionality that is already well defined in existing standards. OpenZR+ allows network operators to achieve all of the benefits of 400ZR with enhanced features that can expand the reach and functionality without sacrificing interoperability. Full OpenZR+ implementation details are available to vendors interested in supporting this market-expanding ecosystem.

References

¹To find out more about the origins and applications of 400ZR, please read [400ZR: Accessible 400G for Edge DCIs and Beyond](#)

About Acacia

Acacia's innovative silicon-based high-speed optical interconnect products accelerate network scalability through advancements in performance, capacity, and cost. Our silicon photonic PICs, DSP ASICs, and coherent modules inside a variety of network equipment products empower cloud and service providers to meet the fast-growing consumer demand for data.



3 Mill and Main Place, Suite 400
Maynard, MA 01754, USA
info@acacia-inc.com
978.938.4896
acacia-inc.com

All trademarks and registered trademarks are the property of their respective companies. WP1119