Optimizing subsea-terrestrial networks

Architecture options to enable subsea cable operators to optimize their end-to-end subsea-terrestrial networks

Application note

Abstract

This white paper provides an overview and benefits of new architecture options for subsea network operators to better optimize their end-to-end subsea and terrestrial network architectures. A first architecture option enables subsea network operators to simplify cable landing stations and reduce end-to-end network cost and space and power usage at the CLS. Other options presented leverage use of L-band WDM spectrum in the terrestrial backhaul network to either lower lease costs for subsea traffic backhaul, or to make maximum use of the terrestrial fiber for serving a range of network traffic types.

Contents

Introduction

The increasing importance of subsea fiber networks in the AI era

Fiber optic networks provide the connectivity ubiquitous to modern society, enabling transmission of data needed for financial transactions, business critical traffic, national defense, entertainment and personal interactions across the globe. Increasingly data transmission is global, inter-connecting continents and island nations, made possible by subsea fiber cables that extend that connectivity across seas and oceans.

The massive data-carrying capacity of optical fiber has led to wide deployment of subsea fiber cables. Today over 570 in-service subsea cables carry >99% of all international traffic and connect data centers around the globe to deliver cloud-based and emerging AI services. This connectivity is being significantly expanded to support the massive global data center build-out currently underway to satisfy the compute needs of AI training and AI-based inferencing, agents and applications.

A large portion of these new subsea cable build-outs are being led by hyperscale, neo-cloud and AI companies in order to provide global connectivity to both their own, partner, and third-party data centers. These data centers, both existing and those that will soon come on-line, will house ever-larger clusters of inter-connected xPUs, which serve as the engines for machine learning and AI training and inferencing.

Modern AI training and large language models (LLMs) span a range of connectivity requirements. The first is a “scale-up” design, where tens to thousands of xPUs are interconnected as a single logical unit, and the resulting interconnect bandwidth is soon expected to exceed 1 pettabit per second per cluster. These xPU clusters are inter-connected to others in a data center using “scale-out” bandwidth. Combined these interconnect requirements multiple orders of magnitude greater intra-data center bandwidth than the historical norm needed by the server-based internet and cloud traffic of traditional data centers.

The explosion of intra-data center connections will also impact the bandwidth needed for inter-data center interconnection (DCI). Limits on maximum DC size and improvements in training protocols will require, and enable, connectivity for AI training and inferencing to be distributed across multiple DCs using “scale-across” architectures. This will drive bandwidth growth across metro and long-haul terrestrial fiber networks, and will extend across trans-oceanic subsea cables as well.

Unlike historical consumer, business and telco traffic, this AI-driven bandwidth, combined with continued growth in cloud services, is delivered from and to data centers (see figure 1). Today fiber optic networks including subsea fiber cables connect over 11,000 data centers, with over one thousand being large “hyperscale” data centers, currently supporting over one thousand petabytes per day of WAN traffic with DC capacity expected to grow at 27% CAGR until 2030. Nokia Bell Labs is forecasting that global AI WAN traffic will add over 1,088 exabytes/month of capacity by 2033, growing at a 24% CAGR on top of continued growth in cloud-based services. Combined this represents a continuation of the industry trend of approximately 30% per year growth in bandwidth transmitted over subsea cables.

Figure 1: An ever-increasing number of data centers powering AI services is driving significant bandwidth growth over subsea fiber optic cables.

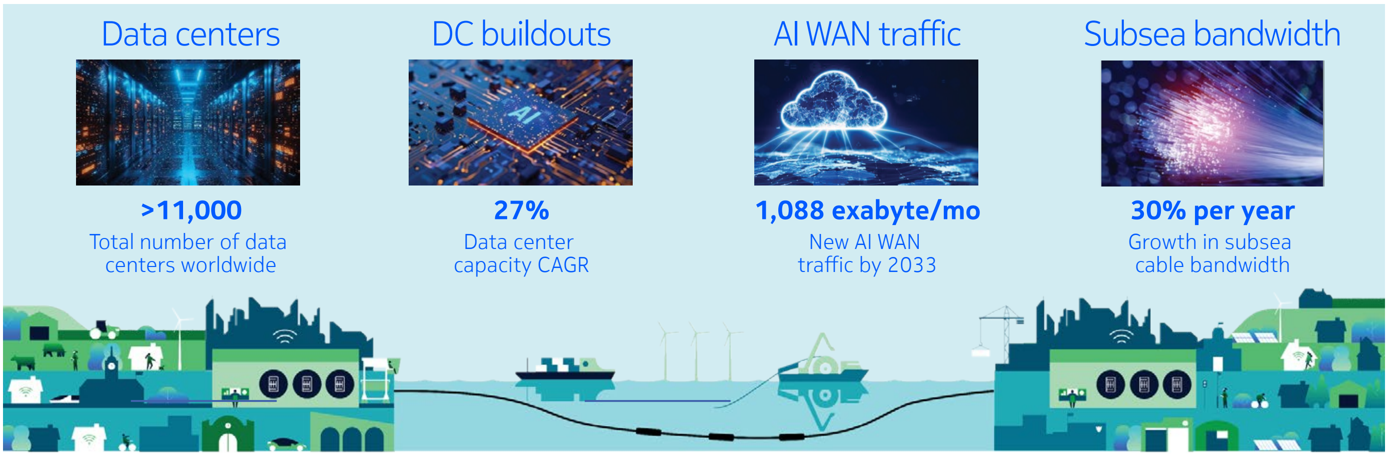

More fundamentally however, the continued growth in cloud-based services, augmented and accelerated by new AI services, is leading to fundamental changes in both the network operators involved and the network architectures optimized for global end-to-end subsea terrestrial networks. By optimizing their subsea and terrestrial backhaul networks using on the latest technologies and SLTE solutions, subsea network operators can simplify their networks, reduce CapEx and operational costs, and maximize their network capacity.

This article will describe these dynamics, and what changes are occurring in subsea networks, and provide an overview, and benefits, of new architecture options for subsea network operators, and enable them to better optimize their end-to-end subsea and terrestrial networks architectures.

How to optimize subsea cable landing stations for the AI era

Taking a step back in time to 1866, when the first practical trans-oceanic copper cable was put into operation connecting Ireland to Newfoundland, Canada, one finds that the cable landing station (CLS) performed a key role. To send data end-to-end, an operator transmitted the telegraph signal at one end of the cable, while another operator received the signal at the other end of the trans-Atlantic cable. Once decoded, the resulting message was then handed off to another telegraph operator, who would re-transmit the message via telegraph over a terrestrial copper cable to its eventual desired destination, such as New York City or London, for example.

Fast-forward 150 years to the mid-2010's, and trans-oceanic cables evolved from copper to fiber. The dominant operators of subsea networks were international telecommunications providers, bandwidth wholesalers or subsea cable consortium, who funded and led the deployment of new cables.

While data rates had increased tremendously since 1866, from approximately 8 words per minute to 12Tb/s per fiber pair, little had changed in the CLS architecture. Subsea traffic continued to be terminated at the CLS with transponders, and was then connected via back-to-back client interfaces to another set of transponders, which re-transmitted the subsea traffic to its ultimate destination over a terrestrial backhaul network.

However the end-to-end bandwidth flows of today's cloud and AI based data traffic is different. First, for the past decade or more, the majority of new subsea fiber cable build-outs have been led by AI and cloud providers to connect data centers spread across different continents. The second major difference derives from the fact that traffic across subsea cables no longer originates or ends at a CLS; it flows end-to-end between data centers that may be located meaningful distances away from shore. This is leading AI and cloud providers to optimize their subsea and terrestrial backhaul networks with these end-to-end connections in mind, encompassing everything between the data centers at each end.

This has led to a change in the end-to-end subsea-terrestrial network architecture, where back-to-back transponders at the CLS can be removed, and subsea traffic is transparently passed through in the optical domain, using reconfigurable optical add/drop multiplexers (ROADMs) at the CLS. With this architecture, wavelengths transiting across the subsea fiber extend directly across the terrestrial backhaul network, all the way to transponders located in the data centers at each end (see figure 2).

Figure 2: Optimizing subsea-terrestrial networking by connecting wavelengths end-to-end between data centers using optical pass-through at cable landing sites.

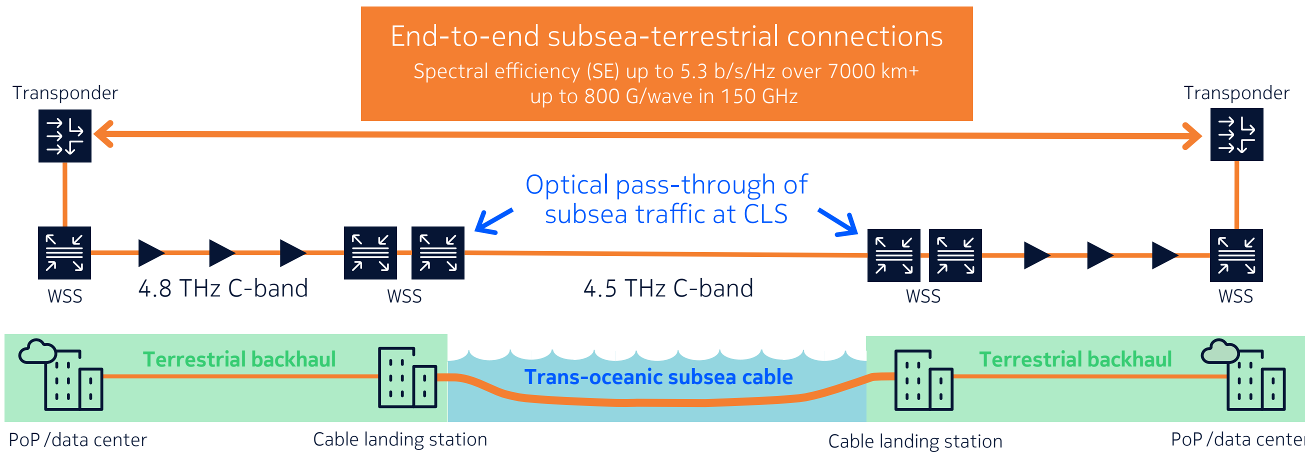

This optical pass-through of wavelengths transmitted across the subsea fiber is enabled using the latest generation of ROADMs supporting subsea/SLTE-optimized features, such as ASE spectrum power insertion into the subsea fiber, spectrum sharing and optical channel monitoring, and features in the terrestrial backhaul network, such as constant power operation for the in-line amplifiers (ILAs) used between the CLS and data center.

The key advantage for subsea network operators implementing an optical pass-through architecture at the CLS sites is a reduction in the total number of transponders needed for end-to-end DC-to-DC connections, and the corresponding cost savings, of up to two-thirds. This also yields savings in operational costs for power and space requirements at the CLS, where these are often at a premium. While the extra backhaul distance may marginally reduce the maximum data rate of each connection, this is partially offset by the performance of the latest generation of 140 GBaud coherent optics.

Thus a CLS-to-CLS connection at 800Gb/s per wavelength across the Atlantic will be able to achieve a speed of 700Gb/s when connecting end-to-end between data centers at each end of the subsea-terrestrial link; a marginal de-rating that yields significant cost and operational savings. This end-to-end optimization provides important benefits to AI and cloud providers that control both the subsea fiber cable and terrestrial fiber routes that inter-connect their data centers.

Leveraging L-band WDM spectrum in the terrestrial backhaul network

In many cases however, an AI and cloud provider seeking to connect to data centers across the world may not own the terrestrial backhaul fiber in other regions, for either economic or regulatory reasons. In many cases, the needed terrestrial backhaul connectivity from a CLS to an out-of-region data center is provided by a Managed Optical Fiber Network (MOFN) from a local CSP, or the subsea network operator may lease terrestrial dark fiber from a local operator. This backhaul capacity, either a MOFN or dark fiber, can sometimes come at high cost, directly proportional to the number of terrestrial fibers needed. For subsea fiber cables having 12, 24 or more fibers, leasing the equivalent number of fibers for the terrestrial backhaul can become cost-prohibitive.

An alternative end-to-end subsea-terrestrial architecture leverages the use of additional bands, such as C+L, in the terrestrial backhaul fiber spectrum, to enable more wavelength division multiplexing (WDM) channels (see figure 3) to be used across fewer backhaul fibers. While a subsea fiber typically provides 4.5THz of WDM spectrum in the C-band, use of C+L bands in the terrestrial backhaul network supports 9.8THz of WDM spectrum, or greater than twice the spectrum capacity of the subsea fiber.

Figure 3: Reducing the cost of subsea traffic backhaul by using more WDM spectrum in the terrestrial network using C+L.

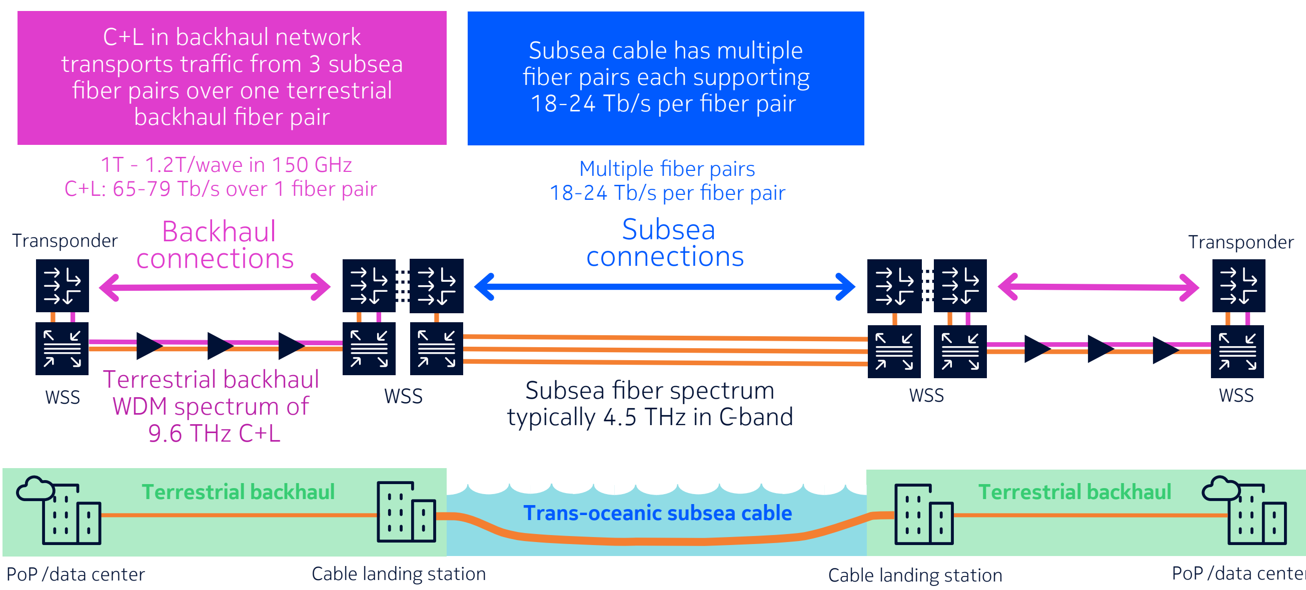

The fact that terrestrial backhaul networks will be of considerably shorter length compared to the subsea fiber also means that coherent wavelengths on the backhaul network can operate at higher data rates; adding an additional capacity boost on the backhaul network of approximately 25% in many cases, and even more for short backhaul links to a data center near a CLS. Thus leveraging C+L band WDM spectrum in the terrestrial backhaul network can reduce the number of backhaul fibers needed by 50-75% compared to the number of subsea fibers that are terminated at the CLS; providing significant overall savings for leased backhaul capacity.

Figure 4: Telecommunications providers operating both subsea networks and national or long-haul terrestrial networks can leverage L-band spectrum to maximize fiber usage.

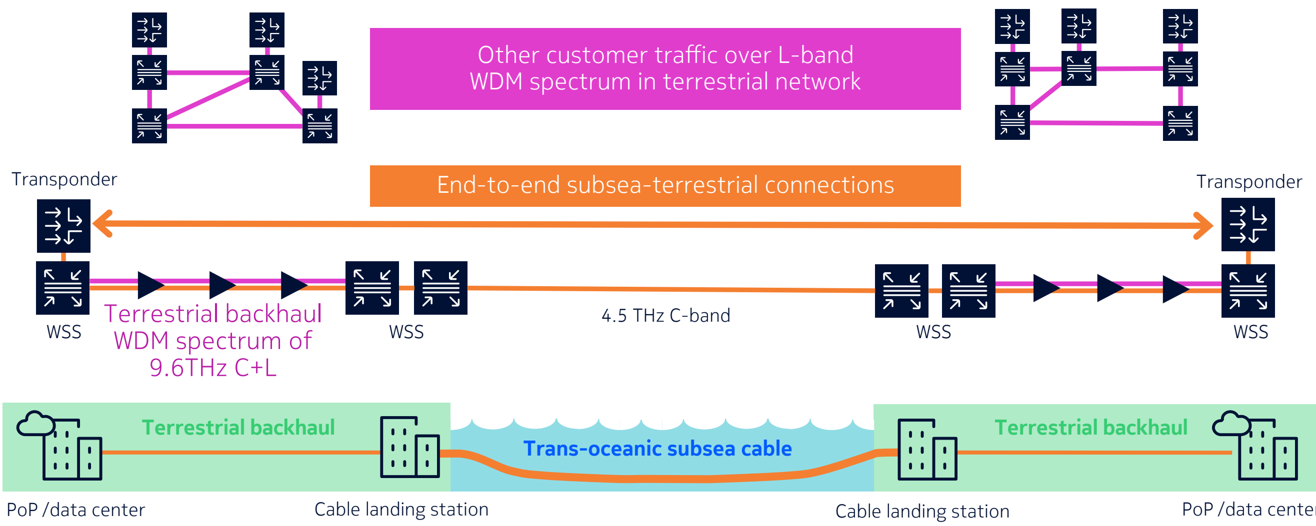

A third architecture approach for optimizing subsea-terrestrial networks combines the two approaches above (see Figure 4) and is well-suited for telecommunications providers that serve customers across their national or long-distance terrestrial networks, in addition to serving international traffic coming from subsea fibers. In that case they can leverage optical pass-through at the CLS in the C-band of the WDM spectrum for their subsea network connections transiting end-to-end between data centers or PoPs, while in parallel using L-band WDM spectrum to carry customer traffic from their national or long-distance network, thus making maximum use of their available terrestrial fiber capacity.

Network automation: A network multiplier

Operators can complement the end-to-end subsea-terrestrial network co-optimizations described above by using a common and unified transport solution across their subsea-terrestrial networks. This can in turn deliver more than lower operational costs and simplicity of common spares, training and management; it also unlocks the potential benefits of end-to-end network automation to plan, optimize, turn-up, monitor and protect the network.

Network automation can enable the rapid deployment of new terrestrial subsea capacity using zero-touch commissioning tools, automate channel optimization and service configuration, reducing the overall manual time and effort needed while also ensuring optimal performance and resource usage.

Network automation also enables real-time planning and provisioning using intent-based service requests for layer 0 and layer 1 connections, help search for new routes while ensuring optical impairment validation and contention resolution, performing optical performance validation, and automating service provisioning and turn-up.

Network automation can also coordinate and manage optical restoration in case of network failures by monitoring alarms and system/connection performance, performing fault isolation and impacted service analysis, and automatically searching for new routes with optical impairment validation, contention resolution and validating resource availability checks and assignment, helping operators to re-provision impacted services across the network.

Conclusion

Today's subsea network are inherently end-to-end; interconnecting data centers across the globe. New network architectures enabled by unified and optimized solutions enable subsea network operators to get the most capacity out of their end-to-end subsea and terrestrial backhaul networks, leading to savings in CapEx and operational costs, lower power, and reduced resource usage at valuable cable landing sites. The architecture options described in this paper cover a wide range of deployment options and associated use cases and benefits.

The well-used expression "it's a small world" takes on a new meaning in today's AI era and increasingly inter-connected world, where ubiquitous network connections bring people and services together across the globe, more than ever enabled by the fabric of end-to-end subsea-terrestrial networks.

Contact our sales team or visit our website to learn more about how Nokia can help you to evolve your subsea networks with a market-leading and proven portfolio of subsea-optimized optical networking solutions.

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs, which is celebrating 100 years of innovation.

With truly open architectures that seamlessly integrate into any ecosystem, our high-performance networks create new opportunities for monetization and scale. Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

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.

© 2025 Nokia

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

Document code: CID215254 (December)