

The Rise of Carrier NaaS

Examining Network-as-a-Service and the New Middle Mile

RESEARCH BRIEF




Table of Contents

Executive Summary..... 1

Introduction and Background..... 1

What’s in a Name? (New Middle Mile and Carrier NaaS).....2

Revisiting Workloads Impacting the Middle Mile..... 4

The Carrier NaaS and Middle Mile Ecosystem5

Recent Trends Impacting the Middle Mile6

Carrier NaaS Standardization Efforts 10

AvidThink’s Observations on NaaS and the Middle Mile 11

Conclusion..... 12

Research Briefs are independent content created by analysts working for AvidThink LLC. These reports are made possible through the sponsorship of our commercial supporters. Sponsors do not have any editorial control over the report content, and the views represented herein are solely those of AvidThink LLC. For more information about report sponsorship, please reach out to us at research@avidthink.com.

About AvidThink

AvidThink is a research and analysis firm focused on providing cutting-edge insights into the latest in infrastructure technologies. Formerly SDxCentral’s research group, AvidThink launched as an independent company in October 2018. AvidThink’s coverage includes 5G infrastructure, enterprise networks, private wireless, edge computing, SD-WAN, SASE, SSE, ZTNA, cloud infrastructure, and infrastructure security. Our clients range from Fortune 500 enterprises and hyperscalers to tier-1 communications service providers, fast-growing unicorns, and innovative startups. AvidThink’s research has been quoted by Forbes, the Wall Street Journal, Light Reading, Fierce Networks, Mobile World Live, and other major publications. Visit AvidThink at avidthink.com.

The Rise of Carrier NaaS

Examining Network-as-a-Service and the New Middle Mile

Executive Summary

The convergence of generative AI, multi-cloud adoption, and evolving enterprise demands are fundamentally reshaping the middle-mile network landscape. As organizations grapple with unprecedented demands for high-performance connectivity between data centers, clouds, and edge locations, traditional network architectures and business models are proving insufficient. This transformation is driving a renaissance in Network-as-a-Service (NaaS) offerings, with major carriers worldwide reimagining their service portfolios to deliver the flexibility and scalability enterprises require.

AI infrastructure investments alone are projected to reach approximately \$1 trillion in capital expenditure over the next few years¹, creating massive demands for network capacity and driving new architectural approaches. Meanwhile, BlackRock projects 60-80% annual AI data center capacity growth through 2024², fundamentally altering traffic patterns and infrastructure requirements. This surge in AI workloads and evolving enterprise needs for multi-cloud connectivity and edge computing capabilities are reshaping middle-mile networks.

Major carriers are racing to position themselves for this transformation, with companies like Lumen securing \$5 billion (with \$7 billion or more pending) in new business driven by AI connectivity demands³. However, success in this evolving landscape requires more than sheer capacity – it demands new approaches to network architecture, orchestration, and service delivery. This report examines how carrier NaaS is emerging as a crucial enabler of AI and digital transformation, offering insights into market dynamics, technological innovations, and strategic implications for stakeholders across the ecosystem.

Introduction and Background

The middle mile landscape has evolved significantly since our inaugural report in April 2023. When we published that analysis, ChatGPT 4 had just succeeded the original ChatGPT (GPT 3.5), and Anthropic's Claude had newly entered the market. Coming off the Metaverse hype, network operators focused on media delivery, cloud applications, mobile 4G/5G traffic, and collaboration tools. Network demands were driven by growth in connected users, proliferating devices, and streaming media consumption, particularly short-form videos (TikTok, YouTube Shorts, Instagram Reels, Snapchat Spotlight).

The past eighteen months have seen generative AI (GenAI) emerge as a dominant driver for network operators – initially for training but eventually for inferencing. Financial analysts project approximately \$1 trillion in capital expenditure for GenAI infrastructure over the next few years, encompassing data centers, power generation and transmission, computing hardware, and networking¹. Network operators across the telecommunications, hyperscale, and infrastructure sectors are racing to increase network capacity and reach while harboring concerns about the long-term efficacy and scalability of GenAI.

Financial analysts project approximately \$1 trillion in capital expenditure for GenAI infrastructure over the next few years ... Network operators across the telecommunications, hyperscale, and infrastructure sectors are racing to increase network capacity and reach while harboring concerns about the long-term efficacy and scalability of GenAI.

¹Goldman Sachs, "Gen AI: too much spend, too little benefit?" Jun. 27, 2024.

²"AI Revolution Will Add Fuel to Data Center Boom, BlackRock Says," Datacenterknowledge.com, Jul. 11, 2024.

³"AI Demand Drives \$5 Billion in New Business and Massive Expansion of the Internet," Lumen.com, 2024.







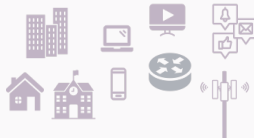



What’s in a Name? (New Middle Mile and Carrier NaaS)

Revisiting the New Middle Mile

In last year’s edition of this report, we introduced the New Middle Mile (NMM) concept, arguing that middle-mile decisions are no longer purely driven by connectivity concerns — computing requirements increasingly influence infrastructure choices in this domain. We demonstrated that distributed computing sites are crucial to enabling intelligent applications, including security gateways, enterprise control systems, and AI-powered video surveillance. These sites complement traditional content delivery networks by enabling general-purpose computing functions closer to users.

We also showed how the role of the middle mile evolved as the Internet matured, from providing basic connectivity to being a place where intelligent functions lived and high-demand content could be stored to improve end user experience.

EVOLUTION OF THE INTERNET (1980 TO 2024)

A. Incubation		Role of Middle Mile Infrastructure Basic connectivity.	
B. Commercialization		More connectivity across more locations. Access to private data centers and remote sites.	
C. Democratization		Lower cost, more capacity, richer connectivity options. Improved access to web providers.	
D. Cloudification (Centralization)		Direct access to clouds, improved access to exchange points, CDNs. Increased resiliency.	
E. Edge-ification (Redistribution))		Edge application hosting, software-defined orchestration, dynamic marketplaces.	



With the focus on AI/ML today, it’s another obvious workload to be accelerated by using the middle mile. While our previous report discussed AI/ML inference as an edge workload facilitated by middle-mile connectivity, infrastructure operators now report interest in hosting generative AI inference services at edge locations. This, combined with AI training driving new data center locations, which in turn drives new middle-mile and first-mile connectivity, supports the thesis of middle connectivity being a joint computing-networking decision. The pattern of centralized training and decentralized inferencing aligns with our original ideas about the Internet’s evolution toward a more distributed architecture.

The Resurgence and Rise of NaaS

Concurrent with these developments, we've witnessed the resurgence of network-as-a-service (NaaS), a concept that has periodically gained attention since the early days of software-defined networking (SDN). While SDN initially focused on data center operations, its goals of disaggregation, programmability, and centralized policy control proved equally relevant to wide-area networks. Carrier NaaS offerings today embody those same SDN principles.

The increasing need for enterprises to connect multiple hyperscale clouds, colocation sites, private clouds, partner services, and remote offices has renewed interest in on-demand connections and NaaS offerings. However, there aren't standards for carrier-centric NaaS today (though industry bodies like the MEF are attempting to do so), and offerings from major carriers worldwide, such as BT, Verizon, Colt, and Lumen, are pretty diverse.

As we've indicated in previous writings⁴, NaaS is fundamentally a business and consumption model. Whether you apply NaaS to campus or carrier networks, the following principles apply:

- **Short-term commitments:** These services are on-demand or only require shorter-term contracts.
- **Quick provisioning:** Designed for rapid deployment, tear-down, scaling, testing, and experimentation.
- **Consumption-based pricing:** Billing aligns with usage, using a pay-as-you-go model.
- **Low capital expenditure:** Minimal upfront costs for enterprises to use these services. The costs may be amortized into the subscription fees.
- **Fully-managed or partially-managed:** Services come with some level of managed support, either partially managed or fully managed.

Enterprise Demands and the Appeal of As-a-Service

Today's enterprises face mounting pressure to maintain profitability in a post-zero interest-rate policy (ZIRP) environment. Despite apparent economic stability in North America and Europe, many technology companies continue to seek productivity gains through cost reduction, return-to-office mandates, and technological innovation in digitization, cloud adoption, and AI adoption.

These pressures have led enterprise customers to prioritize agility, multi-cloud connectivity, and scalability in their networking strategies. NaaS offers an attractive solution by enabling consumption-based network resource utilization. This model particularly appeals to organizations running AI and machine learning workloads, which require high-capacity, low-latency connections between data sources, data lakes, and AI training clusters across private and public clouds. Additionally, enterprises are adopting multi-cloud strategies to find available AI computing resources and avoid vendor lock-in. While NaaS offerings simplify life for enterprises, they introduce greater complexity in middle-mile network management for the network operators.

The rate of NaaS adoption will depend on the evolution and demands of enterprise workloads (which we will cover next). Workload characteristics that will impact the value of NaaS include:

- **Diversity of data sources and destinations**—Where will the data come from, and where is it going? From private data centers to multiple public clouds? To new GPU-as-a-Service destinations?
- **Provisioning timescales** — How fast do these connections need to be made available after being requested?
- **Variability of sources and destinations**—How many sources and destinations will the enterprise need at a point in time? How often will these change? Will it be a relatively static cloud on-ramp, or will the destination change as rare AI computing resources become available?

⁴ R. Chua, "Embracing the Many Faces of NaaS - AvidThink," AvidThink, Sep. 28, 2023.

- **Workload throughput requirements and variability** — What's the size of the data transfer? 100TBs of training data? Or 5-100GB pre-trained models? Or streaming media or IoT feeds? How consistent is the need for peak throughput?
- **Quality of service (QoS) requirements** — Are there any latency or jitter constraints? Any resiliency and uptime requirements?
- **Predictability of bandwidth use and connectivity** — Are there predictable cycles, and can the connections be scheduled in advance? How far in advance? Or is there limited forward visibility into when connections are needed?
- **Data sensitivity and compliance** — Does the data contain sensitive information that needs special handling even with encryption? Are there unique geo-sensitive routing rules?

The higher the variability, the more “on-demand” the workloads, the greater the number of sources and destinations, and the more diversity over time, the more likely NaaS will be appreciated and needed, and along with that, the willingness to pay a premium for the service.

Revisiting Workloads Impacting the Middle Mile

To understand what's shaping the middle mile and the strength of NaaS uptake, we have to examine the workloads riding on top of the middle mile network. Today, the transformation of work, entertainment, and industrial applications continues to reshape network infrastructure requirements. Understanding the evolving workload patterns below is crucial for accurate planning of future investments and service offerings.



- **AI and Machine Learning** — In the last 18 months, no other workload category has impacted middle mile buildout as much as AI and machine learning. The massive datasets required for AI training and the need for distributed inference capabilities have created new challenges for middle-mile infrastructure. High-bandwidth connectivity between data lakes and AI training centers and low-latency paths for inference workloads has driven significant infrastructure investment. Edge computing facilitated by low-latency middle-mile networks presents new opportunities for AI workload optimization. While many generative AI inference tasks may not require edge processing, the “goldilocks” zone of computing embedded within the middle mile remains strategically crucial for specific use cases and applications. We'll have more to say on this in the trends section below.



- **Collaboration and Meeting** — The normalization of hybrid work (despite increasing return-to-office mandates) has fundamentally altered network infrastructure demands. While platforms like Microsoft Teams, Cisco WebEx, and Zoom initially drove network requirements during the pandemic, today's collaboration needs to extend beyond basic video conferencing. Modern workplace collaboration encompasses high-quality video streaming with multiple participants, real-time document collaboration, virtual whiteboarding, and deep integration with productivity applications. The ability of NaaS or middle-middle offerings to improve the routing of collaboration traffic or to provision added bandwidth during large-scale collaboration can help maintain the QoS for such applications.



- **Gaming and Entertainment** — The gaming and entertainment sector represents one of the middle-mile infrastructure's most demanding workload categories. Cloud gaming services have introduced strict latency requirements, while the rise of esports and gaming content creation generates ongoing streaming demands. The continued push toward higher resolution content, with 4K streaming becoming standard and 8K on the horizon, creates unprecedented bandwidth requirements. Even as content quality gets upscaled, short-form videos increasingly dominate access networks, with intelligent middle-mile networks critical in connecting content delivery networks (CDNs) to end-user devices for fast delivery. NaaS capabilities can help support these applications by providing optimized network paths with minimal jitter and packet loss. The ability to dynamically allocate bandwidth during peak gaming periods or commuting hours, combined with direct connectivity to major gaming platforms and content providers, ensures consistent performance for end users.



- **Surveillance, Safety, and Control** — Industrial IoT (IIoT) and Internet of Medical Things (IoMT) applications have introduced a new class of mission-critical workloads to the middle mile. These applications, ranging from real-time video surveillance to industrial control systems and medical device monitoring, demand high performance and absolute reliability and security. The stakes are high in this category, as network failures could impact physical safety or critical medical services. An optimized middle-mile network can address these requirements through dynamic bandwidth allocation and enhanced security through network segmentation. The ability to provide low-latency paths for real-time control applications and redundant connections for high availability makes NaaS well-suited for these demanding use cases.






- **Cybersecurity** — The increasing digitization and connectivity of users and devices enlarge attack surfaces and increase the complexity of software and hardware stacks. With attacks growing and CSPs and others looking to protect the digital infrastructure, the middle mile provides a compelling place to host security gateways for SDWAN and SASE services. Similarly, NaaS offerings are increasingly adding security options on top of basic connectivity; encryption, protection, and visibility help stem the growing threats from adversarial nation-state actors and cyber-criminals.

The Carrier NaaS and Middle Mile Ecosystem

The broader middle mile ecosystem consists of a variety of organizations, including communications service providers (regional and national, wireless, wireline, satellite), data center operators, large technology companies (Web properties), hyper-scale cloud providers, colocation providers, public and investor-owned utilities, grid operators, cooperatives, and academic institutions. For more details, we'd recommend looking at the [2023 edition of this report](#). Below, we will briefly cover the more significant players that make up the middle mile system.

- **Hyperscalers** — The cloud computing giants — AWS, Microsoft Azure, and Google Cloud Platform — operate vast global infrastructures with extensive data center networks. Their middle mile significance extends beyond computing resources to include substantial investments in subsea cables and satellite networks, controlling the source and destination of significant business and consumer traffic flows.
- **Datacenter Operators**—They manage computing facilities at varying scales, from global hyperscalers above to mid-tier providers (Oracle Cloud, Akamai/Linode, Switch, etc.) — there's overlap between categories. Their position in the middle mile stems from hosting essential computing services while facilitating traffic exchange through cross-connects and peering arrangements.
- **Colocation Providers** — Distinguished as specialized data center operators, colocation providers like Equinix and Digital Realty offer shared facility services encompassing space, power, cooling, and connectivity. Their strategic value in the middle mile lies in enabling efficient interconnectivity between diverse networks and hosting crucial Internet Exchange Points (IXP) while increasingly expanding into bare-metal and direct cloud connectivity services.
- **Communication Service Providers (CSPs)** — Traditional telecom operators maintain extensive network infrastructure across first, middle, and last mile segments. While somewhat constrained by regulatory frameworks, these entities are evolving beyond basic connectivity to offer advanced services like SD-WAN and SASE, leveraging their unique position in the network hierarchy to facilitate modern enterprise workloads.
- **Tower Companies** — Infrastructure specialists like American Tower, Crown Castle, and Cellnex support wireless connectivity through tower management (though some have diversified). Their middle mile relevance manifests in providing crucial backhaul and fronthaul network infrastructure, with many now expanding strategically into data center operations to strengthen their market position.
- **Content and Application Distribution Networks**—Well-known CDNs like Akamai, Cloudflare, and others have evolved beyond content delivery to offer comprehensive edge services, including security, distributed databases, and serverless computing. They are vital in optimizing data flow across all network segments, effectively bridging the gap between content sources and end users.



- **Network Equipment and Software Vendors**—The vendor ecosystem spans traditional networking giants (Cisco, Juniper, Nokia, Ericsson, Ciena, Huawei) and innovative disruptors (DriveNets, Graphiant, Arrcus). These organizations supply the fundamental hardware and software components that enable middle-mile functionality, increasingly focusing on virtualization and disaggregation. 
- **Orchestration Vendors**—Orchestration vendors that play in the middle mile and in carrier, NaaS include emerging players like Aarna Networks, which specializes in managing distributed computing environments (and has more recently expanded into serving GPU-as-a-Service clients). Likewise, network orchestration at various levels of the stack (such as Blue Planet, a division of Ciena) enables software-defined networks necessary to offer NaaS services. 
- **Other Players** — The ecosystem encompasses various entities, including government bodies, utility cooperatives, and specialized service providers like GIS companies and marketplace platforms (e.g., Connectbase). Other companies include Corning, which supplies the optical fiber used to connect across the middle mile. These diverse organizations contribute unique assets and capabilities to the middle-mile infrastructure, often filling crucial connectivity and service provisioning gaps. 

Recent Trends Impacting the Middle Mile

Now, we'll explore the key business, social, technology, and geopolitical drivers of the middle mile architecture and push for NaaS. The one dominant force that stands above the others is AI, particularly generative AI, so we'll start with that.

AI Data Center Build — Desperately Seeking Power (and GPUs)

The unprecedented demand for AI infrastructure is driving massive investments in data center capacity, with a particular focus on power availability and distribution — the search for gigawatts of new power for the energy-sucking GPUs critical to AI training (and inference) is top-of-mind for the data center operators and the extended ecosystem. This trend reshapes the middle mile landscape as new data center locations are selected based on power availability rather than traditional factors like population density or proximity to established fiber routes. The resulting network architecture must adapt to connect these new power-advantaged locations (that aren't near existing fiber routes nor in top data center markets) to existing data centers, cloud regions, and enterprise locations. Here's how the various groups within the middle mile are reacting to the AI wave:

- Data center operators have emerged as crucial enablers of the NaaS ecosystem, extending their role beyond traditional colocation and hosting services. Colocation and inter-exchange providers have similarly evolved their offerings, with companies like Equinix expanding to include advanced virtual connection services supporting bandwidths up to 50 Gbps for AI and machine learning workloads.⁵ Likewise, Digital Realty is expanding its capabilities to support AI workloads and innovations in carrier networks, as evidenced by its recent collaboration with Oracle to accelerate AI computing adoption among enterprises⁶. In a similar vein, MegaPort announced the launch of their AI Exchange, which enhances connectivity through a global network platform to support AI workloads⁷.
- Tower companies have expanded beyond their traditional role as passive infrastructure providers. Through strategic acquisitions and investments in edge computing capabilities, companies like American Tower, Crown Castle, and Cellnex position themselves as critical enablers of distributed cloud architecture. American Tower's acquisition of CoreSite in 2021 (along with prior smaller data center purchases) signaled a broader transformation toward integrated digital infrastructure provision, combining traditional tower assets with data center and edge computing capabilities. Their recent joint venture to build a new data center with Stonepeak⁸, an infra-focused alternative investment firm that bought 29% of American Tower's data center business in 2022, indicates an ongoing evolution of the strategy.

⁵"Equinix Announces Faster Virtual Connections to the Cloud; Partners with Google Cloud to Support Workloads Including Artificial Intelligence and Machine Learning," Equinix, Jun. 19, 2023.

⁶"Digital Realty and Oracle Further Strengthen Collaboration to Enable the Next Generation of Cloud Native and AI Enterprise Workloads," Digitalrealty.com, 2024.

⁷"Introducing MegaPort AI Exchange," MegaPort.com, 2024.

⁸"Stonepeak Forms Joint Venture with American Tower to Construct a New Data Center in Denver, Colorado," Stonepeak, Jul. 31, 2024.

- Major carriers have recognized the AI opportunity and are positioning themselves accordingly. Lumen's recent securing of \$5 billion in new business, driven by AI connectivity demands, underscores the scale of this opportunity. Lumen is upgrading and expanding its intercity network as part of its Private Connectivity Fabric offering. Further, Lumen's tie-up with Meta, Microsoft, and AWS⁹, plus their recent agreement with Corning to reserve 10% of Corning's global fiber capacity for each of the next two years to support AI workloads¹⁰, highlights AI's importance. Fellow fiber provider Zayo is upgrading its long-haul and metro networks across the US to 400Gbps (eventually to 800Gbps, then to 1.6Tbps), driven by the demand from hyperscalers for more robust broadband to support AI workloads¹¹. Zayo has highlighted the push for bulk long haul orders of 144 fiber count versus 12 fiber count in the past; meanwhile, within metro areas, orders are pushing 288+ fibers compared to 24-48 in the past.
- Edge computing facilitated by low-latency middle-mile networks presents new opportunities for AI workload optimization. Verizon's recent earnings calls have highlighted the potential for edge-enabled AI inference¹², though questions remain about the practical requirements for such deployments. While many generative AI inference tasks may not require edge processing, the "goldilocks" zone of computing embedded within the middle mile remains strategically crucial for specific use cases and applications¹³. T-Mobile's partnerships with NVIDIA and Ericsson aim to integrate AI capabilities directly into their network infrastructure¹⁴.

This generative AI trend reshapes the middle mile landscape as new data center locations are selected based on power availability rather than traditional factors like population density or proximity to established fiber routes.

Network APIs and NaaS

Traditional communication service providers find themselves at a crucial juncture, balancing the need to maintain reliable infrastructure while adapting to new services like NaaS. UK-headquartered Colt's NaaS solution, Colt Internet On Demand, was originally introduced in 2017 — one of the earlier commercial NaaS offerings. In 2019, the company unveiled their scalable Internet On Demand service globally. More recently, Colt has demonstrated API-based end-to-end service integration with Proximus, a Belgian-based telecoms company, using standard body MEF's LSO Sonata APIs¹⁵ for cross-carrier orchestration.

Tier-1 carriers like BT in Europe likewise push flexible NaaS offerings for the middle mile by upgrading their network portfolio. BT's recent launch of its Global Fabric¹⁶ required streamlining and modularizing its portfolio, developing a cloud-centric API-enabled multi-service edge, and relocating its POPs to service provider edges near cloud scalers, ensuring deep interconnection and operational resilience for AI and cloud consumption.

Disaggregation, Virtualization, and Open Architectures

Disaggregation and modular network architectures are gaining traction, allowing CSPs to replace monolithic systems with flexible configurations from multiple vendors. This approach aligns well with NaaS, enabling dynamic, software-based orchestration of services and resources, leading to greater agility and efficiency. By separating hardware from software, disaggregation simplifies network upgrades and fosters innovation in the NaaS market.

⁹"Lumen and AWS Partner to Unlock the Power of Generative AI Through Enhanced Network Operations and Delivery," Lumen Newsroom, 2024.

¹⁰"Corning and Lumen Reach Supply Agreement on Next-Generation Fiber-Optic Cable to Support Data Center AI Demands," Lumen Newsroom, 2024.

¹¹L. Hardesty, "Zayo says AI is already driving demand for fiber," Fierce Network, Oct. 31, 2023.

¹²Verizon, "Verizon shows mobile customer gains, promises AI opportunities," Lightreading, 2024.

¹³Verizon, "Verizon execs hint at making money from AI infrastructure," Lightreading, 2024.

¹⁴T-Mobile, "T-Mobile Announces Technology Partnership with NVIDIA, Ericsson and Nokia," T-Mobile Newsroom, 2024.

¹⁵"Colt and Proximus trial new Proof of Concept to test API collaboration and End-to-End Network as a Service automation", Colt Technology Services, Jul. 10, 2024.

A new category of platform providers has emerged to serve these needs. Companies like Graphiant (funded by Sequoia Capital, Atlantic Bridge, and others to the tune of almost \$100M) are taking innovative approaches to network service delivery, focusing on software-defined networking and advanced orchestration capabilities. Using a centralized policy approach with stateless software gateways that can be instantiated on on-demand private or public cloud servers, Graphiant enables a high-performance content-aware SDN that supports rich observability and sophisticated policy-based routing that respects compliance, reduces costs, and improves performance.

Meanwhile, DriveNets, which has raised \$587M over multiple rounds from 14 investors, including (Bessemer Venture Partners, D2 Investments, Pitango, and others), is focused on building cloud-native networking software and network disaggregation solutions that run on white-box networking equipment. DriveNets aims to reduce the CapEx and OpEx of carrier networks through disaggregation and increase agility via their converged software platform that can power middle-mile infrastructure.

Other players in the space that take a software-centric approach including Arrcus, providing a converged networking traffic that stretches from carrier core to mobile to data center networking equipment.

Government Funding and the BEAD Program

The Broadband Equity, Access, and Deployment (BEAD) program represents a significant shift in the government's approach to network infrastructure funding. However, its impact on middle-mile development has been more limited than anticipated. While the program's total funding of \$42.5 billion is substantial, only \$1 billion has been specifically allocated to middle-mile infrastructure.

Since our report last year, the NTIA Middle Mile Grant Program has allocated \$980 million to fund projects covering more than 370 counties across 40 states and Puerto Rico¹⁷. While awardees are investing an additional \$880 million in matching funds, the program's scope remains constrained relative to demand, as evidenced by the NTIA receiving over \$7.47 billion in oversubscribed funding requests¹⁸. Some states recognized this gap and initiated complementary programs, such as California's

\$3.25 billion Middle-mile Broadband Initiative. However, this patchwork approach to funding raises concerns about long-term sustainability and regional disparities in network capability.

While awardees are investing an additional \$880 million in matching funds, the program's scope remains constrained relative to demand, as evidenced by the NTIA receiving over \$7.47 billion in oversubscribed funding requests.

Open Access Networks

Open access networks, which allow multiple service providers to leverage shared infrastructure, have gained traction as a cost-effective way to expand middle-mile connectivity. This model reduces entry barriers for new providers and fosters competition, helping to close connectivity gaps, particularly in rural areas. BT's OpenReach is one such example in the UK, meanwhile, in the US, Zayo's open-access fiber network in Nevada, backed by \$153M in funding and spanning 800 miles¹⁹, is an example of how this approach can advance both commercial and public policy goals, making middle-mile infrastructure more accessible across a broader region. AT&T and BlackRock's Gigapower joint venture is another example where they're looking for multiple ISP tenants on their network. In an example of the embrace of open access, AT&T has signed contracts with other open-access providers: Boldyn Networks, Digital Infrastructure Group, PRIME FiBER, and Ubiquity²⁰.

¹⁶ "BT readies customers for Global Fabric infrastructure," ComputerWeekly.com, 2024.

¹⁷ Masha Abarinova, "NTIA awards last middle mile grants to four states," Fierce Network, Sep. 07, 2023.

¹⁸ NTIA, "Enabling Middle Mile Broadband Infrastructure Program | BroadbandUSA," Doc.gov, 2023.

¹⁹ "Zayo Granted \$153 Million for Middle Mile Initiative in Nevada," Zayo.com, Jun. 26, 2024.

²⁰ L. Hardesty, "The fiber race gets a billion-dollar boost with Corning-AT&T deal," Fierce Network, Oct. 28, 2024.

Hybrid and Multi-Cloud Connectivity

As enterprises adopt hybrid and multi-cloud strategies, the middle mile is becoming increasingly vital for connecting diverse cloud environments securely and efficiently. Enterprises will likely pursue hybrid multi-cloud strategies for existing cloud and AI workloads. CoreSite, in its 2024 State of the Data Center report, indicated that 98% of organizations have adopted or plan to adopt a hybrid IT model²¹. This trend is driving demand for flexible, high-performance connectivity solutions that can seamlessly integrate various cloud environments while maintaining security and compliance requirements.

Intelligent Edge with NaaS: SD-WAN, SASE, and more

Secure Access Service Edge (SASE) combines SD-WAN with security functionalities, providing a secure, high-performance connectivity solution for middle-mile networks. MEF's standards for SASE, in conjunction with dynamic SD-WAN capabilities, encourage carriers to deliver unified security and networking services as part of an overall NaaS blueprint. This trend is significant for NaaS providers, as it offers a holistic approach to network management and security across distributed infrastructures. For now, most of the ecosystem (including AvidThink) view Carrier NaaS as primarily lower-layer services (IP layer and below) with optional privacy added via encryption, but we acknowledge the value of a more expansive definition.

In line with this trend for converged networking and security and aligned with our New Middle Mile concept of combined networking and computing, companies like Akamai and Cloudflare now support various services, including security, distributed databases, and edge workload hosting in serverless form factors. Their extensive global networks and experience managing distributed infrastructure make them valuable partners for carriers looking to enhance their NaaS offerings with advanced edge capabilities.

Geopolitical Influences, Data Sovereignty, and Compliance

Data sovereignty requirements increasingly influence network infrastructure, necessitating careful consideration of data location and routing decisions. These regulations present new challenges for NaaS providers, as they must ensure compliance with evolving international data governance laws while maintaining high service quality. This trend is particularly relevant for middle-mile infrastructure, where flexible traffic routing and workload placement can help providers adapt to diverse regulatory environments. NaaS offerings from carriers and solutions from players like Graphiant allow intelligent routing policies that can stipulate countries or regions to be avoided while maintaining appropriate redundancy and resiliency SLAs. Similar mechanisms can also be used for sustainability (which has become less of a driver due to the over-focus on AI and "find computing resources at all cost" sentiment) – for example; Colt offers enterprise customers the ability to stipulate a "green" path for their end-to-end connectivity that has a reduced carbon footprint.

Security and Digital Resiliency

As cyber threats escalate, middle-mile infrastructure security has evolved to incorporate robust, multi-layered defenses. DDoS protection at network aggregation points has become standard, while zero-trust architecture and quantum-safe encryption are emerging as essential elements in future-proofing middle-mile networks. Digital resiliency is also increasingly critical, with operators investing in redundant paths, failover mechanisms, and infrastructure hardening to protect against environmental challenges and maintain essential connectivity.

Post-Quantum Cryptography (PQC) and Quantum Key Distribution (QKD)

With quantum computing on the horizon, middle-mile providers are exploring post-quantum cryptography (PQC) to secure data against future quantum-based attacks. This includes implementing quantum-resistant encryption algorithms and developing QKD infrastructure to ensure that NaaS services remain secure as quantum technologies advance. Backward compatibility with existing cryptographic standards is also essential, enabling a seamless transition as quantum-safe encryption standards are adopted.

²¹"CoreSite Releases 2024 State of the Data Center Report Highlighting Acceleration of Hybrid IT Model Adoption to Support Today's IT Workloads," Coresite.com, Jul. 16, 2024.

Non-Terrestrial Networks: Integrating LEO Satellites

Low Earth Orbit (LEO) satellites are increasingly integrated into middle-mile infrastructure to enhance coverage and resilience. In conjunction with 5G non-terrestrial network standards, LEO satellite networks provide additional redundancy and enable connectivity in underserved regions. This convergence with terrestrial networks presents new opportunities for edge computing, IoT, and latency-sensitive applications, enhancing middle-mile network capabilities. LEOs can also play a role in flexible network architectures, supporting dynamic routing and, depending on source and destination locations, providing lower latencies than terrestrial and subsea fiber connections.

AIOps and AI-Driven Network Orchestration

Artificial Intelligence for IT Operations (AIOps) is emerging as a vital tool for automating middle-mile network management. By leveraging AI and ML, AIOps enables real-time network monitoring and adaptation, reducing operational costs and improving response times. As NaaS deployments expand, AI-driven orchestration will be essential for managing complex, distributed networks.

The orchestration layer has become increasingly critical as carriers seek to manage complex middle-mile networks spanning data center and cloud on-ramps, aggregation points from access networks, 5G mobile network towers, and more. Many networking vendors have expanded their solution suites to encompass flexible NaaS solutions, from OSS providers like Amdocs and Netcracker to network equipment providers' software groups. For example, Blue Planet, a division of Ciena, provides an open and standards-based NaaS framework that enables telecommunication providers to dynamically create a cloud-like experience across their networks. These solutions increasingly incorporate AI capabilities into their operations (AIOps) and enable intent-driven declarative frameworks coupled with closed-loop assurance for improved manageability and scalability.

Carrier NaaS Standardization Efforts

Industry standardization efforts, led by organizations like MEF, are important in developing a cohesive framework for NaaS deployment and operation. As a global industry association comprising network, cloud, security, and technology providers, MEF is actively working with carriers and other organizations to accelerate enterprise digital transformation through a collaborative ecosystem approach.

The organization's recent Network-as-a-Service 2025 Industry Blueprint²² serves as a guide for service providers and enterprises. It defines NaaS as the comprehensive combination of on-demand connectivity, application assurance, cybersecurity, and multi-cloud-based services across a standards-based automated ecosystem. This proposed definition provides a foundation for consistent service development and delivery across the industry.

MEF's related development of Lifecycle Service Orchestration (LSO) APIs represents an advancement in facilitating automation between ecosystem partners. These APIs provide interoperability capabilities for business and operational functions, offering context clarity for attributes and including product and service models for various NaaS services. The organization's partnership with TM Forum also strengthens these efforts, aligning MEF's LSO APIs with TM Forum's Open Digital Architecture and Open API standards.

The orchestration layer has become increasingly critical as carriers seek to manage complex middle-mile networks spanning data center and cloud on-ramps, aggregation points from access networks, 5G mobile network towers, and more.

²²"Network-as-a-Service," MEF, Nov. 2024.

AvidThink's Observations on NaaS and the Middle Mile

As we look ahead to 2025, several key developments will likely shape the middle mile landscape. We'll start with some general observations and then call out a few elements that we think deserve particular attention.

High-Level Observations for 2025

The generative AI boom will continue driving unprecedented demand for high-capacity middle-mile connectivity between data centers, though potentially at a more measured pace as enterprises better understand their actual AI infrastructure needs. We expect increased investment in dark fiber and wavelength services between major AI/ML training facilities, alongside growing demand for flexible, on-demand connectivity to support hybrid AI architectures.

Network-as-a-Service adoption will accelerate, but challenges around standardization and interoperability will persist. Enterprise acceptance of consumption-based networking models will grow, though the market will likely experience continued fragmentation of NaaS offerings before meaningful consolidation begins. Carriers will face significant operational challenges as they adapt to support true NaaS delivery models.

The middle mile funding gap will become more apparent as last-mile BEAD deployments increase demand for backhaul capacity. This situation may drive more state-level middle-mile funding initiatives and increased private investment in open-access middle-mile networks. Creative public-private partnerships will likely emerge to address rural connectivity needs. It's possible that the AI boom in driving middle-mile connectivity buildout could help offset this; it's unclear whether access on-ramps for underserved communities will improve as this depends on where new data centers get built out.

Security and compliance requirements will add complexity to middle-mile architectures, with a growing focus on quantum-safe encryption capabilities and enhanced requirements for traffic inspection and threat mitigation. More stringent data sovereignty controls will increasingly affect traffic routing decisions and network architecture.

Parting Thoughts

- **Known Unknowns and Unknown Unknowns of Generative AI** — AI and specifically, Generative AI is driving a lot of the boom in middle mile connectivity (as well as data center interconnect and data center networking investments). GenAI is still an early technology and not fully understood (unlike traditional predictive AI). Many AI infrastructure investments today are driven by fear of missing out. With tens to hundreds of billions of dollars pouring in, we do not know if today's GenAI approach will scale with computing resources, training data size, and existing model architectures (transformers, derivations, and other variations). We believe OpenAI's GPT-5 will be better than GPT-4o, just as GPT-4o is better than GPT-4, and GPT-4 was better than GPT-3.5 — but this is far from guaranteed.

Likewise, today's open-source training corpora are around 100TB (compressed), which will likely not be moved around once transferred to a data center cluster. Further, because GPUs and other AI accelerators are in short supply, enterprises must sign 12-month or longer contracts, reducing the need to move data between data center locations. Synthetically generated data for training will likely sit within a data center cluster instead of being transferred regularly between sites. On the other hand, new multimodal frontier models might require ongoing large transfers of image, audio, and video training data. Time will bring increased visibility, but infrastructure builders and their financiers run the risk of over-investment, as we did recently with 5G infrastructure, and right before the telco crash in 2000.

- **NaaS Confusion**—There's a lack of clarity around NaaS, and while MEF is trying to help define a blueprint, this will take time. Some early telco NaaS offerings saw little to no traction because enterprises did not know how to consume them. We suggest building backward from enterprise needs and involving them in the early stages of the process and have shared more detailed thoughts on this in a [blog post on MEF's site](#).
- **Layers of Abstraction** — Multiple players in the ecosystem are touting the need for flexible services at different layers for different workloads. As in computing, which has bare metal servers, containers on bare metal, virtual machines, containers in virtual machines, and serverless, along with their parallel IaaS (infrastructure-as-a-service), PaaS (platform-as-a-service), CaaS (container-as-a-service), SaaS (software-as-a-service) consumption taxonomy, NaaS will have the same.

Whether enterprises pick dark fiber, wave, or ethernet services that provide varying degrees of “rawness” of connectivity will depend on the nature of the workload. AI training workloads that demand high throughput may push an enterprise towards a wave-type service that’s nailed up for 12-24 months between enterprise private data centers and GPU-as-a-Service operators, while cloud workloads may be fine with on-demand high-speed direct internet access (DIA) link.

- **Dangers of Over-Pivoting**—While GenAI is the shiny new object today, the bulk of traffic on the middle mile (and last mile) will be streaming media. Most end-user interactions over the Internet will be with cloud-based applications. Even as we try to accommodate the potentially huge AI workloads, the ecosystem should not give today’s proven revenue-generating cloud-based workloads the short shrift.
- **Agility and Flexibility**—Regardless of how the middle mile evolves, infrastructure providers (carriers, data center operators, hyperscalers) will want to build more agility into their underlying capabilities. Orchestration and flexible software-centric, cloud-centric architectures should be areas of investment.

Conclusion

The middle-mile network infrastructure is at an inflection point. The convergence of AI workloads, multi-cloud adoption, and the shift toward NaaS delivery models are reshaping the technical requirements and business models for middle-mile connectivity.

While standardization efforts by MEF and others are helping to define the future of carrier NaaS, significant challenges remain. The apparent underfunding of middle-mile infrastructure through programs like BEAD could constrain growth as demand accelerates, though the AI drive to increase middle-mile capacity could help close the gap. Additionally, the complexity of supporting emerging technologies like quantum-safe encryption and non-terrestrial networks will require sustained investment and innovation.

Success in this evolving landscape will require close collaboration between traditional carriers, hyperscalers, and an emerging ecosystem of specialized platform providers. Those who can navigate the technical, operational, and business model transitions ahead will be well-positioned to capture value from the ongoing digital transformation of enterprise networks. The middle mile’s strategic position between the access edge and core infrastructure assets like data centers makes it critical to enabling the next generation of digital services, particularly as enterprises increasingly embrace hybrid IT architectures and AI-driven applications.

We welcome feedback on our reports; please reach out to us at research@avidthink.com

SPONSOR INTERVIEWS SECTION FOLLOWS

Please support our sponsors — check out in-depth interviews with their experts on Carrier NaaS the the middle mile. We encourage you to visit their websites to learn more about their solutions.

Q&A with Khalid Raza, CEO and Founder of Graphiant



Q: Telcos have traditionally focused on data transit, but you suggest they should shift focus. Could you elaborate on why this is crucial for telcos now?

A: Absolutely. For decades, telcos have been primarily concerned with data transit, building networks to transport data efficiently. However, as services and applications evolved, many began bypassing telco infrastructures or using over-the-top (OTT) platforms for content and service delivery. This shift has left telcos out of the value chain, losing ground to tech companies. Now, telcos have a real opportunity to transform and leverage their infrastructure for data processing rather than merely transit. They're in the closest proximity to where data is generated, giving them a unique advantage. By focusing on private data processing at the network edge, they can tap into entirely new revenue streams.

Q: What strategic lessons do you think are essential for telcos to learn from tech giants like Amazon, Google, and Microsoft?

A: Telcos need to rethink their entire business model, moving from being traditional telecommunications companies, or "telcos," to becoming tech-driven organizations, or "techcos." Leading tech companies centralize services in the cloud, simplifying service definitions across devices. Telcos can adopt a similar model, where the service definitions reside in a centralized infrastructure and localize the actual execution to their network edge. This shift could open doors in business-to-business connectivity, especially considering the anticipated data growth. The current model of backhauling data to centralized cloud locations isn't feasible long-term due to bandwidth limitations. With compute becoming cheaper than bandwidth, there's a strong case for telcos to transform their middle mile and points-of-presence (POPs) into data-processing hubs.

Q: How can telcos leverage their existing POP infrastructure to enhance their value in the AI economy?

A: POPs are a hidden gem for telcos. They're often located where the last mile of connectivity lands, providing a unique opportunity to create dedicated business connectivity services. By modernizing their POP infrastructure, telcos could offer these locations as secure, isolated DMZs (demilitarized zones) for data exchange between business partners, especially for ephemeral, AI-driven workloads. Instead of enterprises building their own DMZs, they could use these decentralized DMZs from telcos for secure, business-class connectivity, creating a new revenue stream centered around B2B data exchanges.

Q: Software-defined networking (SDN) was touted as a transformative technology, but adoption has stalled. What direction do you believe SDN should take for telcos?

A: The initial promise of SDN was indeed transformative, but its implementation has veered off course. SDN should be about state abstraction rather than just provisioning and visibility. With true state abstraction, telcos could separate network state from physical devices, enabling changes without the need for constant reprovisioning. Without proper SDN principles, we risk overloading the middle mile with excessive compute and networking functions, like issues seen in MPLS VPNs, where customer-specific data in transit drove up costs. By moving toward a state-abstracted SDN, telcos can simplify network management and open infrastructure up for data processing, positioning themselves as leaders in the AI economy.

Q: What's unique about the Graphiant approach and why is it a fit for today's enterprise needs?

A: Graphiant's unique innovation and architecture combines the advantages of traditional MPLS, SD-WAN and SRv6 to create a new wide area network that is entirely programmable using metadata. The Graphiant Stateless Core is a globally distributed private backbone that utilizes a new metadata protocol to remove all customer state information from the data plane. It delivers private connectivity to resources and services (applications, clouds, business partners) with direct any-to-any paths and lower latency. The Stateless Core provides programmable pathing & SLAs for customers via policy and metadata. With Graphiant, enterprises can consume a flexible NaaS solutions that brings improved control, flexibility, and reliability while providing visibility.



Interview with Gabriele Di Piazza, VP of Products and Alliances, Blue Planet, a division of Ciena

Q: What are some leading challenges that carriers face today, and how does NaaS help address these?

A: Carriers are under significant pressure to reduce costs while improving agility in service delivery. NaaS, or Network-as-a-Service, provides a way to streamline service creation, deployment, and fulfillment through automation. Operators can achieve better economies of scale, allowing services to be composed and delivered more flexibly and efficiently.

Q: Can you explain what Network-as-a-Service (NaaS) encompasses and how Blue Planet approaches it?

A: NaaS encompasses various networking domains, including wide area networks, cloud networks, campus networks, mobile, and 5G networks. For Blue Planet, NaaS has three essential components. First, there's intent-based declarative orchestration, which exposes the NaaS services. Second, dynamic inventory capabilities give a precise view of network assets and allow them to be reconciled across services. Finally, AI-driven operations and assurance ensure service quality and reliability across these different systems, which is critical for seamless service delivery.

Q: What makes NaaS appealing to carriers today?

A: Carriers are increasingly drawn to NaaS because it enables them to evolve their service offerings to an on-demand, consumption-based model. This approach aligns with the industry's shift toward cloud-based models, enabling faster creation and delivery of new services. For carriers, NaaS enhances service flexibility and creates opportunities for upselling and cross-selling, making it an attractive model for future growth.

Q: How does AI play a role in NaaS implementations?

A: AI is crucial in NaaS, especially for operations and assurance. AI helps operators shift from reactive resolution to preventative maintenance by tapping into network state and performing automated network re-configuration or providing early warning. From designing services to resolving potential issues, AI ensures that network operations are proactive, reliable, and optimized.

Q: How important are APIs to the success of NaaS?

A: APIs are fundamental to NaaS. The complexity of new dynamic network services demands an API-first approach, where different functions and operations can be decomposed and integrated seamlessly. Industry organizations like MEF, Linux Foundation, TM Forum, and GSMA have promoted prescriptive APIs, which help create a standardized, end-to-end service delivery framework like Open Gateway/CAMARA. Blue Planet aims to drive interoperability with these APIs to ensure a consistent and composable service model across multiple vendors and technology domains.

Q: Is NaaS exclusive to Tier-1 carriers, or is it accessible to others?

A: NaaS is not limited to Tier-1 carriers like AT&T or Verizon. There's a broad spectrum of NaaS use cases, including NaaS-like SD-WAN and SASE services, where smaller players or pure-play vendors can provide valuable offerings. What matters most is the service's packaging and simplicity, which makes NaaS accessible to a broader range of providers. In some markets, smaller operators have found innovative ways to bring NaaS to customers without needing to be at a Tier-1 scale.

Interview with Dudy Cohen VP of Product Marketing DriveNets



Q: What's the main objective for carriers when looking to modernize their middle mile infrastructure?

A: Carriers today need flexibility above all. The middle mile has become pivotal in enabling service agility, which lets operators respond to new market demands and rapidly roll out services. This agility isn't just about adding bandwidth or expanding reach; it's about creating programmable networks that allow carriers to introduce new services quickly and expand their revenue streams. Achieving such flexibility requires converged infrastructure, where different service silos—like fixed and wireless, residential and business—are integrated into a unified architecture. This convergence allows a network to adapt to diverse and evolving customer needs.

Q: Convergence sounds complex. Why is it essential, and how does it benefit carriers?

A: Convergence transforms the network into a holistic system, simplifying operations and eliminating the barriers of isolated, disparate systems. Imagine managing everything—different access types, service levels, and customer requirements—from a single infrastructure platform. That's the vision. Such a unified approach reduces operational complexity and, by eliminating silos, helps carriers respond faster to market demands. It also enables better automation and programmability, critical elements in modern networks where flexibility and rapid deployment capabilities are essential.

Q: You previously discussed a "turbo button" concept for on-demand services. Could you explain how this might work?

A: The "turbo button" is an idea rooted in the need for instant service modification. Years ago, telecom companies tried to create a button that allowed customers to boost their connectivity at will, but they faced challenges like rigid OSS/BSS systems and complex multi-vendor environments.

Today, with advances in cloud and programmable networks, we're closer to realizing this concept. Modern programmable infrastructure eliminates the need to pre-provisioned every service scenario and allows, instead, dynamic, on-the-fly configuration changes, allowing customers to easily toggle between service profiles. This approach addresses many past barriers, letting carriers offer seamless, self-service upgrades similar to cloud services and thus approach new target markets and create new revenue streams.

Q: How does programmability help carriers achieve these goals?

A: Programmability fundamentally shifts how networks operate. It allows carriers to configure network services to resemble cloud environments—where a single click can adjust resources and optimize performance. This isn't just a feature but a necessity for modern carriers. To reach this level, carriers need a converged, software-centric, disaggregated network, where the network's intelligence lies in its software. This architecture ensures that the network is flexible and future-ready, enabling carriers to quickly adapt to market shifts without the limitations of traditional hardware-centric designs.

Q: What's next for carriers as they continue evolving their infrastructure for NaaS?

A: Automation will be the next major phase, making networks "autonomous-ready." Heavy automation, possibly incorporating AI, will minimize manual intervention and make operations more efficient. Carriers are moving toward almost autonomous networks, with the flexibility to offer dynamic, on-demand services and self-optimization capabilities. Ultimately, a fully automated, software-driven network infrastructure is the goal, bringing carriers closer to a seamless, service-oriented model that aligns with the expectations set by cloud service leaders.





AvidThink, LLC
1900 Camden Ave
San Jose, California 95124 USA
avidthink.com

©2024 AvidThink LLC. All Rights Reserved.
This material may not be copied, reproduced, or modified in whole or in part for any purpose except with express written permission from an authorized representative of AvidThink LLC. In addition to such written permission to copy, reproduce, or modify this document in whole or part, an acknowledgment of the authors of the document and all applicable portions of the copyright notice must be clearly referenced.