Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING
CUSTOM REPORTS

Network Automation for Cloud Providers: Interviews and Analysis

A custom Heavy Reading (now part of Omdia) report produced for Nokia

AUTHOR: STERLING PERRIN, SENIOR PRINCIPAL ANALYST, HEAVY READING (NOW PART OF OMDIA)

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY

Over the past several years, the communications industry has seen an increase in network automation, including in transport networks. The growing uptake of IP and optical convergence and artificial intelligence (AI) has further driven this trend. With its greater complexities—particularly in the optical layers—transport automation is not always easy to implement, but operators understand it is a crucial component of a successful automation strategy.

Since 2022, Heavy Reading (now part of Omdia) has been surveying communications service providers (CSPs) annually about their views and plans for automation in their transport networks. The annual study serves as a valuable barometer for the direction of network automation in the global telecom industry, including fixed, mobile, and converged network operators (*Open, Automated, and Programmable Transport Networks: 2024 Heavy Reading Survey*).

That survey, however, does not necessarily represent the requirements of cloud providers and carrier-neutral providers (CNPs). These are service providers that operate data centers, provide optical connectivity between those data centers (i.e., data center interconnect [DCI]), and use automation and AI/machine learning (ML) extensively within and between data centers.

With an expectation that cloud providers and CNPs have network automation interests and requirements that can be very different from telecom operators, Nokia partnered with Heavy Reading (now part of Omdia) to complete a custom, interview-based study of the cloud connectivity market. Specifically, we conducted six in-depth, one-to-one interviews. (See the ***Methodology*** section for more details.)

Key areas of investigation included the following:

- Leading use cases for network automation, both near term and longer term
- Primary benefits of network automation in IP networks and optical networks
- Views on IP and optical convergence through software-defined networking (SDN) and physical convergence (IP over DWDM)
- Performance and capacity requirements for coherent pluggable optics
- Optical layer automation opportunities and challenges
- Use cases to benefit the most from AI and ML
- Early use cases for generative AI in IP and optical networks

1.1 Key findings

The key findings of this report are as follows.

1.1.1 Drivers and use cases

- **Cloud providers interviewed by Heavy Reading (now part of Omdia) are realizing a wide range of benefits from network automation, and their high level views of benefits align closely with those of their CSP counterparts.** Interviewees are strongly focused on reducing human errors, simplifying operations, accelerating revenue, optimizing infrastructure usage, and improving customer experience/quality of service (CX/QoS).
- **Network as a service (NaaS), in which end customers configure services on demand via self-service portals, stands out as a means to meet customer needs, differentiate connectivity services, and generate new revenue.** NaaS can be applied across OSI Layers 1, 2, and 3, but most cloud providers interviewed for this project are largely focused on Layer 3 and Layer 2 NaaS, as this is where they see the strongest customer demand.
- **AI exchanges are another revenue opportunity for cloud providers that can be unlocked by investment in network automation.** AI exchanges are an evolution of internet exchange and cloud exchange models through which customers access AI resources (such as GPU processing) on demand. Cloud providers with existing internet exchange and cloud exchange products are particularly interested in AI exchanges.
- **In addition to top-line revenue/service automation, interviewees are strongly focused on "getting hands off the network" with automation.** Reasons cited include reducing human error, reducing time for configuration, upgrades, and provisioning elements, saving opex with reduced manual labor, increasing accuracy, and freeing employee time on rote activities to work on higher value tasks.

1.1.2 IP and optical convergence

- **Based on interviews, cloud providers behave much more like the hyperscalers (albeit with much lower volume requirements) and have been far more aggressive adopters of IPoDWDM architectures compared to CSPs.** While CSPs continue to plan how to adopt IPoDWDM architectures, the cloud providers are already doing it—and many have been doing so for years.
- **In terms of data rates, cloud providers are focused on migrating from 100G to 400G for the near to medium term and planning for an 800G future longer term.** These capacity requirements align more with CSPs than hyperscalers. That said, cloud providers are interested in adopting the new generation of 800G pluggable optics that can run 400G at ultra-long haul (ULH) and have the flexibility to run at 800G where needed.
- **Layer 3 disaggregation is a crucial pillar of IP and optical convergence for some cloud providers.** Two of the six companies interviewed have adopted IPoDWDM strategies based on disaggregated routers running on white box hardware. Disaggregated routing hardware and software vendors have come a long way in the past several years in building features and functions that are competitive with traditional routers. Examples include support for coherent pluggables optics up to the current generation of 800G technologies.

1.1.3 Optical layer automation

- **Although inter-data center capacity requirements can be large and are growing quickly, capacity alone is not a significant driver of optical layer automation, according to interviews.** There are three key factors that cloud providers weigh in automating their optical layer: degree of connectivity complexity, volume of changes to connectivity, and the need to support wavelength or spectrum services.
- **Optical layer automation is a lower priority for many cloud providers when compared to their CSP counterparts.** The reason is not that the cloud providers are slow-moving or behind, but rather, that many cloud providers simply do not see significant value in automating an optical layer that is based on point-to-point connectivity between data centers (i.e., DCI) and largely static. The lack of urgent need for optical layer automation was stated by half of the interviewees, each of which uses fixed point-to-point DCI for internal connectivity. Notable exceptions include the introduction of ROADMs (which add connectivity complexity and make cloud networks look more like CSP networks) and the delivery of wavelength services to enterprises and other service providers. In the latter case, NaaS is an important revenue-generating opportunity.
- **In addition to managing and automating optical wavelengths, two of the cloud providers interviewed expressed interest in automating the interconnection of fibers within the data center as a growing opportunity.** Variations of fiber interconnection have been around for decades, but the application is experiencing strong renewed interest driven by hyperscalers' fiber-dense AI factories. Some smaller cloud providers, like those interviewed, also see a market need for these functions.

1.1.4 AI and ML

- **Based on interviews, cloud providers are aware of the promise of AI/ML technologies in networking but are mostly in the education and evaluation stages of building their strategies across all network layers.** If hyperscalers are at one end of the spectrum as leading AI/ML technology adopters and CSPs are at the other end, we would place cloud providers closest to the CSPs along that spectrum.
- **Among cloud providers interviewed, the top transport use cases expected to benefit from AI/ML are network troubleshooting and root-cause analysis, network optimization and traffic engineering, predictive analytics and network health analysis, network performance monitoring, and capacity/ service provisioning/activation (when wavelength services are sold).** These AI/ML views are largely consistent with CSP views.
- **Challenges to AI/ML for networks include a lack of maturity in open AI products, a lack of open APIs among vendors (including the complex optical layer), and security and data sovereignty concerns.** All these concerns are also aligned with concerns of CSPs as they evaluate the application of AI/ML for networks.

1.2 Methodology

This analysis is based on data and findings from a set of six one-to-one interviews of data center operators, CNPs, and cloud connectivity providers conducted by Heavy Reading (now part of Omdia). Interviews were based on a common *Interview Guide* that was jointly developed by Heavy Reading (now part of Omdia) and Nokia. The six interviews were conducted from February through June 2025 by Sterling Perrin, a senior principal analyst at Heavy Reading (now part of Omdia).

We generically use the term "cloud providers" to describe these companies and to distinguish them from the CSPs, a.k.a. telcos, targeted in our previous network automation survey projects. "Cloud provider" is not an industry standard term and can be used to describe many types of companies. In this report, the cloud providers interviewed all own/operate data centers, connect to their own and to partner data centers with optical networks (i.e., DCI), sell space/power/cooling and connectivity to other cloud providers, and (in some cases) sell their own cloud services to enterprises.

Moreover, *cloud provider* and *CSP* are not mutually exclusive terms. Some CSPs are also cloud providers, and some cloud providers sell traditional CSP services. In this project, EXATEL is an example of an enterprise-focused CSP that also owns data centers and operates a DCI network. Note that hyperscalers can act as cloud providers, but they are specifically excluded from this study.

Five of the companies interviewed have agreed to have their names used in the report. One company asked to remain anonymous. The six service providers interviewed for the project are as follows:

- Telehouse America (US)
- Netrouting (Netherlands)
- DE-CIX (Germany)
- OpenColo (US)
- EXATEL SA (Poland)
- Tier 1 CNP (North America)

2. DRIVERS AND USE CASES

Cloud providers interviewed by Heavy Reading (now part of Omdia) expect—and are realizing—a wide range of benefits from network automation. Their high level views of benefits align closely with those of their CSP counterparts. Mapping discussion points to broad benefit categories, interviewees are strongly focused on reducing human errors, simplifying operations, accelerating revenue, optimizing infrastructure usage, and improving CX/QoS (see **Figure 1**). All these benefits are also important to CSPs.

Figure 1: Automation benefits and interview references

Automation benefits	Examples referenced
Reduced human error	• Low-touch and zero-touch operations • Operations consistency • Improved accuracy
Operational simplification	• Repeatable playbooks • Consistent service delivery
Faster service turn-up and revenue acceleration	• Layer 1, 2, and 3 NaaS/bandwidth on demand • AI exchange and GPUaaS • Faster service delivery • Automated service provisioning
Infrastructure usage optimization	• Network performance monitoring • Automated troubleshooting • Predictive analytics • Automated backup
Mitigate skilled labor shortage	• No direct references
Improved QoS and end-customer experience	• On-demand service • Predictive analytics • Network performance monitoring • Faster setup time
Enable open networking	• Rapid protection/restoration in multi-vendor networks

Source: Heavy Reading (now part of Omdia), 2025

Cloud providers are particularly focused on new revenue opportunities enabled by automation, including NaaS, and emerging AI exchanges, which include GPU as a service (GPUaaS). These emerging opportunities come from a combination of automation, data center assets in prime locations, multi-tenant occupancy of those data centers, and fiber-optic connectivity to the data centers.

2.1 NaaS

Mplify (formerly called MEF) defines NaaS as:

The combination of on-demand connectivity, application assurance, cybersecurity, and multi-cloud networking within a standards-based automated ecosystem. NaaS adopts a cloud-inspired operational model that delivers cloud-like flexibility, automation, and scalability to network services, while extending capabilities such as dynamic connectivity, enhanced application awareness, real-time performance visibility, advanced cybersecurity, and seamless integration across partner networks and multi-cloud, edge, and hybrid environments.

– Mplify 2025 NaaS Industry Blueprint, November 2024

Key NaaS service attributes include automation, programmability, flexibility, and visibility. While the early focus for NaaS was Ethernet and IP layer services, today the NaaS concept applies to many service types across many OSI layers and includes wavelength services and dark fibers, among others. NaaS also encompasses connectivity within data centers and campuses, between data centers, and across the WAN. Today's NaaS ecosystem is broad and includes enterprises, CSPs, hyperscalers, data center operators, and software as a service (SaaS) providers.

Cloud providers—including those interviewed for this report—believe they are uniquely positioned as the physical hubs in which all the NaaS ecosystem players locate and interact. Since they serve as crucial connectivity hubs, cloud providers have a prime opportunity to grow revenue as the NaaS market expands. And because automation is central to the NaaS definition, automation is also central to the success of cloud and data center operators.

The DE-CIX quote below explains the importance of NaaS and the role of automation.

First, automation enables the Network-as-a-Service (NaaS) business model. Transitioning from a traditional internet exchange business to a NaaS model means our customers expect a self-service portal where they can configure the services they need—whether peering services or cloud connectivity. They're accustomed to the seamless experience of cloud platforms like Azure and AWS, and they now expect similar functionality in networking. Automation is essential here because, without it, offering self-service provisioning in seconds wouldn't be possible.

– Thomas King, CTO, DE-CIX

While NaaS can be applied across multiple layers, including the optical layer, most cloud providers interviewed for this project are largely focused on Layer 3 and Layer 2 services for NaaS and not the optical layer. The latter is primarily used for internal DCI and is largely static.

The main exception among the interviewees is EXATEL. It is looking at optical layer NaaS specifically to differentiate optical wavelength services and sell them at higher prices than current wavelength services, which are somewhat commoditized in EXATEL's markets. On-demand connectivity is the key to this differentiation. The quote below explains EXATEL's motivation for optical layer NaaS.

It's all about new products and services and how we expose them to customers. The problem right now is the massive erosion in pricing—especially for lambdas. You can get a 100G or 200G Lambda at a very low cost. So, our goal is to add value on top of these basic offerings in order to maintain or even increase pricing. We're looking at value-added services in the optical domain, and automation is a big part of this product development strategy.

– Michal Szczęsny, Director of the Department of Architecture and Network Planning, EXATEL

It is important to keep in mind, however, that EXATEL offers the broadest range of services of all the companies interviewed for this project, including connectivity services sold to cloud providers as well as government and enterprises. Thus, EXATEL's optical NaaS strategy covers more than cloud connectivity applications alone.

2.2 AI exchanges

AI exchanges are new and are an evolution of the earlier models of internet exchanges and cloud exchanges. Whereas cloud exchanges build an ecosystem around public cloud providers, like Amazon, Google, and Microsoft, AI exchanges are built around high performance connectivity to AI resources, including AI services and infrastructure. Like cloud exchanges and internet exchanges, AI exchanges have a natural fit in data center hubs that house broad AI ecosystems that include AI infrastructure, producers, users, and network connectivity.

GPUaaS is an early subset of AI exchanges targeted specifically at on-demand access to GPU/TPU compute resources for both training and inference. Hyperscalers can afford to build their own AI factories housing 1,000s of GPUs, but most other industry participants do not have these resources and can benefit from renting GPUs on demand and paying for compute only when needed. These companies include enterprises, AI startups, and research institutions, among others.

Cloud and CNPs are well-suited to rent space and provide exchanges for these companies. As stated by the Telehouse America interviewee:

This middle market consists of small to mid-sized AI cloud providers that allow users to access GPUs on a time-based, pay-as-you-go model. Their services are much more affordable compared to major providers like Google …These AI cloud providers could become our customers. If we serve enterprise customers, they will need to connect to AI components for training and inference. However, they may not need continuous connectivity. They might only require access for a specific period—say, for one Monday—and then want to disconnect. That's why automation is crucial for data center providers like us.

– Akio Sugeno, Vice President of Internet Engineering, Operations, Business Development, Telehouse America

2.3 Other drivers and use cases

In addition to the top-line revenue/service automation drivers described above, interviewees for this project are strongly focused on "getting hands off the network" with automation. Reasons cited include reducing human error, reducing time for configuration, upgrades, and provisioning elements, saving opex with reduced manual labor, increasing accuracy, and freeing employee time on rote activities to work on higher value tasks. These motivations for reducing manual tasks are very consistent with views expressed in the 2024 CSP survey. The quote below from OpenColo encapsulates some of the thinking on automation and opex:

We're very budget conscious. I wouldn't say we're cheap, but we operate with strict financial limitations. I have just one developer on my team who helps build out the automation ideas I come up with, so we're not spending huge amounts of money on third-party solutions or a large development team. From an operational cost perspective, automation saves us a lot of money.

– Scott Brookshire, CTO, OpenColo

Lastly, interviewees did not spend much time talking about automation as an enabler for open networking and did not mention mitigating human labor directly as a motivation. Open interfaces (e.g., APIs) are essential for multi-vendor networks, and interviewees did discuss the importance of multi-vendor networks for NaaS, AI exchanges, GPUaaS, and network-wide restoration/recovery. Here, however, automation is not the enabler for open networking. Rather, open networks are an essential requirement for the automation use cases. Open interfaces and APIs are viewed as a crucial means of achieving cloud providers' automation goals.

Regarding labor, there is a relationship to automation, but not exactly one of mitigating a shortage of skilled labor. At least at this stage, the easiest wins in automation come from removing humans from routine, highly repeated tasks. Doing so frees up skilled employees to dedicate their time to more complex (and higher value) tasks. One interviewee expressed this function as a driver. In time, AI could change the relationship, but interviewees are just getting started with AI technology in their networks.

3. IP AND OPTICAL CONVERGENCE

Among the topics covered in this project, IP over DWDM (IPoDWDM)—defined as the physical integration of coherent pluggable optics on routers—is the one with the greatest divergence in views between CSPs and cloud providers interviewed in this project (again, excluding hyperscalers). Nokia uses the term coherent routing.

Hypercalers, including Google, Microsoft, and Amazon, pioneered 400ZR coherent pluggable optics for routers. They account for the vast majority of the 500,000-plus 400ZR units that have been shipped to date.

Although CSPs have shown strong and consistent interest in IPoDWDM in Heavy Reading (now part of Omdia) surveys, they have been slow to adopt for a couple of reasons, including requirements for higher performance ZR+ and ZR++ modules and network management complexity. The performance challenges have largely been addressed, but network management remains a significant challenge among CSPs.

Based on interviews, cloud providers behave much more like the hyperscalers (albeit with much lower volume requirements) and have been far more aggressive adopters of IPoDWDM compared to CSPs. While CSPs continue to plan how to adopt IPoDWDM architectures, the cloud providers are already doing it—and have been for years.

As the DE-CIX interviewee stated:

This is the direction the entire industry is moving toward. Managing optical transport directly from the router simplifies operations. It eliminates the need for separate DCI boxes, where traditionally you had to deal with colored optics and additional grey optics to interface with the router … With fewer hardware components, the system becomes less complex and easier to manage. This is why everyone is adopting pluggable optics, and we've been using them for years.

– Thomas King, CTO, DE-CIX

IPoDWDM was the general trend among interviewees, but there was one notable exception that found value in keeping the optical and router layers/equipment separate. Internet exchange (IX) provider Netrouting stated that interconnect customers do not want to depend on Netrouting Layer 2/3 equipment when purchasing optical interconnect services. Separation allows Netrouting to sell interconnect services at the optical layer, with customers adopting Layer 2/3 equipment of their choosing—whether it is Nokia or a different vendor.

The strategy is particularly interesting because Netrouting uses Nokia for both routing and optics. While a single-vendor architecture makes IP and optical layer network management much simpler, Netrouting chose not to converge the layers.

3.1 400G and beyond

The CSPs surveyed view 400G as the main currency for transport networks for the next several years, even as hyperscalers quickly move to 800G and race to standardize 1.6T coherent pluggable optics. Cloud providers surveyed for this project are much more aligned with the CSPs on this issue. They are focusing, for the most part, on migrating from 100G to 400G for the near to medium term and planning for an 800G future longer term.

As the Telehouse America interviewee explained:

From my perspective, 800G is still about five years away, based on customer feedback—especially for peering. Right now, 100G is still sufficient. Some customers need four or five 100G connections, and that is still sufficient. However, many are looking to move to 400G because cross-connects in New York are extremely expensive. If you need 400G, you require four cross-connects. But with 400G as a single connection, you only need one, which saves customers a significant amount of money.

– Akio Sugeno, Vice President of Internet Engineering, Operations, Business Development, Telehouse America

Still, the OpenColo interviewee even made a case for remaining at 100G in the near term for cost-efficient redundancy reasons. As he explained, a 400G link that goes down is a major capacity hit to a network, affecting traffic from multiple customers. Yet, if one of four 100G links goes down, capacity can be rerouted efficiently to keep services up and running. The cost per bit for 400G optics is lower than for 100G, but the costs of assuring reliability and redundancy for a 400G-based network are higher in OpenColo's case. As long as its customer demand focuses on 100G circuits, OpenColo prefers 100G connectivity.

In discussing data rates, it is also important to distinguish between the modules and the routers that host the modules. As described above, most cloud providers interviewed are interested in 100G and 400G modules today. However, future planning is driving a trend toward 800G-capable routers today, particularly for new builds or as part of the router upgrade cycle. The Netrouting interviewee explained the rationale for 800G-capable routers:

So, we bought the optical gear at 400G, knowing we could upgrade it later if needed. We purchased some IP equipment a little over a year later, and we just received that. It's already equipped with 800G ports. During the purchasing and delivery process, we actually issued a change order with Nokia to upgrade some of the IP gear to be ready for the next generation of optics.

– Savvas Bout, Founder and CEO, Netrouting

These new 800GE-capable routers are fully compatible with 400G coherent optics and can be equipped partially or fully with 400ZR/ZR+ modules on day one. There are additional options for 800G-capable routers. Service providers can run 2x400G breakouts to get two 400G connections from a single 800G router port, or they can configure the 800G optics using QPSK or 16QAM to achieve 400G ULH reaches.

3.2 IPoDWDM with disaggregation

In an interesting twist to the discussion, two of the six companies interviewed have IPoDWDM strategies based on disaggregated routers with white box hardware, including the large North American cloud provider and EXATEL.

Disaggregated routing hardware and software vendors have come a long way in the past several years in building features and functions that are competitive with traditional routers, including support for coherent pluggable optics. For the CNP interviewee, the key appeal is that disaggregated routers offer all the required routing functions at lower prices. The interviewee is especially impressed with IP Infusion, as noted in the quote below.

IP Infusion has probably made the most progress in our space, as we see it. Without naming names, I can say that they've dethroned some of the major players—whether it's our top competitors or other data center operators. They're replacing solutions from Arista, Juniper, and others who traditionally filled the switch/router role. IP Infusion does this using off-the-shelf hardware, their own OS and support, and pluggable optics.

– Senior executive, large North American CNP

While the large North American cloud provider is planning to migrate to white box router-based IPoDWDM for its next generation of transport, EXATEL is already using IP Infusion and DriveNets routers on white box hardware from UfiSpace and Edgecore. These routers are equipped with 400G OpenZR+ pluggable optics, with the largest white boxes having 36 x 400G ports. EXATEL has an internal development team that has built an SDN controller based on an open network operating system (ONOS) and plans to expand the functionality into orchestration.

4. OPTICAL LAYER AUTOMATION

Through the interviews, however, Heavy Reading (now part of Omdia) found that optical layer automation is a much lower priority for many cloud providers when compared to the CSPs. The reason is not that the cloud providers are slow-moving or behind in any way, but rather, that many cloud providers simply do not see a compelling driver to automate an optical layer that is based on point-to-point connectivity between data centers (i.e., DCI) and is largely static.

The lack of urgent need for optical layer automation was stated by three of the six interviewees. Their explanations were very consistent, as highlighted in the set of quotes below:

On the optical transport layer, where we don't sell optical services, our network is relatively static. Changes are typically related to upgrades, such as lighting up new dark fiber or adding more DWDM channels. As a result, automation in this area has been less of a priority.

– Thomas King, CTO, DE-CIX

Our setup is relatively simple—it's all point-to-point. We don't have to do much in the way of complex configurations.

– Scott Brookshire, CTO, OpenColo

We don't deploy those [optical layer] connections daily. You put that infrastructure in place, and while you manage and monitor it, there's not a lot of change happening there. We do implement customer circuits from time to time, but it's nothing compared to the volume of internet circuits or EVPN circuits we provide to our customers.

– Savvas Bout, Founder and CEO, Netrouting

Note that the key factors for automating the optical layer are as follows:

- Degree of connectivity complexity
- Volume of changes to connectivity
- Optical network services (e.g., wavelength services)

Capacity between data centers can be significant, and interviewees referenced the near-term need for 400G and planning for 800G, including purchasing 800G-capable routers for future capacity needs. But static capacity—even when large—does not necessarily drive optical layer automation.

The introduction of ROADMs transforms the network from a simple point-to-point interconnect to a more complex, mesh network where traffic can be dynamically routed. Added complexities in ROADM-based networks include provisioning various transponder technologies with different baud rates on the same add-drop structure, power management to ensure consistent performance across all channels (whether in point-to-point connections or multi-hop configurations), and path diversity configurations for data center interconnection resilience. Automation is an important tool to simplify all these important and complex ROADM network functions.

As one example, ROADM deployments are driving DE-CIX's initial adoption of optical layer automation. The DE-CIX interviewee sees automation as valuable for network restoration and incident management by rerouting traffic at the wavelength level.

Note that ROADM topologies also make cloud provider networks look much more like CSP networks, and CSPs are adopting optical layer automation to add flexibility and to manage complexity. Still, based on the interviews, it is clear that cloud providers will continue to make their heaviest automation investments at the IP layer for three primary reasons:

- As a routed layer, IP has significant connectivity complexity that can be efficiently managed through automation, using Ethernet virtual private networks (EVPNs) for the services layer and segment routing for the IP transport layer.
- IP connections are high volume and require frequent changes.
- IP is the layer at which most of their customer services are delivered.

As noted earlier, EXATEL's strong focus on optical layer automation is a notable exception among the service providers interviewed. The company is doing limited automation at the optical layer today, but a major goal over the next two years is to expose optical domain capabilities directly to its customers (i.e., optical layer NaaS).

The goal is to give customers flexible, on-demand access to optical services. Automation in the optical layer will not only support new products but also improve customer relationships and service efficiency.

– Michal Szczęsny, Director of the Department of Architecture and Network Planning, EXATEL

One key difference is that other data center operators interviewed for this project use DCI only for internal connectivity, while EXATEL sells optical wavelength connectivity to these types of service providers. Other cloud providers, along with enterprises, are target customers for EXATEL.

4.1 Intra-data center automation

In addition to managing and automating optical wavelengths, two of the cloud providers interviewed expressed interest in automating the interconnection of fibers within the data center as a growing opportunity. Variations of fiber interconnection have been around for decades, but the application is experiencing strong renewed interest driven by hyperscalers' fiber-dense AI factories. Some smaller cloud providers, like those interviewed, also see a market need for these functions.

In the past, optical circuit switches (OCS) were primarily based on MEMS, liquid crystals, or piezoelectric steering and found niche use in lab testing environments and as fiber patch panels. Today, a new crop of vendors has introduced innovative technologies, including robotics, to map fiber input and output ports on demand. This OCS renaissance was sparked by hyperscalers, and by Google in particular, but it is spreading.

The large CNP interviewee sees revenue opportunities in "physical NaaS," in which data center cross-connects are set up for customers through automation, on demand. For this provider, the physical NaaS layer will become an extension of its end-to-end NaaS strategy. This strategy includes Layer 3 and Layer 2 services, as well as a move into Layer 1 optical services. The Telehouse interviewee also referenced potential applications for fiber automation, though with some reservations.

Multi-layer NaaS combined with physical NaaS enables end-to-end on-demand services, but it also requires industry standardization. The large CNP owns its data center and campus DCI infrastructure but relies on other service providers for WAN connectivity. WAN providers must support on-demand networking and also share standard APIs for provider-to-provider interworking.

For physical NaaS specifically, another challenge for cloud providers is vendor and technology maturity. The Telehouse interviewee was concerned by a new OCS market dominated by untested startups and technologies (e.g., robotic fiber switching), as noted in the quote below.

My concern is that I don't know how long these [fiber automation] companies will last. If they disappear in five or ten years, we're left with unsupported technology, and that's a major risk. If there were multiple stable vendors—say, Company A, B, and C—then I'd feel more comfortable. That way, if one disappears, we have alternatives.

– Akio Sugeno, Vice President of Internet Engineering, Operations, Business Development, Telehouse America

5. AI AND ML

Cloud providers interviewed for this report and CSPs surveyed by Heavy Reading (now part of Omdia) are largely aligned around their AI/ML strategies in terms of the following:

- Their current levels of AI adoption for networking
- Transport use cases where AI and ML are expected to have the biggest impacts

Based on the interviews, cloud providers are aware of the promise of AI/ML technologies in networking but are mostly in the education and evaluation stages of building their strategies. If hyperscalers are at one end of the spectrum as leading AI/ML technology adopters and CSPs are at the other end, Heavy Reading (now part of Omdia) would place cloud providers closest to the CSPs along that spectrum. The quote below captures the early-stage nature of AI/ML adoption among cloud providers:

I suspect it will be over time. But I don't think we're quite ready for it yet. A lot of our infrastructure and services—and pretty much everything we do—is still a work in progress … To get back to your question: I don't think we've matured enough operationally to start looking at tools like AI. We need a better understanding of our own environment before we start handing off responsibilities to something we have less control over.

– Savvas Bout, Founder and CEO, Netrouting

While not widely adopted, interviewees recognize the potential for AI/ML for their networks moving forward, including the application of AI/ML technologies together with network automation. In the 2024 Heavy Reading (now part of Omdia) CSP survey, respondents identified the top transport use cases to benefit the most from AI/ML as follows:

- Predictive analytics (selected by 47% of respondents)
- Network troubleshooting (selected by 44%)
- Energy consumption optimization (selected by 41%)
- Network optimization and traffic engineering (selected by 40%)

In this interview project, we asked cloud interviewees to identify their targeted use cases for AI/ML. Responses were open-ended, but we were able to map interviewee responses and descriptions to the broad use case categories defined in the CSP survey.

Figure 2 lists transport use cases and how these use cases were referenced and described in the in-depth interviews.

Figure 2: Transport use cases to benefit most from AI/ML

Transport network use case categories	Examples referenced in interviews
Predictive analytics and network health analysis	Failure prediction
Network troubleshooting and root-cause analysis	Incident diagnosis, root-cause analysis, mitigation strategies, log analysis, proactive problem resolution
Network optimization and traffic engineering	Network resilience, resource optimization, operational efficiency
Network performance monitoring	Quality control
Capacity/service provisioning/activation	Predictive network capacity requirements
Network design/planning	Resource-focused ML to predict future investment requirements

Source: Heavy Reading (now part of Omdia), 2025

Based on this mapping exercise, top transport use cases expected to benefit from AI/ML are

- Network troubleshooting and root-cause analysis
- Network optimization and traffic engineering
- Predictive analytics and network health analysis
- Network performance monitoring
- Capacity/service provisioning/activation (when wavelength services are sold)

As shown in **Figure 2**, network troubleshooting and root-cause analysis and network optimization and traffic engineering garnered the most references. These were followed by predictive analytics and network health analysis and network performance monitoring.

These responses are largely in line with CSP views, and they make sense as AI/ML focus areas. Each of these use cases benefits from AI's ability to rapidly analyze massive amounts of data and make decisions and predictions based on the data.

An area of possible divergence is AI/ML for energy consumption optimization, which was listed as highly important in the CSP survey but not directly referenced during in-depth interviews. Energy consumption is a critical issue for data center operators, and companies could benefit greatly from applying AI/ML technologies.

It is possible that Heavy Reading (now part of Omdia) simply did not speak with enough companies in this project for energy consumption optimization to rise to the surface as an AI/ML use case. Yet, several interviewees mentioned high energy consumption in data centers as a significant issue.

5.1 AI adoption challenges

As they plan for using AI/ML tools in the future, openness and standardization are of high importance to the cloud providers interviewed. In fact, a lack of openness among vendors and a lack of maturity in open AI products were both cited as inhibitors to adoption for AI/ML technologies in the network.

One view expressed is that single-vendor AI functionality is not useful in multi-vendor networks because the aim is to apply AI-enhanced functions—such as network monitoring and predictive analytics—to the full network, not just to individual elements. Multi-vendor interoperability requires standardization. The OpenColo interviewee summarized the role of open AI as follows:

The goal should be to stick to open standards and make things as generic as possible. One system to control it all—that would be amazing. We already have companies like Google with OpenConfig. That's a great start, but we need smart people to make it easier for network engineers to understand how to implement these things … A Nokia-only or Cisco-only solution isn't enough. It needs to work across an open network, aggregating data from multiple vendors in a meaningful way.

– Scott Brookshire, CTO, OpenColo

As the industry moves toward open AI tools, there is also concern about maturity, as summarized by the DE-CIX interviewee:

Open AI tools are still in the early stages, making it difficult to determine which vendors to trust. We have to carefully decide which tool to invest in and integrate into our stack, as that commitment comes with significant costs and long-term implications.

– Thomas King, CTO, DE-CIX

A third concern highlighted specifically in the EXATEL interview is the security of data used in AI models. EXATEL's security team wants everything to be built and hosted internally, with applications to be built by EXATEL and data physically hosted within EXATEL data centers. But the DIY model adds tremendous costs to both software development and EXATEL-owned GPUs. For these reasons, EXATEL has been conducting a slowly progressing feasibility study on domain-specific AI for nearly two years. As the interviewee states:

We're trying to determine what type of model we should use. If we host everything locally, it will be very expensive. But using external platforms requires sharing sensitive data. There's a trade-off between control and cost.

– Michal Szczęsny, Director of the Department of Architecture and Network Planning, EXATEL

A final note: EXATEL's security requirements are at the extreme end of the spectrum because the operator is government-owned and provides security services. However, concerns around data sharing, security, and data sovereignty are frequent topics in other discussions with operators about telco-specific AI models. AI domain data model concerns are not exclusive to transport, but AI transport use cases are affected by these decisions.