Optical line system evolution for the AI era

White paper

Optical networks are entering a new phase of capacity growth driven by cloud, video, data center interconnect (DCI) and—most significantly—AI infrastructure connectivity. At the same time, the Shannon limit is capping spectral efficiency gains, optical engines are evolving to higher baud rates and pluggable form factors, minimizing space and power is increasingly critical and customer applications and requirements are diversifying. These forces are pushing optical line systems (OLS) to evolve in both the short to medium term and the long term.

Contents

OLS evolution drivers	3
Network traffic growth	3
The Shannon limit	3
Coherent optical engine evolution	4
Fewer wavelengths (per band, per fiber)	4
Space and power constraints	4
Point-to-point architectures	5
Increasing application and customer requirement diversity	5
Short- to medium-term evolution	5
Enhanced performance	5
More spectrum: C+L, Super-C, Super-L	5
Multi-rail	6
Direct attach and low-port-count CDC	6
Reduced footprint	7
Extended temperature range	8
Operational features	8
Open optical networks	8
Potential longer-term evolutions	9
Beyond C and L (O, E, S, U)	9
Support for new fiber types: HCF and MCF	9
Full-spectrum transponders with integrated OLS functions	10
Fiber switching and hybrid wavelength/fiber switching	10
One size does not fit all	10
Meeting the diverse needs of the AI era	11

OLS evolution drivers

Network traffic growth

According to a recent Nokia Bell Labs study, while traditional non-AI consumer and enterprise traffic is growing at a CAGR of 16%, AI-related consumer and enterprise traffic is growing at a CAGR of 24%. In addition, AI and cloud providers are seeing dramatic traffic growth related to AI infrastructure. AI clusters are growing to millions of GPUs, which, due to power and space constraints, are being distributed across multiple data centers. Scale across connects AI backend networks across multiple data centers for training and inference (Figure 1), and is therefore being adopted rapidly. Scale across bandwidth requirements are typically an order of magnitude greater than traditional DCI.

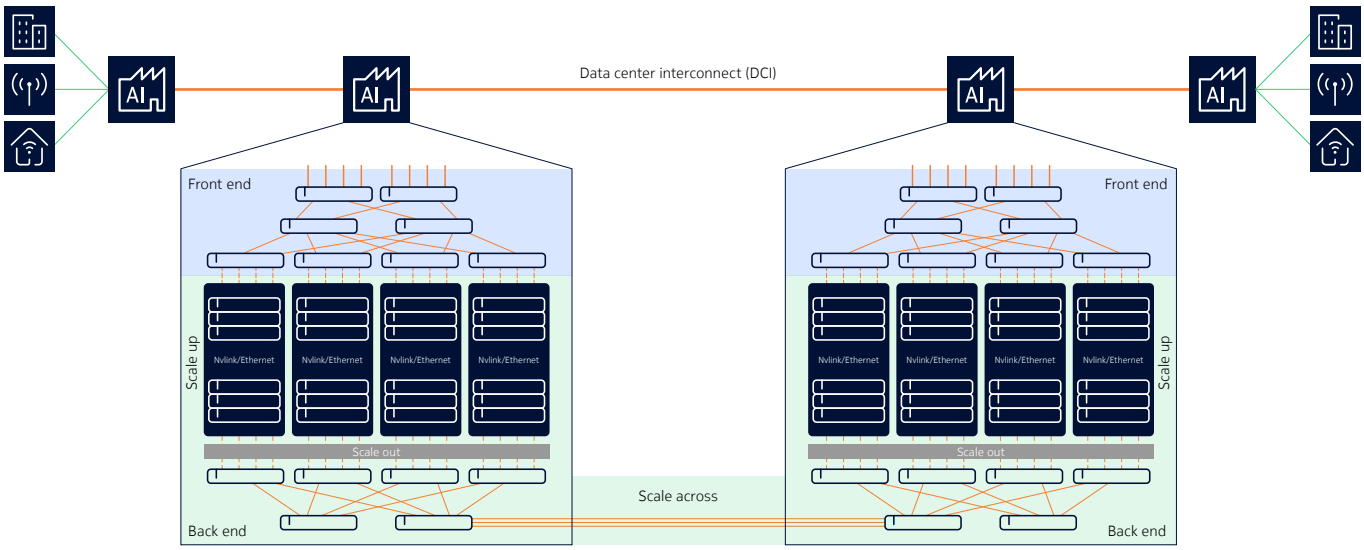


Figure 1 – Scale across connects the GPU backend while DCI connects the frontend

AI and cloud providers' core networks are also feeling the impact when data is moved to AI clusters and when trained models are transferred to inference clusters for monetization. Furthermore, service requests that were previously handled locally by caches now have to travel back to more centralized inference data centers. This was highlighted by a leading AI and cloud provider who presented data at optical conferences in 2024 and 2025 showing their backbone traffic was growing at a CAGR of 100%.

The Shannon limit

Shannon's law, named after Claude Shannon of Bell Labs (now part of Nokia), defines the maximum amount of information that can be communicated over a channel with a given bandwidth and noise. This is described by the equation C = B x Log2(1+SNR), where C is the maximum achievable channel capacity in bits/s, B is the bandwidth in Hz, "Log2 (1+SNR)" is the spectral efficiency in bits/s/Hz, with SNR the ratio of signal power (S) to noise power (N).

At very large SNR values, where we can ignore the 1 in the equation, increasing spectral efficiency by adding 1 bit/s/Hz requires us to double (+3 dB) the required SNR, thus halving the reach. This provides an upper bound on the spectral efficiency that can be achieved. There is also a lower bound based on nonlinearities, arrived at by simulations, sometimes controversially called the "nonlinear Shannon limit". Today's state-of-the-art, high-performance embedded optical engines are typically between 1 and 2 dB from the (linear) Shannon limit, meaning any future spectral efficiency gains are likely to be incremental at best.

Coherent optical engine evolution

Baud rates for Nokia embedded coherent optical engines have increased from 28 GBaud in 2010 to 140+ GBaud in 2025, with channel sizes going from 50 GHz to 150 GHz (Table 1). At the same time, coherent pluggables have evolved from ~30 GBaud in 2014 (100 Gb/s CFP) to 60+ GBaud in 75 GHz with 400G pluggables around 2020, then to 130+ GBaud in 150 GHz with 800ZR/ZR+ in 2025. As the boundary between pluggable technology and embedded engine blurs, the next generation for both form factors will target 250+ GBaud, typically with 300 GHz channel sizes.

Year	Name	Max data per λ	Symbol rate	Typical channel	Spectral efficiency[1]
		Gb/s	GBaud	GHz	Bits/s/Hz
2010	First gen	100	28	50	2
2013	PSE	200	32	50	4
2016	PSE-2s	400	45	62.5	6.4
2018	PSE-3s	600	67	87.5	6.86
2020	PSE-Vs	600	90	100	6
2020	ICE6	800	96	112.5	7.11
2022	ICE6 Turbo	800	100.4	125	6.4
2023	PSE-6s	1200	130+	150	8
2025	ICE7	1200	140+	150	8

1. Typical spectral efficiency based on max data rate and typical channel spacing

Table 1 - Nokia embedded coherent baud rate evolution

While embedded engines will have a role long into the future, the industry is shifting towards coherent pluggables. These pluggables may have lower transmit power relative to embedded engines. Consequently, the OLS must be able to support different generations of pluggables and embedded engines, including heterogeneous deployments with multiple generations of each. Coherent pluggables also typically have lower spectral efficiency relative to embedded engines, which puts additional pressure on the OLS to increase capacity in other ways.

An additional consideration is that as baud rates increase, transmit power of coherent interfaces typically does not. Transmit power per Hz and net transmit power coming into the OLS is therefore decreasing. The terminal OLS, therefore, needs to provide more per-channel amplification in the add/drop or line booster amplifiers.

Fewer wavelengths (per band, per fiber)

As coherent optical engines evolve to higher baud rates, wavelength data rates are increasing while the number of wavelengths required to fill a band (C or L) or the whole fiber (C+L) is decreasing. For example, with ~30 GBaud and 50 GHz channels, 192 wavelengths are required to fill the 9.6THz of C+L. With 250+ GBaud and 300 GHz channels, a 9.6 THz C+L fiber will be filled with only 32 channels.

Space and power constraints

Space and power constraints are also driving OLS evolution, especially at inline amplifier (ILA) sites, which are typically small huts with room for only a single 40 RU rack and a maximum power availability of 3 KW. In addition, there is a continued need to lower power consumption, footprint and cost across the optical network to reduce total cost of ownership.

Point-to-point architectures

Another OLS evolution driver is the transition to point-to-point optical architectures by several large AI and cloud providers for applications such as long haul, which traditionally used ROADMs for optical passthrough. These networks instead rely purely on the IP layer for switching between degrees.

Removing ROADMs from the wavelength's path also has performance benefits. This is due to reduced wavelength selective switch (WSS) filter narrowing and improved SNR, as noise from the amplification that compensates for ROADM component losses (e.g. 2~8 dB per WSS, 6dB for a 1:4 splitter) is also reduced.

Increasing application and customer requirement diversity

Another key trend driving OLS evolution is the growing diversity in terms of applications and customer requirements. The typical requirements of large AI and cloud providers and telecommunication providers have diverged. One driver for this is the different services they offer. For example, telecommunication providers typically offer wavelength services, which often need to be segregated and protected at the optical layer. Conversely, AI and cloud providers usually build large IP networks, which provide protection at the IP layer.

Another driver is the sheer volume of traffic, with AI and cloud providers typically requiring orders of magnitude higher capacity for data center interconnect and AI infrastructure. Mission-critical enterprise requirements are often different again. Applications have also evolved from metro, long haul and submarine, with the additions of DCI, managed optical fiber networks (MOFNs) and, most recently, scale across.

Short- to medium-term evolution

Enhanced performance

As embedded coherent optical engines approach the Shannon limit, OLS have evolved to best use the existing spectrum with multiple performance enhancements. Erbium-doped fiber amplifiers (EDFAs) have been developed with lower noise, especially at longer span losses/higher gains. ILAs have integrated dynamic gain equalizers (DGEs), with a DGE/ILA typically deployed every five ILAs in long-haul networks to boost performance in longer ILA chains.

ROADMs have evolved flexible grid granularity from 12.5GHz to 6.25GHz, which in some scenarios enables better alignment of the wavelengths' spectrum and channel sizes, resulting in improved spectral efficiency. Link control algorithms that are responsible for optimizing the wavelength power levels and amplifier gains have been enhanced to more optimally trade-off OSNR and nonlinear penalties on a per wavelength basis.

More spectrum: C+L, Super-C, Super-L

As we approach the Shannon limit, OLS have advanced to enable higher fiber capacity by expanding the available spectrum, increasing the B in Shannon equation. This option is especially attractive when fiber availability is constrained and the cost of new fibers is prohibitive. Beyond the 4.8 THz of the extended C-band, the most widely deployed option is to light the L-band, giving a total of 9.6 THz. The L-band has only marginally higher loss than the C-band and EDFA amplifiers work in the L-band with only marginally higher cost and lower power efficiency. One key challenge with C+L line systems is stimulated Raman scattering (SRS) tilt, with power shifting from the C-band to the L-band. Solutions include using ASE noise to fully load both bands from day one and/or advanced link control that can quickly rebalance the power after any network changes.

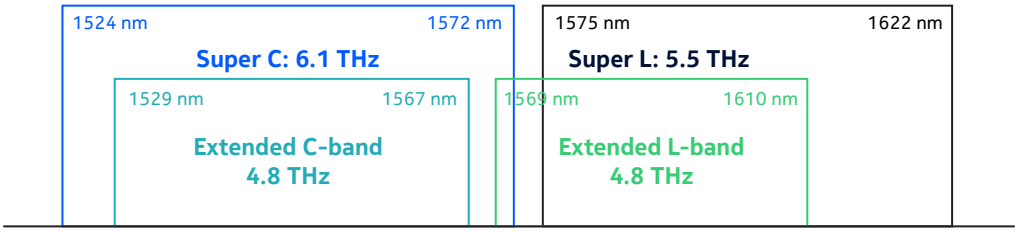


Figure 2 - Increased spectrum with L, Super C and Super L

A second option is Super C, expanding the C-band from 4.8 THz to 6.1 THz. This can provide a cost-effective option for increasing capacity without the full cost (i.e., typically >2x) and complexity (e.g., SRS tilt) of C+L. Furthermore, Super L provides the option for an additional 5.5 THz, taking the total spectrum in a Super C plus Super L system to 11.6 THz, a greater than 20% increase relative to the 9.6 THz of C+L.

Multi-rail

These additional spectrum options cannot do much more than double fiber capacity. For the leading AI and cloud providers, AI-related capacity requirements demand much greater scalability. Their regional data centers are now typically connected with 128 or more fiber pairs, which can again be viewed as increasing the B in the Shannon equation. Lighting these fibers becomes especially challenging at space- and power-constrained ILA huts. To address this, OLS vendors such as Nokia are introducing multi-rail ILAs that support bidirectional amplification for multiple fiber pairs in a single ~1RU card.

Cost, power and footprint gains come from sharing components, such as optical time-domain reflectometry (OTDR), optical channel monitoring (OCM) and digital gain equalization (DGE) and from using low-power uncooled and/or multi-chip pump lasers. For example, the multi-rail ILA Nokia announced at OFC 2026 enables up to 160 C+L ILAs in a 40 RU (600 mm depth) rack with 4 ILAs per RU. This compares to current ILA technology, which typically requires 8 RU (300 mm depth) for four C+L ILAs. At the same time, power consumption per ILA is reduced by more than 60%.

Multi-rail ILAs, therefore, deliver significant density improvements while also providing power consumption and cooling benefits. In addition, multi-rail ILAs have a diverse range of requirements: 600 mm or 300 mm, C-band only or C+L, EDFA or hybrid EDFA/Raman. And while the initial focus is on multi-rail ILAs, multi-rail terminals and ROADMs will also be required. One option for multi-rail terminals is to combine multi-rail ILAs with wide passband mux/demux filters. One option for multi-rail ROADMs is to partition a single physical WSS component (e.g., an LCoS array) into multiple, typically four or eight, “virtual WSSs”.

Direct attach and low-port-count CDC

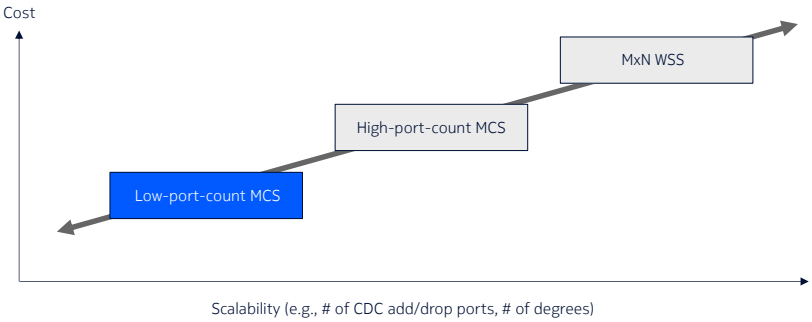
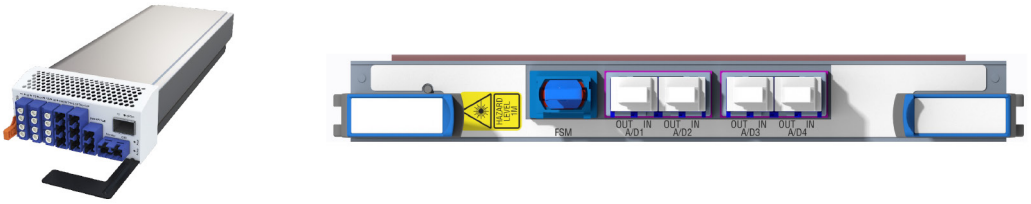


Figure 3 - CDC technology options: cost vs. scalability

As baud rates increase and number of wavelengths per fiber decreases, new ROADM add/drop architectures become attractive. For example, cost-effective “low port count” colorless directionless, contentionless (CDC) leveraging unamplified multicast switch technology. Nokia examples include the compact CDC8D6 for the 1830 GX with 8 directions and 6 add/drops in a single slot card, and the MCS44 with 4 directions and 4 add/drops for the Nokia 1830 PSS/PSI-L (Figure 4).



CDC8D6 for 1830 GX

MCS44 for 1830 PSS/PSI-L

Figure 4 – Low-port-count CDC examples

Another newly attractive option is colorless direct attach, where the transponder is directly attached to the ROADM’s WSS port without an intermediate add/drop layer. This architecture is also enabled by increases in the number of WSS ports to 20, 32 and 66. In addition to a low-cost, low-footprint add/drop option, this offers the benefits of colorless, flexible grid add/drop and improved optical performance. For example, the RD66 of the 1830 GX hyperscale OLS solution (Figure 5), uses a high port count C+L WSS to support direct attach of C or L band coherent interfaces on any of its 66 add/drop ports.



Figure 5 – Nokia 1830 GX RD66 with 66 C+L flexible grid, colorless direct attach ports

Reduced footprint

OLS are continually evolving to reduce footprint. This has been possible because the size of key ROADM components have shrunk, including WSSs and amplifiers, and the availability of functions such as OSC and OTDR as compact pluggables. The evolution to twin WSS on a single unit, as well as “ROADM-on-a-blade” architectures have also contributed. Furthermore, it has been enabled by the availability of compact modular platforms with a 600 mm depth.

For example, around 2010, long-haul ROADMs typically required around 6R U per degree, while metro ROADMs were around 3 RU per degree. Modern compact modular platforms can deliver two ROADM degrees in 1 RU. As discussed, multi-rail ILAs are dramatically reducing ILA footprint, while direct attach and low-port-count CDC are reducing add/drop footprint. The Nokia 1830 GX OLS with the RD66, for example, provides C and L EDFA pre-amplifiers, C and L EDFA booster amplifiers, a twin C+L 1x66 WSS, integrated ASE noise, OSC, OCM and OTDR, all in 2 RU, and with an additional 2 spare slots that could be used for Raman amplifiers or transponders.

Extended temperature range

Another evolution vector for OLS evolution is extended temperature range (ETR). This has a number of benefits and applications. As OLS are deployed in metro access networks, including street cabinets that are not air-conditioned, ETR is necessary. A second application for ETR is reducing environmental controls in the ILA huts of long-haul networks. Higher availability is an additional benefit, as devices can still operate even if the cooling system fails. ETR is supported in the Nokia 1830 XTM, and in Nokia 1830 PSS configurations including new ILA and four-degree ROADM options.

Operational features

Several innovations reduce operational cost and enhance manageability. Installation has been simplified, with zero-touch provisioning, cabling verification, auto-discovery of transponders and coherent pluggables in routers and switchable gain amplifiers. High-performance link control (in-band signalling between OLS nodes to set amplifier gain and wavelength power levels) no longer requires complex planning and configuration. OTDRs have become a standard feature, with use cases that include identifying the location of fiber cuts, detecting increased fiber loss, and intrusion detection. OTDRs are now available as OTDR modules, compact pluggables or integrated into line system modules such as ROADM-on-a-blade or ILAs.

Open optical networks

OLS have become more open, with integrated OCMs and WSS- or DGE-based attenuation simplifying support for third-party wavelengths, while flexible-grid capabilities provide a path to spectrum services. Management interfaces such as TL1 and SNMP have evolved to open APIs (e.g., NETCONF and RESTCONF) with OpenConfig and OpenROADM YANG data models. Pull-based management protocols such as SNMP are evolving to push-based (e.g., gNMI, gRPC) with streaming telemetry.

Potential longer-term evolutions

Beyond C and L (O, E, S, U)

As discussed, C and L bands can provide up to 11.6 THz with Super C plus Super L. Amplifier innovation may yet increase Super L spectrum from 5.5 THz to 6 or more THz, taking total capacity to 12.1 or more THz. However, for fiber-constrained network operators, maximizing the capacity of deployed single-mode fiber (SMF) by using additional bands could be very attractive. These bands include the O-band, E-band, S-band and U-band, with their SMF properties summarized in Table 2.

Band	O	E	S	C	L	U
Wavelengths (nm)	1260-1360	1360-1460	1460-1530	1529-1567	1569-1610	1625–1675
Loss (dB/km)	0.32~0.38	0.24~0.30	0.20~0.24	0.18~0.20	0.20~0.22	~0.25~0.40
CD (ps/nm/km)	~0	5~10	10~15	16~18	18~21	21~25
Amp Dopant(s)	Bismuth	Bismuth	Thulium, Bismuth	Erbium (EDFA)		Bismuth
Amplifier Cost	250%	250%	200%	100%	120%	-
Amp Efficiency	2%	10%	13%	25%	12%	-

Table 2 - Band comparison: O, E, S, C, L and U (for SMF)

Adopting these additional bands, however, faces significant challenges. For the O-band, near-zero chromatic dispersion (CD) limits coherent performance, as high CD is needed to mitigate nonlinearities including cross-phase modulation (XPM) and self-phase modulation (SPM). The O-band, E-band, U-band and, to a lesser extent, the S-band have high loss relative to C and L bands. Furthermore, the E-band can suffer from very high loss (2~3dB/km) in older SMF with high water peak absorption. The U-band is also limited by low bend tolerance.

Amplification is also a challenge. Bands beyond C and L require dopants such as bismuth and thulium rather than the erbium used for C and L. These amplifiers have much higher cost and deliver lower power conversion efficiency (Table 2).

A key challenge with using these bands alongside C and L is SRS tilt. For example, just as the C-band pumps the L-band (as discussed), the S-band pumps the C-band. SRS tilt is therefore likely to limit which combination of bands can be used, and in turn the total useable spectrum.

Existing deployments add further complication, as some already use these bands for the OSC (1510 nm in the S-band) and OTDR (1625 nm in the U-band). Going beyond C and L also faces significant ecosystem challenges.

The next section discusses how hollow core fiber could be the catalyst for using bands beyond C and L.

Support for new fiber types: HCF and MCF

Unlike traditional fiber which has a glass core, hollow-core fiber (HCF), as the name implies, is hollow and filled with air or, more typically, a gas. Its refractive index is around 1, meaning light travels 50% faster and latency is reduced by approximately 30 percent compared to solid core fibers such as SMF. In the short-term, this lower latency is the primary driver for HCF adoption. However, HCF also has the benefit of low loss, with 0.04 dB/km the current state of the art. Nonlinearities are minimal, which has the potential to enable much higher launch power into the fiber. This higher power could be used to extend the distance between ILAs and require fewer ILA sites overall.

Alternatively, if traditional ILA spacing (e.g., 60 - 80 km) is maintained, this increased launch power could be used for higher OSNR at the receiver, enabling increased wavelength capacity and better spectral efficiency. HCF also offers much wider low-loss spectrum, and the absence of SRS tilt simplifies the use of bands beyond C and L.

Today's OLS can be deployed with HCF using SMF-HCF adaptors, which may come integrated with the HCF. However, to take full advantage of the benefits of HCF, future OLS will need to add the following capabilities:

- Higher power amplification, to benefit from HCF's low nonlinearities.
- OTDRs with much higher transmit power (and/or dynamic range), to overcome HCF's low backscattering.
- Support for bi-directional single fiber, to benefit from HCF's low backscattering.
- Support for additional or different bands (O, E, S, U), including WSS and amplification.

Another novel fiber type gaining traction is multi-core fiber (MCF). These are typically uncoupled MCFs with between two and four cores. To adopt MCF, OLS could evolve with native MCF connectors, efficient MCF amplification (e.g., shared cladding, cladding-pumped) and core selective switching (CSS).

Full-spectrum transponders with integrated OLS functions

As the larger AI and cloud providers deploy multi-rail line systems, their capacity requirements are reaching levels where each fiber pair is fully loaded on day one. This is driving the industry to develop full-spectrum transponders (FSTs), which are a single transponder card with multiple coherent engines capable of filling an entire band (e.g., 16 x 300 GHz 1.6T) or full fiber (e.g., 32 x 300 GHz 1.6T) in a compact shelf.

A key driver for FSTs is faster and simpler deployment. They can also take advantage of integration (e.g., co-packaged optics, client/line DSP integration, double-sided pluggables/media converters, large scale photonic integrated circuits, comb lasers), and innovative multi-wavelength modulation techniques to deliver significant cost, power and footprint savings. From an OLS perspective, FSTs typically integrate line system functions such as mux/demux and EDFAs, significantly simplifying the OLS at the terminal sites.

Fiber switching and hybrid wavelength/fiber switching

As the industry shifts to multi-rail and FSTs, switching will evolve from wavelength-based ROADMs using WSS to fiber-based switching using optical circuit switches (OCS). Hybrid wavelength/fiber switching is a likely transitional step, offering significant scalability, Capex and power consumption benefits over wavelength-only switching as capacity scales in high-capacity core networks.

One size does not fit all

As discussed, OLS applications and customer requirements are becoming increasingly diverse. Some customers want 300 mm depth with side-to-side airflow, while others request 600 mm and front-to-back. Some customers are evolving to multi-rail and FSTs, while others find 400 Gb/s still exceeds their capacity requirements. For many customers, multi-degree ROADM provides the flexibility they need while others are evolving to point-to-point architectures, keeping switching at the IP layer.

Some operators want open APIs and streaming telemetry to create their own network management solutions while others prefer to buy the network management platform from their optical equipment vendor as part of a holistic solution. For some operators, disaggregation is a key enabler for fast innovation and avoiding vendor lock-in, while for others minimizing space and power with greater functional integration is a higher priority.

Meeting the diverse needs of the AI era

Traffic growth is accelerating, amplified by AI infrastructure buildouts. At the same time, spectral efficiency gains are starting to run into the Shannon limit. In response, OLS are evolving over the short- to-medium term and long term to increase capacity, reduce cost, space and power, simply operations and increase availability (Table 3). In the near term, OLS innovation is focused on scaling capacity by improving coherent performance, increasing spectrum and by adopting multi-rail architectures. Additional near-term innovations include direct-attach and low-port-count CDC add/drop, extended temperature ranges, operational enhancements and open optical.

		More capacity	Reduced cost/ space/power	Simplified operations	Higher availability	Other
Short-medium term	Enhanced performance	•				Coherent alignment
	Super C, L-band, super L	•				Leverages existing fiber
	Mult-rail	•	•			
	Direct attach and LPC CDC		•	•		Coherent alignment
	Extended temperature range		•	•	•	Enables edge deployments
	Operational enhancements			•	•	
	Open optical	•	•	•		Faster innovation
Long term	Bands beyond C & L	•				Leverages existing fiber
	Multi-core fiber (MCF)	•				Reduced cable weight/size
	Hollow-core fiber (HCF)	•	•			Lower latency
	Full-spectrum transponder	•	•	•		
	Fiber/hybrid switching	•	•		•	

Table 3 - OLS evolution benefits

Longer-term evolutions may include bands beyond C and L, optimizations to maximize the value of new fiber types (HCF, MCF), FSTs with some integrated line system functions and a gradual shift toward hybrid wavelength/fiber switching and eventually fiber-only switching. Nokia recognizes these trends and has the R&D resources to evolve its broad OLS portfolio to meet the diverse requirements and applications of AI and cloud providers, telecommunication providers and mission-critical enterprises.

About Nokia

Nokia is a global leader in connectivity for the AI era. With expertise across fixed, mobile, and transport networks, we're advancing connectivity to secure a brighter world.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2026 Nokia

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: CID 215480 (June)