

SOLUTION BRIEF

Datacenter Interconnection Supporting Cloud and AI Scale- Across Connectivity

Artificial Intelligence is Redefining the Datacenter

The growth of Artificial Intelligence is causing a new datacenter construction boom. AI GPU clusters use more electrical power than a single datacenter can support requiring them to be distributed across multiple, geographically separated facilities.

The amount of power these facilities now use requires that they be on separate power grids to handle the load. In fact, some of the larger datacenters—hundreds of megawatts to gigawatts of capacity—are so big that they have their own dedicated electrical generation facilities including oil, wind and even nuclear-powered plants (see Figure 1).




Figure 1 - Datacenter capacity growth; Source: Gartner, IDC, and Nvidia via McKinsey

As a result, datacenters are being built farther apart, driving a need for high-capacity connectivity over longer distances than ZR optics can effectively support. While the performance of pluggable coherent technology has improved markedly since its introduction in 2020 and 400ZR+ can now close many links it couldn't before, demand is now shifting to 800G, 1.6T and anticipating 3.2T to match the increasing rate of Ethernet client interfaces.

Because of their buying power and insatiable demand for bandwidth, hyperscalers will have the inside track for access to these next-generation technologies with the Tier 1 service providers left to compete for whatever inventory remains.

For the rest of the market, options are fewer. Beyond the hyperscalers are several other classes of datacenter operator ^[1]:

- Carrier-neutral colos with multiple service providers delivering connectivity services to enterprises and Cloud providers
- Edge datacenters supporting IoT, 5G, and autonomous systems for critical, real-time processing
- Neoscale (i.e., High Performance Computing or HPC) datacenters focusing on GPU-as-a-Service and AI inferencing
- Telecom datacenters supporting ISP backbone infrastructure with high network density, meet-me rooms and peering exchanges

[1] https://www.linkedin.com/posts/ramesh-babu-siddavatam_types-of-data-centers-understanding-activity-7431325354532806656-bkKd/?utm_source=share&utm_medium=member_desktop&rcm=ACoAAANXekBgJ5AVk70disCQdFm0yTrPy6jdaI

Facts & Figures:

- Global datacenter interconnect market revenue reached \$16.3 billion in 2025 and is expected to reach \$65 billion by 2035, a CAGR of 14.8%.^[2]
- Interconnection is growing at twice the rate of the internet and is 20x the size.^[3]
- Total interconnection bandwidth is expected to reach 33.6 Petabits per second (33,600 Terabits), 62% of which is used by service providers.^[4]
- As of 2024, 400G wavelengths make up the largest total terabits of bandwidth purchased, more than both 10G and 100G solutions.^[5]

The Challenge

Since its inception in the early 2020's, AI has largely been in the training stage, perfecting Large Language Models (LLMs) from which to optimize the algorithms that make AI useful. This activity takes place in the core of the AI network which can exist anywhere and is typically in remote locations where massive datacenters can be built more easily at lower cost and with cheaper power to drive the massive number of GPUs needed to make AI a reality.

Now that the LLMs are more 'intelligent', they are moving from the training stage to the inference stage ready for commercial use. This means the inference function needs to be closer to the end user for scalability, availability and latency reasons. This is driving the development of more edge and neoscale datacenters that are typically located close to and within metropolitan areas that can be tens, hundreds or even thousands of kilometers from the core site.

Unlike hyperscalers that typically build their own private fiber optical networks, edge, neoscalers and other types of datacenter providers depend on service providers for their connectivity needs. With the explosion of AI and Cloud-based services, demand for both lit services (i.e., wavelengths) as well as dark fiber across both the metro and long haul segments of the network is outpacing service providers' ability to keep up (see Figure 2).




Figure 2 - Extending the service provider network to the datacenter

^[2] Precedence Research
^[3] Equinix GXI report

^[4] Ibid
^[5] Zayo Bandwidth Report, 2025

Impact

All of this market evolution means that datacenter interconnection—and particularly scale-across—is becoming increasingly critical to the success and deployment of AI that relies on highly efficient data distribution between learning and inferencing sites.

Service providers must efficiently scale their networks to maximize capacity using denser transport platforms and higher rate channels.

Since the late 2010's, the growth of interconnection bandwidth has led to the development of a separate class of optical transport equipment optimized for datacenter environments called Compact Modular.

Competition for rack space and power with the datacenter is meaningful. Every rack unit of space and watt of power used for transport equipment is one fewer that can be used for revenue-generating compute, servers and storage. Compact Modular devices have several advantages over traditional WDM solutions for datacenter applications:

- Lower profile resulting in fewer RUs required to achieve equivalent capacity
- High port density and capacity—10s of Tbps per shelf
- Lower power in terms of watts per Gbps
- Datacenter form factor—fits in the same rack as servers and switches
- Front-to-back airflow to conform to hot aisle/cool aisle configurations

With its decided advantages in density and power, Compact Modular is also being increasingly adopted by service providers for all applications from metro access to subsea connectivity (see Figure 3).

Despite its advantages, Compact Modular platforms also run into certain barriers to entry:

- **Cost**—Compact Modular systems utilize the latest generation of coherent technology (e.g., 120Gbaud+) that lower the cost per bit transported but require additional CAPEX and OPEX to replace telco standard WDM systems.
- **Disruption**—In most cases, Compact Modular systems use different hardware, software and management than telco standard WDM systems even when they are from the same vendor. They require separate training and can result in significant changes to network operations and procedures.
- **Obsolescence**—Because of their architecture, Compact Modular systems are often specific to a generation of coherent technology. Once a new generation becomes available, the old one is immediately obsolete and all development stopped. To take advantage of the latest technology then requires a complete rip-and-replace to upgrade your network.




Figure 3 - Compact Modular adoption. Source: Signal AI Transport Hardware Report and Ekinops estimates

So, if you are not a hyperscaler or Tier 1 operator, what alternatives do you have to address your interconnection issues? Ekinops is here to help you answer that question.



Ekinops PTM300

Ekinops PTM300 provides the ideal mix of features, functionality and operability not found in other Compact Modular systems. Based on Ekinops' existing FlexRate™ technology, it provides single chassis capacity up to 32 Tbps in 3RU at 800G per channel with 1.6T channels coming in a future release.

Features

- Small 3RU form factor with up to 32 Tbps capacity—40 × 800G channels
- Space efficient—combine up to 32 800G channels with an amplifier and ROADM for a fully loaded, full C-band terminal in 3RU
- 'Evergreen' design allows technology migration within existing units by simply replacing modules—eliminates rip-and-replace
- 450mm depth—conforms to standard datacenter rack depths
- Ultra-low power consumption—0.06W per Gbps (incl. chassis, commons and optics)
- FlexRate™ programmable line interface—400G, 600G or 800G
- No license fee to upgrade bandwidth
- Reach over 2000 Km at 400Gbps to connect distant datacenters
- Automatic power balancing
- Easy to install and commission—lowers OPEX
- Lowest equipment cost per 100G
- Compatibility with the existing 360 telco standard platform line system allows PTM to be used at terminal sites and take advantage of existing infrastructure to minimize both CAPEX and OPEX

Simply by setting basic operating parameters such as OSNR tolerance and channel spacing, FlexRate sleds automatically select the optimal modulation format and baud rate to achieve the highest performance possible.

Management

Managing the PTM is simple and direct, in particular for basic point-to-point DCI and scale-across applications. In this case, no NMS is necessary, the equipment can be managed directly via standard interfaces based on OPENCONFIG and YANG data models.

For larger networks with hybrid services over more complex ring and mesh architectures, the PTM300 can also be managed by the same Celestis NMS network management system as the 360 telco standard platform. This means having a 'single pane of glass' interface for both

platforms and minimal training requirements for existing Ekinops customers.

With an open software architecture, the PTM300 is fully interoperable with all other standards-based network equipment and can be easily integrated into any third-party controller network environment.

Celestis NMS also provides a standard northbound interface (NBI) for seamless integration with your existing higher layer OSS/BSS platform.



Conclusion

The datacenter interconnection market is at a crossroads where more capacity and greater performance is required at a lower cost in less space and with lower power than ever before. New datacenter models outside of the hyperscalers are becoming more prevalent with differing needs and expectations about DCI and scale-across connectivity.

Service providers are becoming increasingly relied upon to provide that connectivity meaning substantial network upgrades are necessary to meet demand. This will require additional capital investment in their networks that must produce a deterministic ROI to make the business case work.

Ekinops PTM300 provides all the performance, scalability, manageability and 'evergreen' upgrade capability to justify the investment. High-density FlexRate coherent technology combined with complete flexibility of configuration minimizes the TCO while providing a future-proof platform that extends the economic value of the platform. Ekinops PTM300 is purpose-built to meet the needs of the DCI future.

About Ekinops

Ekinops is a European leader in optical networking and secured connectivity solutions for service providers, enterprises and public organizations around the world.

Ekinops designs, develops and manufactures high-performance solutions with cutting-edge software used to deliver high-bandwidth connectivity and managed services in over 70 countries worldwide.

Ekinops leads in the areas of:

- Advanced coherent optical networking for all metropolitan, regional, long-distance and datacenter interconnection applications
- Secured connectivity solutions for enterprises, from voice and data routing devices to Secure Access Service Edge (SASE) capabilities, including ZTNA (Zero-Trust Network Access), SD-WAN, and SSE.

Ekinops operates globally and is listed on Euronext Paris (FR0011466069 - EKI).

Contact us

sales.us@ekinops.com | www.ekinops.com