RCR Wireless News
NEWS. INSIGHT. IMPACT.

RCR Tech
NEWS. INSIGHT. IMPACT.

Scaling Optical Networks for the Hyperscale and AI Era

By Juan Pedro Tomas, Editor, RCRTech

CONTENTS

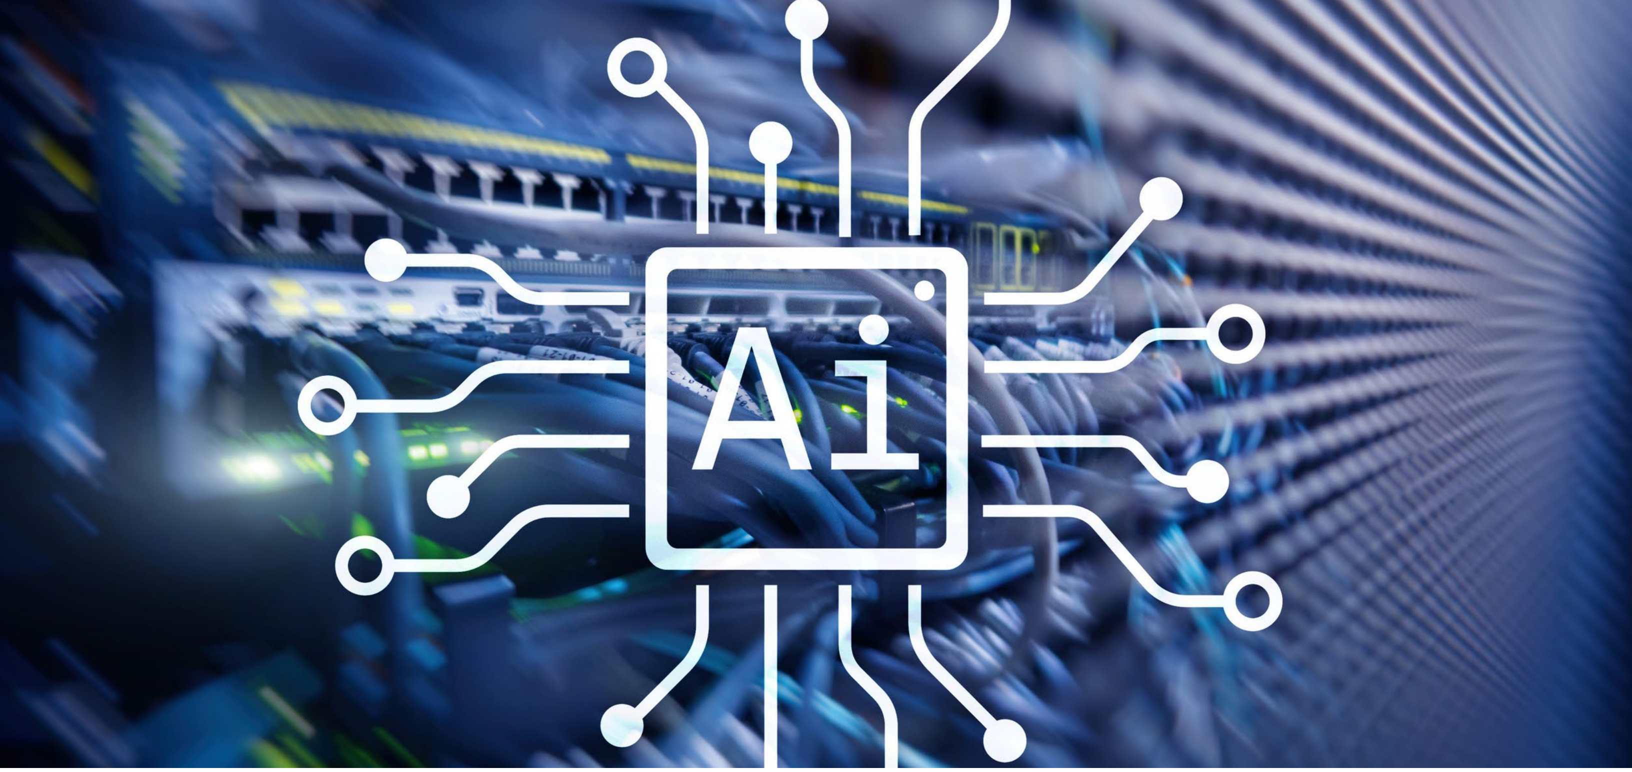

INTRODUCTION

The rapid expansion of artificial intelligence (AI) infrastructure is reshaping how optical and fiber networks are designed, deployed, and operated. While cloud computing has driven steady bandwidth growth over the past decade, AI workloads are introducing a fundamentally different traffic profile—one characterized by massive east-west data flows, tightly synchronized compute clusters, and unprecedented requirements for low latency and deterministic performance.

As hyperscalers, cloud providers, and enterprises scale AI training and inference deployments, optical transport infrastructure is becoming a strategic foundation for the next generation of digital infrastructure. The growth of distributed AI architectures, coupled with power constraints at traditional data center hubs, is driving a surge in demand for new fiber routes, coherent optics, and high-capacity data center interconnect (DCI) networks.

The implications extend across the optical ecosystem. Metro, long-haul, and subsea networks are all facing increased pressure as operators prepare for traffic growth measured not in incremental percentages, but in multiples. At the same time, the industry is approaching physical limits in spectral efficiency, forcing a broader shift toward multi-fiber architectures, new optical line systems, and increasingly automated operations.

AI is also changing how operators think about networks. Historically, network infrastructure was often viewed as a supporting layer that scaled alongside compute. Increasingly, however, networking is becoming a central determinant of AI performance. Large language model training, distributed inference, and emerging agentic AI architectures depend on high-capacity, highly synchronized connectivity between accelerators, storage platforms, and geographically distributed facilities.

This report examines how AI-driven demand is changing traffic patterns and network requirements, the role of coherent optics and pluggable technologies in scaling capacity, the growing importance of DCI and scale-across architectures, and the physical, operational, and economic challenges associated with deploying AI-era optical infrastructure.

AI WORKLOADS ARE RESHAPING NETWORK TRAFFIC

One of the clearest themes emerging from industry interviews is that AI traffic differs fundamentally from traditional cloud workloads.

Stephen Callaghan, senior principal technologist, network-core at AWS, explained the contrast: "Traditional cloud traffic is distributed and bursty. It's made up of millions of independent customers, each generating modest, uncorrelated demand. AI workloads are the opposite: a single AI training cluster could represent a single customer, involve over 100,000 accelerators communicating simultaneously, generating petabits of tightly coupled traffic across a handful of datacenters."

Bruno Berti, SVP of global product management at NTT Global Data Centers, described a similar transformation. "Training and large-scale inference generate enormous volumes of traffic inside the data center—GPUs talking to GPUs—with far higher bandwidth and much tighter latency tolerances. In some cases, microseconds matter."

According to Berti, these characteristics require networks that are "flatter, more deterministic, and far more resilient," combining higher-capacity fabrics with advanced optical interconnects and tighter integration between compute, storage, and networking.

Mike Jack, director of product marketing at VIAVI, emphasized that AI traffic is not only larger in scale but also fundamentally different in behavior.

"AI workloads are fundamentally transforming network behavior by introducing traffic patterns that are both more intense and more synchronized than traditional cloud environments. While conventional applications generate short, bursty flows tied to user transactions, AI training produces sustained, high-bandwidth 'all-to-all' communication between thousands of accelerators."

He added that these synchronized exchanges create strict operational requirements because "any delay or packet loss can block an entire training cycle, making low latency, lossless performance, and tight synchronization critical requirements."

The implications extend beyond bandwidth alone. Noah Taylor, head of market intelligence and growth strategy for broadband at AFL, noted that AI is changing the balance of network traffic itself.

"AI is flipping the traffic model from the traditional heavy downstream pattern to a much more symmetric—and often

"Traditional cloud traffic is distributed and bursty... AI workloads are the opposite: a single AI training cluster could represent a single customer, involve over 100,000 accelerators communicating simultaneously, generating petabits of tightly coupled traffic across a handful of datacenters."

— Stephen Callaghan, Senior Principal Technologist, Network-Core at AWS upload-heavy—profile. Inferencing at scale, especially real-time IoT use cases, is driving massive upstream traffic as data flows back to models for instant decisions. Bandwidth alone is no longer enough; ultra-low latency has become equally critical," Taylor said.

This shift is creating unprecedented demands on both intra-data-center and inter-data-center connectivity.

Rob Shore, head of optical networks solution marketing at Nokia, observed that AI demand is driving: "Significantly more bandwidth between locations (data centers, regions, and clouds), often exceeding what a single fiber pair can support."

Phil Wong, technology principal at KPMG US, believes the next wave of network demand will increasingly come from inference.

"As we move into the Agentic AI era, traffic coming from inference workload will drive demand for high speed, low latency connectivity."

Wong added that inference traffic could migrate toward the network edge as physical AI applications become more widespread.

Brodie Gage, senior vice president, global products and supply chain at Ciena, argued that AI is forcing the industry to rethink network design entirely.

"AI is not just another enterprise application riding on top of the network—it requires new architectures to address changing traffic patterns and new solutions to handle changing network requirements," Gage said.

According to Gage, AI-driven workloads require "an order of magnitude of higher capacity than traditional metro DCI networks."

He also pointed to multimodal AI as an important accelerator of bandwidth growth. "Text-based AI had less of

"AI is flipping the traffic model from the traditional heavy downstream pattern to a much more symmetric–and often upload-heavy–profile. Inferencing at scale, especially real-time IoT use cases, is driving massive upstream traffic as data flows back to models for instant decisions. Bandwidth alone is no longer enough; ultra-low latency has become equally critical."

— Noah Taylor, Head of Market Intelligence and Growth Strategy for Broadband, AFL

"AI is not just another enterprise application riding on top of the network–it requires new architectures to address changing traffic patterns and new solutions to handle changing network requirements."

—Brodie Gage, Senior Vice President, Global Products and Supply Chain at Ciena an impact; multimodal AI is what is changing everything."

Historically, optical innovation concentrated on extracting more capacity from individual fiber pairs. Shore believes priorities are shifting toward architectures that can scale horizontally.

"Innovation priorities [are moving] away from purely spectral efficiency toward solutions that help operators and AI/ cloud providers scale capacity across multiple fibers with lower cost, power, and space," Shore said.

As power availability increasingly dictates where AI campuses are built, geography is changing as well.

According to Shore: "New facilities are now being deployed primarily where power is accessible versus where fiber already exists or near population centers."

Taken together, these trends suggest that AI is not simply increasing traffic volumes. It is changing the fundamental assumptions underlying network architecture, synchronization, and physical infrastructure planning.

"Innovation priorities [are moving] away from purely spectral efficiency toward solutions that help operators and AI/ cloud providers scale capacity across multiple fibers with lower cost, power, and space."

— Rob Shore, Head of Optical Networks Solution Marketing at Nokia

"As we move into the Agentic AI era, traffic coming from inference workload will drive demand for high speed, low latency connectivity."

— Phil Wong, Technology Principal at KPMG US


NOKIA
Lighting the way for AI infrastructure
Scale faster, operate more reliably and deploy with confidence
Built for the demands of AI, our innovative optical networking solutions deliver performance, efficiency and agility.
From coherent pluggable optics and high-performance embedded optical engines to next-generation optical line systems and automated operations across network domains, we help you maximize capacity, improve power efficiency and lower TCO.
Learn more at
nokia.com/optical-networks

DATA CENTER INTERCONNECT AND THE RISE OF SCALE-ACROSS

Perhaps the most significant architectural consequence of AI is the emergence of scale-across networking.

Historically, GPU clusters were largely confined to individual facilities. Today, power-density limits and campus constraints are pushing deployments across multiple buildings and, increasingly, across metropolitan regions.

Shore explained: "The largest GPU cluster now exceeds 500,000 GPUs and the industry is headed toward a one million GPU cluster size perhaps by the end of 2026."

He added: "About 12 months ago a new term—scale-across—was introduced as a form of DCI for backend networks to support connecting GPU/XPUs across multiple physical data centers at 100s or even 1,000s of kilometers."

Jimmy Yu, vice president of optical transport market research at Dell'Oro Group, described the practical implementation of these architectures.

"The current architecture being deployed for scale-across DCI is IP-over-DWDM (IPoDWDM), where large quantities of 800 Gbps ZR+ pluggable optics are installed on a data center switch/router and multiplexed into hundreds of fiber pairs between AI data centers."

Berti said AI is transforming expectations for inter-data-center connectivity.

"We're now seeing much higher sustained bandwidth, tighter latency tolerances, and far greater sensitivity to consistency and jitter—especially when AI training, inference, and data pipelines are distributed across multiple facilities."

He concluded: "The net result is that interconnect is no longer just a connectivity layer—it's becoming a core enabler of AI performance."

Gage similarly described inter-data-center networking as becoming "an ultra-scalable, lossless, deterministic extension of the intra-data-center fabric."

AWS has adapted its own architecture to reflect this evolution. According to Callaghan: "Because we're connecting buildings, not just stamping out isolated ML clusters linked only by the front-end network, our DCI carries significantly more traffic."

He added: "The scale of individual DCI requirements today would have been unthinkable a few years ago."

The physical infrastructure supporting these deployments is expanding rapidly as well. Taylor noted that hyperscalers increasingly require multiple redundant paths between facilities.

"The current architecture being deployed for scale-across DCI is IP-over-DWDM (IPoDWDM), where large quantities of 800 Gbps ZR+ pluggable optics are installed on a data center switch/router and multiplexed into hundreds of fiber pairs between AI data centers."

— Jimmy Yu, Vice President of Optical Transport Market Research at Dell'Oro Group

"Our hyperscaler partners are now routinely asking for multiple fully diverse, redundant interconnects per data center. That drives major changes in the outside plant: more fiber counts, diverse routing, and far higher expectations around build quality."

He added that traditional hub-and-spoke topologies are giving way to mesh-oriented designs. "What used to be a hub-and-spoke model is quickly turning into more of a mesh network architecture design for outside plant," he said.

This surge reflects both dense east-west traffic within facilities and hyperscalers' growing insistence on resilient, geographically diverse connectivity between campuses.

"The net result is that interconnect is no longer just a connectivity layer—it's becoming a core enabler of AI performance."

—Bruno Berti, SVP of Global Product Management at NTT Global Data Centers

More Fiber. Tighter Timelines. Longer Distances. One Answer.

AFL solutions for campus, metro and long-haul DCI. Built for the AI era.

Every AI cluster multiplies the demand on the interconnects around it - campus, metro and long-haul alike. Campus links demand more fiber than ever in restricted duct space. Metro rings need to deploy at speed through congested underground routes. And long-haul links must deliver capacity across vast distances, where every splice point and every dB of loss counts.

AFL answers all three. Ultra-high-density cabling packs maximum fiber into every campus pathway. Deployment-ready cable and splicing solutions accelerate metro turn-up. And reel-length, reduced-OD designs extend long-haul reach with fewer splice points. All from one integrated ecosystem of cable, connectivity, splicing and test from a single partner.

The result: networks that scale with AI demand and deliver a faster path to revenue, wherever your data centers sit.

Planning your AI network strategy?
Download the white paper:
Architecting AI at Scale
>

Ready to see the solutions?
Explore our DCI Solutions
>

Visit our website www.AFLglobal.com or call 864.433.0333

COHERENT OPTICS AND CAPACITY SCALING

As AI infrastructure scales, coherent optical technologies are becoming central to network expansion strategies. The industry is increasingly deploying 400G and 800G interfaces while preparing for 1.6T technologies to maximize the capacity of existing fiber infrastructure. At the same time, interviewees consistently emphasized that faster optics alone will not eliminate the need for substantial investment in additional fiber.

Shore of Nokia noted that optical innovation has already significantly reduced "the cost, space, and power required for data center connectivity," but warned that the industry is approaching physical limits.

"Due to the laws of physics, we are approaching the theoretical limit of how much information can be carried by a single pair of fibers."

Yu of Dell'Oro Group reached a similar conclusion. "The current optical technologies, such as 1.6 Tbps-capable wavelengths and the use of both C-band and L-band, help reduce the need for new fiber deployments by improving spectral efficiencies. However, as we have nearly reached Shannon's Limit, the benefits are diminishing."

According to Yu, this means AI-driven infrastructure growth will require a combination of higher-capacity optics and new physical fiber construction rather than relying on spectral-efficiency gains alone.

Berti of NTT Global Data Centers echoed that assessment. "Higher-speed optics significantly extend the life of existing fiber. But as AI traffic scales, optics alone won't eliminate the need for new fiber—the winning model is smarter optics plus thoughtful, targeted fiber expansion."

AWS has already invested heavily in coherent technologies. Callaghan explained: "We've used coherent optical technology for many years to maximize capacity per fiber pair, and we were an early large-scale adopter of ZR pluggable optics."

He added: "Where the industry is now scrambling to move from 400G to 800G coherent, we were already there."

Gage of Ciena expects coherent pluggables to become increasingly important as AI deployments expand. "Coherent optics, including coherent pluggables, will play an increasingly critical role in future AI-driven network architectures," he said.

At the same time, he believes traditional coherent transponders will remain important where operators need to maximize spectral efficiency over constrained routes.

As deployment scenarios become more diverse, operators are optimizing for different constraints rather than converging on a single architecture. Helen Xenos, senior director of portfolio marketing at Ciena, said: "Different constraints — whether spectral efficiency, power efficiency, density, thermal scaling, operational simplicity, or deployment velocity — are driving different optimal implementations."

Taylor of AFL argued that the challenge is increasingly shifting from wavelengths to physical fiber density. "400G, 800G, and emerging 1.6T transceivers are buying us valuable time and spectrum efficiency, but they cannot fully solve the problem alone. The raw fiber count shortage is the limiting factor."

He noted that this reality is driving interest in high-density cable designs and future technologies such as multi-core fiber and hollow-core fiber, which could increase capacity without requiring proportional increases in conduit space.

FIBER EXPANSION AND PHYSICAL CONSTRAINTS

Although coherent optics continue to improve transport efficiency, nearly every participant agreed that AI growth will require significant fiber expansion. Yet many interviewees argued that the biggest near-term bottleneck is not fiber but power availability.

Nokia's Shore stated plainly: "The primary physical constraint for AI buildouts today is access to power."

Because AI campuses require extraordinary electrical capacity, operators increasingly select locations based on available power rather than existing telecommunications infrastructure, creating demand for entirely new long-haul fiber routes.

Shore also highlighted the scale of modern deployments. "Connections are now requiring thousands of fiber strands between locations."

Wong of KPMG similarly observed: "Currently, access to power, on or off-grid, is the biggest challenge, followed by supply chain delays and availability of labor."

Berti likewise concluded: "If you step back, the biggest physical constraint today isn't fiber—it's power, followed closely by deployment timelines."

He explained that grid interconnections, substations, transformers, and permitting frequently struggle to keep pace with hyperscale AI demand. "Speed-to-power has become just as important as speed-to-market," he said.

Within AI campuses, interconnect density is also changing dramatically. "AI training environments require massive, continuous east-west communication between GPUs. That drives dramatically higher fiber density inside halls and campuses, flatter network layouts, and far more attention to signal integrity and routing," the executive added.

From AFL's perspective, the physical-layer implications are equally significant. Taylor warned that deployment schedules are increasingly constrained by fiber availability.

"Network infrastructure is becoming a limiting factor in the short term. Ribbon fiber lead times have stretched beyond 60 weeks in many markets, directly impacting AI build schedules."

He added that hyperscalers are demanding faster deployment methods because "their revenue depends on rapid scaling."

Taylor also emphasized the growing importance of cable density. "The increase is dramatic—often 5–10× or more, and in some large AI campuses it's approaching an order of magnitude higher."

According to AFL, this surge in demand is driving interest in higher-density cable designs, modular deployment approaches, and new fiber technologies capable of maximizing existing conduit capacity.

AWS has attempted to reduce deployment timelines through tighter integration between logical and physical planning. Callaghan explained: "Fiber and power are available, but the constraint is when they come online."

He noted that AWS reduced deployment time substantially by linking structured cabling directly to network design. "At our first machine learning site using this approach, we cut network build time by 48%—from 16 weeks to just over 8."

Gage argued that urgency itself has become a defining characteristic of AI infrastructure deployment. "The industry is moving to turn up networks faster than we have ever seen before because AI factories cannot deliver value until the connectivity is fully operational," he said.

The urgency to deploy AI infrastructure rapidly is also influencing transport architectures themselves. As Xenos observed: "The urgency to rapidly turn up massive amounts of network capacity is becoming a defining requirement for many AI infrastructure deployments."

METRO, LONG-HAUL, AND SUBSEA NETWORKS UNDER PRESSURE

AI demand is affecting every segment of the optical ecosystem, although not equally.

According to Nokia's Shore: "All areas of the network—metro, long-haul, and subsea—are all experiencing pressure for more capacity, more fibers, and faster service delivery."

Callaghan noted that AWS has already observed dramatic changes in metro planning assumptions. "The headroom we'd planned for 18 months may be consumed in six."

He added that AWS's newest backbone generation expanded rapidly. "Our newest backbone generation grew from zero to 3× the capacity of our legacy network in under a year."

Ciena's Gage identified metro DCI as one of the most immediate pressure points. "The most immediate pressure building is in the metro, particularly for DCI, and in AI multi-cloud access networks," he said.

Taylor offered a somewhat different perspective, emphasizing long-haul infrastructure. "Right now, the greatest pressure is on long-haul routes that connect the exploding number of AI data centers."

He expects metro networks to experience greater stress as inference workloads become more distributed.

"Metro networks will feel the next wave as AI shifts from training to widespread inferencing and edge use cases," Taylor added.

Meanwhile, subsea deployment remains constrained primarily by construction timelines. As Callaghan observed: "Subsea is a timeline challenge: you can't compress the physics of laying ocean cable."

TESTING, VALIDATION, AND NETWORK OBSERVABILITY

As networks become integral components of AI systems, validating their performance is becoming as important as deploying bandwidth.

Jack of VIAVI argued that traditional testing methodologies are no longer sufficient. "Testing is shifting from a traditional, component-level approach to a holistic, system-level model that reflects the behavior of real AI workloads."

Rather than validating isolated devices, operators increasingly need assurance that complete systems perform correctly under synchronized AI traffic patterns, congestion events, and failure scenarios.

Jack explained that validation is moving toward ensuring "behavior, trust, and resilience at scale," rather than simple conformance to specifications.

The industry's migration to 800G and 1.6T interfaces compounds these challenges.

"The move to 800G and 1.6T introduces significant challenges at both the physical and system levels," the VIAVI executive said.

According to Jack, higher speeds increase susceptibility to signal degradation while also making interoperability across multi-vendor ecosystems increasingly important.

He also stressed that conventional throughput measurements alone are insufficient.

"AI workloads require validation of synchronized performance, low latency, and reliability across the entire fabric."

Observability requirements are evolving in parallel. Jack noted: "AI workloads require a step change in observability, moving from coarse, aggregated monitoring to fine-grained, real-time visibility into network behavior."

Operators increasingly need telemetry capable of correlating network conditions directly with GPU utilization, job completion times, and application performance. "This shift is driving the adoption of advanced telemetry and analytics that can correlate network conditions directly with application outcomes," Jack said.

Traditional assurance methodologies are also evolving. "Traditional testing approaches are no longer sufficient for AI-scale environments," the executive added. Instead, operators are adopting "AI workload emulation, digital twins, and continuous validation strategies." These approaches allow networks to be tested under realistic conditions without the need to deploy large physical clusters, reducing cost and complexity while improving accuracy.

Power and thermal behavior are becoming testing priorities as well. "Testing must extend beyond performance metrics to include power consumption, thermal behavior, and system stability under real operating conditions," Jack added.

"Testing is shifting from a traditional, component-level approach to a holistic, system-level model that reflects the behavior of real AI workloads."

— Mike Jack, Director of Product Marketing at VIAVI

As AI infrastructure scales toward terabit optics and increasingly dense deployments, validation is expanding from bandwidth verification into comprehensive systems assurance.

Powering the future of high-capacity DCI

With the rapid expansion of cloud and AI applications, the data center interconnect (DCI) landscape is facing growing demands for scalability, efficiency, and performance.

As the global leader in high-speed connectivity, Ciena delivers robust DCI solutions with the optical systems, automation software, and services to scale networks with ease. These solutions meet diverse connectivity use cases from metro links to transoceanic routes.

Metro DCI: Achieve fast, energy-efficient connectivity with coherent pluggables and intelligent photonic line systems that make deployment and management simple

Backbone DCI: Scale with industry-leading WaveLogic™ 6 Extreme 1.6 Tb/s coherent optics and Ciena's C&L-band open line systems, designed to maximize efficiency and fiber capacity

Submarine DCI: Extend reach with easy-to-deploy SLTE solutions and high-performance coherent optics that maximize fiber capacity for the longest subsea distances

With proven expertise in DCI networking, Ciena simplifies AI-driven growth and innovation—providing scalability and reliability. Additionally, with investments in new scale-across architectures for geographically distributed AI training and multi-rail amplification solutions, Ciena ensures infrastructure is ready for next-generation AI and cloud applications.

Discover how Ciena can transform your DCI strategy and pave a path to success.

AUTOMATION AND OPERATIONAL COMPLEXITY

The growing complexity of AI infrastructure is driving increased reliance on automation.

Shore believes automation will be essential: "To speed up operational turn-up, optimize network resources for peak performance, and protect SLAs through proactive resiliency," he said.

Gage similarly argued: "Software automation is becoming equally important as the physical infrastructure itself."

AWS emphasizes end-to-end operational ownership. Callaghan summarized the company's philosophy succinctly: "Our philosophy is that you can't optimize what you don't own end-to-end."

Berti likewise stressed systems-level integration. "You can't design interconnects, power, or cooling in isolation anymore. High-density AI environments only work when those elements are engineered as a single system."

ECONOMICS AND INVESTMENT PRIORITIES

AI is reshaping network economics alongside technology. Phil Wong observed: "Most of the current CapEx is focused on building compute (GPU) capacity."

However, every expansion of compute requires corresponding investment in connectivity.

Gage argued: "Network CapEx is no longer discretionary. It's a prerequisite for realizing compute investment."

Berti believes operators are increasingly front-loading spending on networking because retrofitting later is inefficient. "We're seeing much more front-loaded investment in high-capacity fabrics, optical interconnects, and resilience," he said.

He added that compute, networking, power, and cooling are increasingly planned together as integrated systems.

From AFL's perspective, AI is fundamentally changing infrastructure economics. Taylor observed: "Hyperscalers are spending roughly $700 billion on data centers—more than double the entire global telecom CapEx."

As a result, priorities increasingly favor performance and deployment speed over minimum cost. "This shifts CapEx focus from 'least cost' to 'high performance, fastest deployment'," he added.

AWS pursues similar objectives through vertical integration. According to Callaghan, the company owns: "Every layer of the network stack—from the fiber in the ground and the components in the optical transceivers, through to the network hardware, firmware, operating system, and management software."

Wong cautioned that operators must nevertheless evaluate the long-term economics of new routes carefully, particularly where AI demand may not be complemented by traditional enterprise traffic.

CONCLUSION

AI infrastructure is fundamentally reshaping optical networking. The transition from traditional cloud architectures to distributed AI systems is driving unprecedented demand for deterministic performance, synchronized communications, and high-capacity data center interconnect.

As GPU clusters scale beyond the limits of individual facilities, scale-across architectures are extending AI fabrics across campuses, metropolitan regions, and long-haul networks. This evolution is accelerating the adoption of coherent optics and higher-speed interfaces while simultaneously increasing demand for additional physical fiber and more resilient outside-plant infrastructure.

The industry's biggest constraints increasingly lie outside the optical layer itself. Power availability, deployment timelines, and physical infrastructure coordination often present greater challenges than transmission technology. Meanwhile, approaching practical limits in spectral efficiency mean that faster optics must be complemented by higher fiber counts, multi-fiber architectures, and targeted expansion strategies.

Testing, validation, and observability are likewise becoming strategic capabilities. As AI workloads become more synchronized and more sensitive to latency, operators are moving toward continuous assurance, system-level validation, and real-time telemetry to ensure that increasingly complex infrastructures perform as intended.

Perhaps most importantly, networking is no longer viewed simply as an enabling technology. Across the interviews, participants consistently described connectivity as an integral component of AI infrastructure itself. Whether through advanced coherent optics, dense fiber deployments, software automation, or new approaches to testing and operations, optical networks are evolving into the foundation upon which hyperscale AI systems will be built over the coming decade.

ACKNOWLEDGEMENTS

AFL GLOBAL

AFL is a global manufacturer providing end-to-end solutions to the energy, broadband, data center, AI and industrial markets. From fiber optic cable to connectivity solutions and test and inspection equipment, AFL delivers what operators and data center providers need to scale capacity without sacrificing efficiency or cost control.

Learn more

CIENA

Ciena is the global leader in high-speed connectivity. We build advanced networks to support exponential growth in bandwidth demand—empowering our customers, partners, and communities to thrive in the AI era. With unparalleled expertise and innovation, our networking systems, interconnects, automation software, and services revolutionize data transmission and network management.

Learn more

NOKIA

Nokia is a global leader in connectivity for the AI era, providing the critical network infrastructure the world relies on. Every day, we're powering our customers with advanced connectivity across fixed, mobile, and transport networks, delivering the performance and security they need to meet the demands of an AI-enabled future.

Learn more

VIAVI

VIAVI Solutions is a global leader in communications test and measurement and optical technologies. VIAVI enables customer innovation in industries ranging from communication networks, hyperscale and enterprise data centers to consumer electronics and mission-critical avionics, transportation, aerospace and anti-counterfeiting systems.

Learn more

RCR Wireless News
NEWS. INSIGHT. IMPACT.

RCR Tech
NEWS. INSIGHT. IMPACT.

Scaling Optical Networks for the Hyperscale and AI Era

By Juan Pedro Tomas, Editor, RCRTech