

XR Optics

Innovative Point-to-Multipoint Coherent that Slashes Aggregation Network TCO

Many networks, including metro aggregation, are inherently hub and spoke, with a large number of spoke devices (cell site, fiber node, RPD, DSLAM, OLT, etc.) connecting to a smaller number of hub devices (router, EPC/NGC, CMTS/CCAP, BNG, etc.). However, as shown in Figure 1, the optical technology used to build these networks has always been point to point, with an individual optical transceiver of the same speed required at each end. This mismatch results in a large number of inefficiently used optical transceivers and router ports, and multiple layers of packet aggregation, resulting in both high CapEx and OpEx. Furthermore, this problem will be exacerbated as access network evolutions, including mobile to 5G, cable to distributed access architecture (DAA), and passive optical networks to next-generation PON technologies, drive significant growth in both bandwidth and number of spoke devices.




Figure 1: Point-to-point optics vs. hub-and-spoke traffic

Introducing Point-to-Multipoint Coherent with Digital Subcarriers

Conventional optical technology provides a point-to-point wavelength between two locations, with each end required to operate at the same speed (1G, 10G, 25G, 100G, etc.), and two transceivers required for each connection, one at each end. For a hub router to communicate directly with N access nodes, 2N pluggables are required, N in the router and 1 in each of the N access nodes. Digital subcarrier technology takes a single-carrier wavelength and divides it up into multiple lower-bandwidth subcarriers generated by a single coherent laser/transceiver leveraging advanced digital signal processing.

For example, a 16 GHz 100G wavelength could become four 25 Gb/s subcarriers, a 64 GHz 400G wavelength could become 16 x 25 Gb/s subcarriers, and a 128 GHz 800G wavelength could become 32 x 25 Gb/s subcarriers, with the 25 Gb/s subcarriers each having a symbol rate of 4 GBaud with 16QAM modulation while occupying 4 GHz of spectrum. These subcarriers can now be routed to and from access transceivers in a variety of ways. A 400G hub transceiver could support 16 access nodes with a single 25 Gb/s subcarrier each, or four access nodes with four 25 Gb/s subcarriers each. It is also possible to have different numbers of subcarriers assigned to different access nodes depending on their bandwidth needs, as shown in Figure 2. XR optics implements this concept in industry-standard pluggable form factors such as SFP28, QSFP28, QSFP-DD, OSFP, and CFP2.

BENEFITS OF XR OPTICS

MINIMIZE the number of optical transceivers with the ability to aggregate multiple spoke devices onto a single hub transceiver

MAXIMIZE router efficiency, density, and simplicity by replacing large numbers of low-speed ports with far fewer high-speed ports, and with the ability to use these same high-speed ports as both aggregation and network interfaces

ALIGN CapEx with actual bandwidth requirements while still maintaining the ability to quickly adapt to changing bandwidth demands and traffic patterns

ELIMINATE intermediate packet aggregation stages while leveraging larger, more efficient switching devices at centralized sites

REDUCE OpEx in terms of power consumption, footprint, the number of aggregation sites, product support costs, and truck rolls

INNOVATION AWARDS

ECOC 2020 MOST INNOVATIVE PRODUCT

2020 LIGHT READING LEADING LIGHTS

2021 LIGHTWAVE INNOVATION REVIEW 4.5/5.0

2021 NGON & DCI WORLD BEST NEW GAMECHANGER OR INNOVATION

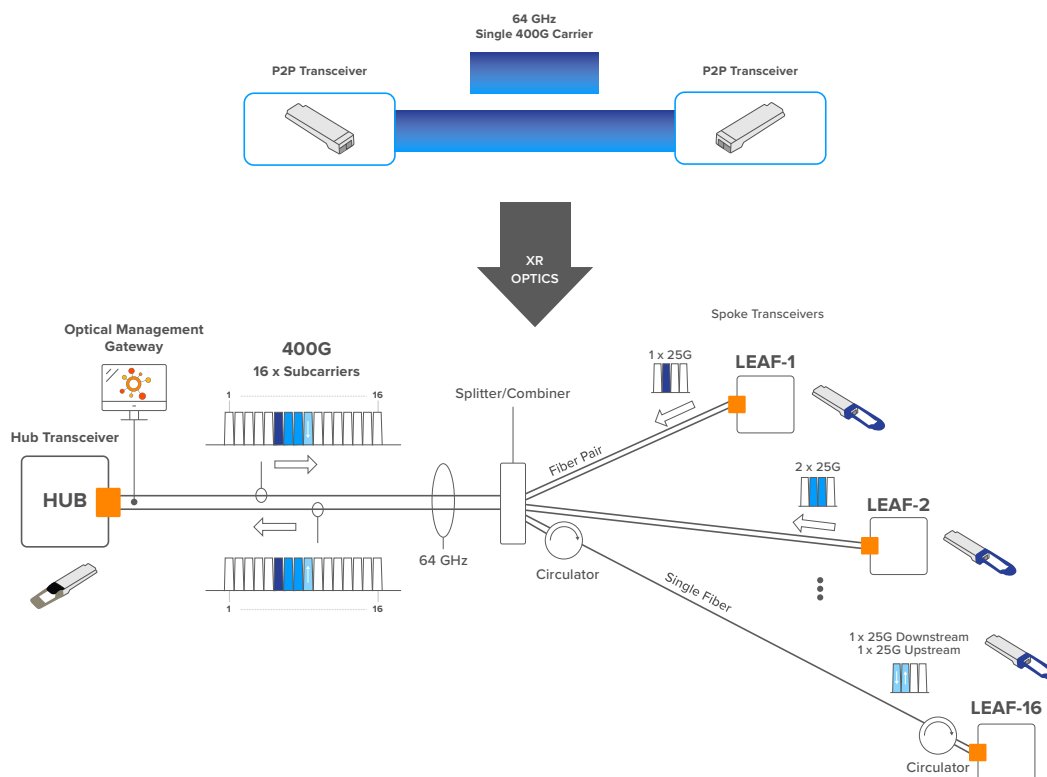


Figure 2: Point-to-multipoint coherent with 25 Gb/s digital subcarriers

Ideal for Hub-and-Spoke Applications Including 5G X-Haul, Cable DAA, and Next-Gen PON

XR optics with digital subcarrier technology provides an ideal solution for many inherently hub-and-spoke use cases. These include 5G fronthaul, where multiple radio units (RUs) connect to a distributed unit (DU); 5G midhaul, where multiple DUs connect to a centralized unit (CU); and 5G backhaul, where the CU or converged cell site (RU/DU/CU) is backhauled to the mobile core (NGC). Cable MSO examples include DAA, where many Remote PHY devices (RPDs) connect to a virtual Converged Cable Access Platform (vCCAP), or many Remote MAC-PHY devices connect to a router. Other possible use cases include backhaul for next-generation passive optical technologies such as XGS-PON and NG-PON2, and high-speed business services.




Figure 3: Example applications

Deployable Over Hub-and-Spoke, Ring-and-Chain Fiber Topologies

XR optics can leverage cost-effective splitter/combiner optical infrastructure to provide point-to-multipoint logical connectivity over hub-and-spoke, ring-and-chain physical fiber topologies, as shown in Figure 4. Once the subcarriers are combined as a multi-subcarrier wavelength, this wavelength can be transported over a flexible grid ROADM network, or a fixed grid DWDM infrastructure that aligns with the required spectrum (i.e., 100 GHz fixed grid for a 64 GHz 400G carrier). Distances of up to 1,000 km based on 25 Gb/s 16QAM subcarriers, and up to 4,000 km with 12.5 Gb/s QPSK subcarriers, enable the cost-effective centralization of hub resources, as shown in Figure 5.




Figure 4: Supported access fiber topologies

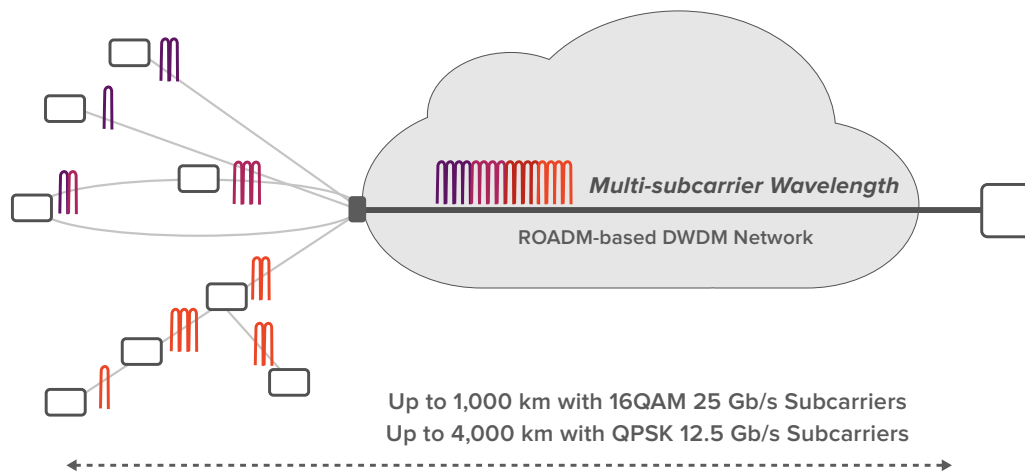


Figure 5: Multi-subcarrier wavelength transport over flexible grid ROADMs or fixed grid DWDM

Dramatically Reduce the Number of Optical Transceivers

Compared to scenarios where access nodes are connected directly to the hub router with conventional point-to-point WDM optics, the total number of optical interfaces reduces from $2N$ to $N + 1$, while the number of optical interfaces at the hub location can be reduced by a factor of up to 16 with a 400G interface and 25G subcarriers, and up to 32 with an 800G interface and 25G subcarriers. At the hub site, this can deliver CapEx savings of up to 80%, power savings in excess of 80%, and space savings in excess of 90%.

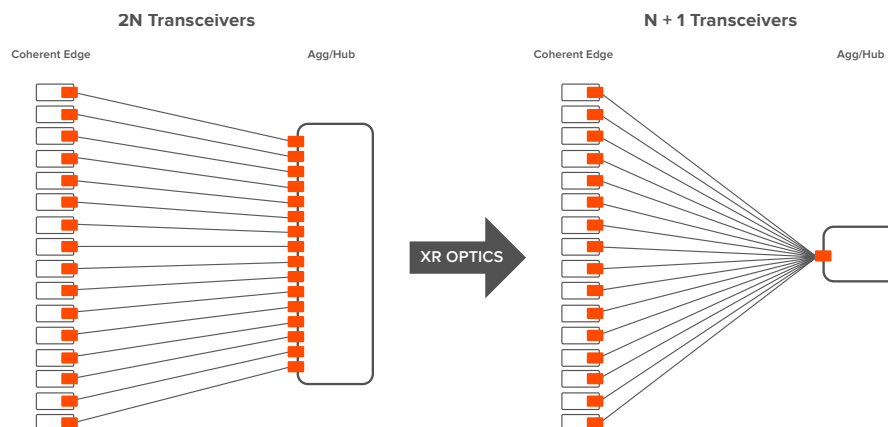


Figure 6: Reduce the number of optical interfaces from $2N$ to $N + 1$

Maximize Router Efficiency, Density, and Simplicity

Router CapEx, footprint, and power consumption are all optimized more efficiently by using router slots and ports. A large number of low-speed ports (i.e., SFP+, SFP28) can be replaced by a far smaller number of high-speed ports (i.e., QSFP-DD, OSFP), maximizing router faceplate density and processing efficiency. XR optics can also greatly simplify the hub router, as the same high-speed ports (i.e., 400G QSFP-DD or OSFP) can be used for both high-speed router-to-router network interfaces and aggregating multiple access nodes, with up to 16 access nodes per 400G port based on 25 Gb/s subcarriers.

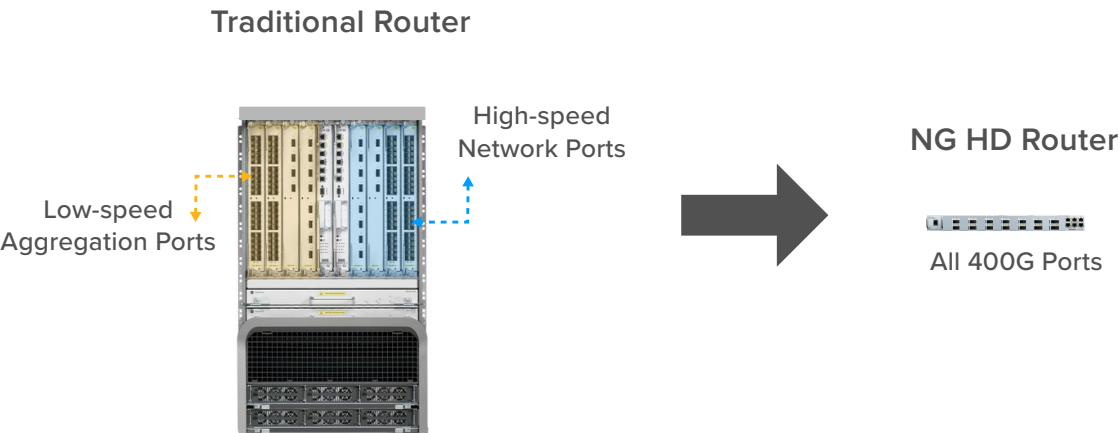


Figure 7: Maximize router efficiency, density, and simplicity

Align CapEx and Actual Bandwidth Requirements

By decoupling transceiver speed and bandwidth, XR optics enables the hub interfaces to more closely match actual bandwidth requirements rather than the sum of spoke transceiver speeds. Bandwidth to each access node can be sized based on current peak utilization rounded up to the nearest subcarrier value (i.e., 25 Gb/s) rather than expected future peak utilization rounded up to the next highest available line rate. For example, if current peak capacity is 18 Gb/s, and this is expected to grow to 85 Gb/s over the coming 24 months, conventional optics would require a 100 Gb/s interface for that one spoke at the hub device, while XR optics could be provisioned with a single 25 Gb/s subcarrier initially, then later upgraded to two and eventually four 25G/s subcarriers for a full 100 Gb/s of capacity, as shown in Table 1.

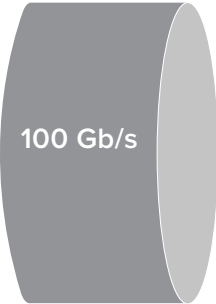

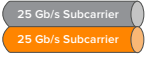
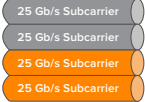
Peak Data Rate Over Time	Conventional Optics Solution	XR Optics with 25G Subcarriers
Day 1: 18 Gb/s		 } 25 Gb/s
12 Months: 48 Gb/s		 } 50 Gb/s
24 Months: 85 Gb/s		 } 100 Gb/s

Table 1: Hub capacity for an access node with 18 Gb/s growing to 85 Gb/s

This can save substantial CapEx at the hub site in terms of both transceivers and router ports. For example, 16 spoke sites, each with a 100G transceiver but initially needing only 25 Gb/s, can be aggregated on a single 400G hub transceiver/router port, as shown in Figure 8. This provides a more cost-effective solution compared to the alternative of 16 x 100G transceivers at the hub, yet the 100G spoke transceivers can be upgraded in the future without the need for truck rolls or new pluggables, as will be discussed in more detail in the next section.

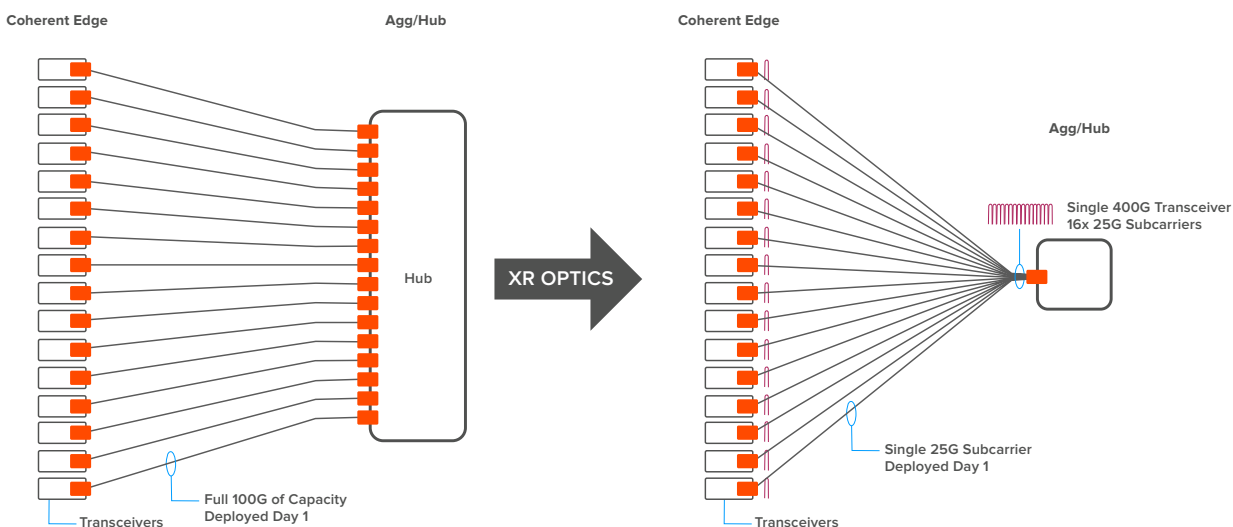


Figure 8: Aggregate 16 x 100G onto 1 x 400G

Seamlessly Adapt to Changing Bandwidth Demands and Traffic Patterns

In addition to the significant transceiver and router CapEx savings discussed in the previous section, decoupling transceiver speed and bandwidth provides the ability to quickly add or move bandwidth as traffic patterns change, while also significantly reducing the number of truck rolls and improving end-user service quality.




Figure 9: Upgrade scenario with spare hub port capacity

For example, as shown in Figure 9, where there are unused subcarriers at the hub, individual spoke nodes can be upgraded very quickly without the need for any truck rolls. In this way it is also possible to remotely reassign subcarriers between access nodes as bandwidth needs change.




Figure 10: Upgrade scenario with no spare hub port capacity

An alternative scenario is shown in Figure 10. Eight spoke nodes with 100G XR optics have 50 Gb/s each, fully utilizing the 400G hub interface. In this scenario, upgrading the spoke nodes to 100 Gb/s each only requires a single truck roll to the hub site to add a second 400 Gb/s XR pluggable, avoiding the need to visit any of the spoke sites to swap out pluggable optics.

Eliminate Aggregation Stages and Leverage Larger, More Efficient Aggregation Devices

XR optics provides the option to simplify the aggregation infrastructure by consolidating packet aggregation at a smaller number of sites with larger, more efficient packet aggregation devices, eliminating intermediate stages of aggregation. This is illustrated by the simplified example in Figure 11, which shows a present mode of operation with 25G aggregated onto 100G at eight aggregation sites, then multiple 100G aggregated onto 400G at the hub site. These aggregation stages provide two key functions: interworking different interface speeds (25G<->100G and 100G<->400G) and statistically multiplexing packet traffic. In this example, statistical gain, or oversubscription, is 8:1, with 4:1 at the aggregation sites and 2:1 at the hub. With XR optics, the need to interwork different interface speeds goes away, and statistical multiplexing can be centralized at the hub location, leveraging the economies of scale of the latest multi-terabit network processors and fabric ASICs. In the Figure 11 illustration, eight 500 Gb/s switches and one 1.2 Tb/s switch, totaling 5.2 Tb/s, can be replaced by a single 3.6 Tb/s switch with 8:1 statistical gain at the hub, with only simple splitter/combiner optical layer equipment required at the eight aggregation sites. 145 (128 x 25G + 16 x 100G + 1 x 400G) interfaces on the packet switches are replaced by only 9 x 400G interfaces, an almost 94% reduction, while the total number of transceivers is almost halved, going from 273 (256 x 25G + 16 x 100G + 1 x 400G) to 137 (128 x 25G + 9 x 400G).

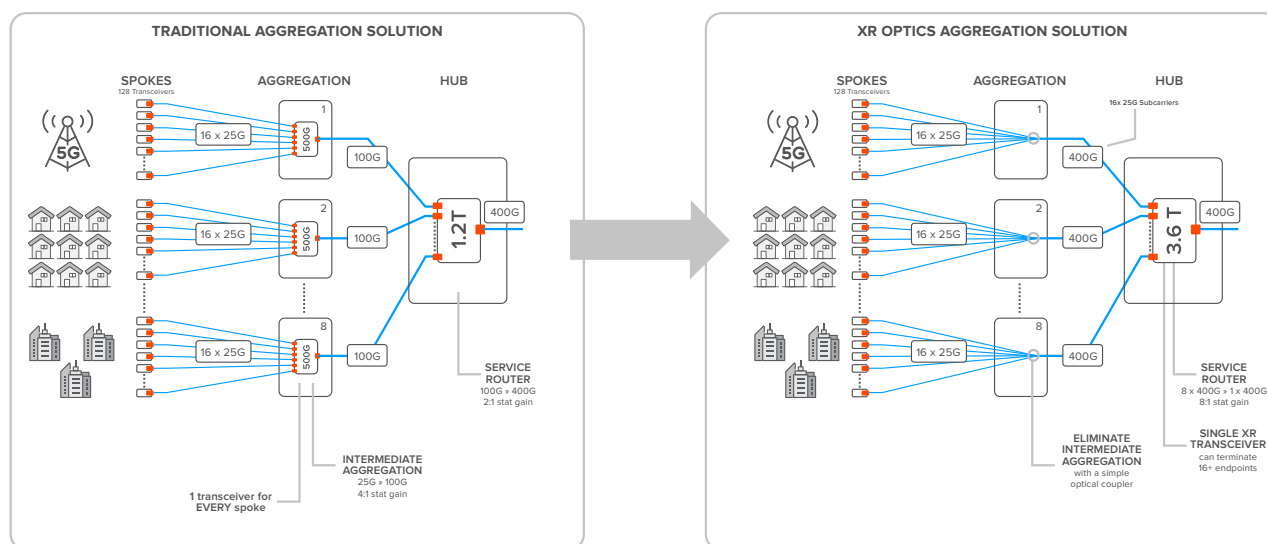


Figure 11: Simplified packet aggregation architectures

Reduce Operational Costs, Including Power, Space, Product Support, and Truck Rolls

Reducing the number of hub optical transceivers, maximizing router efficiency and simplicity, and eliminating intermediate stages of packet aggregation will significantly reduce power consumption and footprint. Additional OpEx savings relate to the product support costs paid to the equipment vendor, which are typically charged as a percentage of the initial price, and truck rolls, which are significantly reduced for bandwidth upgrades and reassignments. Other OpEx costs related to the amount and types of equipment, including planning, installation, commissioning, management, and maintenance, will also be reduced.

Summary: Gain a Competitive Edge and Accelerate the Deployment of Next-generation Services

As access networks evolve to 5G, DAA, and next-generation PON architectures, aggregation networks need to deliver an order of magnitude more bandwidth to a far larger number of locations and devices. Building these inherently hub-and-spoke aggregation networks with conventional point-to-point optics will result in large numbers of inefficiently used optical transceivers and router ports, as well as unnecessary intermediate layers of packet aggregation. Point-to-multipoint coherent based on digital subcarriers provides an innovative alternative, with XR optics implementing this concept in pluggable form factors such as SFP28, QSFP28, QSFP-DD, OSFP, and CFP2.

Combining CapEx and OpEx savings, both in excess of 70%, with the ability to quickly adapt to changes in traffic, XR optics will enable network operators to gain an edge over their competition and accelerate the deployment of these next-generation access networks and the new services they enable.