MARKET BRIEF

NaaS: The Automated Network Supply Chain for Agentic AI

November 2025

Created by:
Stan Hubbard, Principal Analyst

MPLIFY

MPLIFY

Market Brief

NaaS: The Automated Network Supply Chain for Agentic AI

November 2025

Key Findings

- AI is shifting from single-response transactions (generative) to continuous reasoning (agentic), driving massive inference growth across a multi-trillion dollar AI infrastructure buildout.
- Networks transform from utility to foundational coordination layer, enabling reasoning flows across distributed environments.
- Agentic AI requires a deterministic, programmable, on-demand fabric with lifecycle automation across providers.
- Network-as-a-Service (NaaS) is *the* automated supply chain uniting certified connectivity, automation, cybersecurity, and new revenue services.
- NaaS depends on alignment between buyers and sellers through standardized Lifecycle Service Orchestration (LSO) APIs, payloads, and processes that enable consistent automation of business and operational functions across ecosystems.
- Mplify certification validates service performance, LSO automation conformance, and cybersecurity, distinguishing strategic AI infrastructure partners from commoditized providers.

Executive Summary

The global race to build AI infrastructure, from hyperscale data centers to national digital sovereignty programs, marks the most consequential industrial expansion in decades. While focus remains on GPUs, data center growth, and power supply, connectivity has emerged as the decisive bottleneck and strategic opportunity for service providers. NVIDIA projects inference demand will rise a billionfold as AI evolves from single-response transactions to continuous reasoning loops, requiring networks that sustain real-time intelligence flows across globally distributed networks.

Agentic AI marks a fundamental departure from the 2023–2024 generative AI wave. These systems maintain persistent context, orchestrate distributed workflows, and coordinate autonomous actions. Intelligence is shifting from model-level analytics to infrastructure-wide operation, with AI factories, data hubs, and edge environments acting as a unified reasoning fabric. **Connectivity is no longer a supporting utility; it is the foundational coordination layer of intelligence production**.

Traditional architectures built for predictable traffic cannot support AI systems that dynamically reconfigure in real time. McKinsey estimates $6.7 trillion in data center investment by 2030, including $5.2 trillion for AI-ready capacity. Without programmable connectivity, compute investments risk under-delivering. As enterprises deploy agentic AI, **intelligence production requires federated connectivity ecosystems, not isolated provider networks**.

This transformation gives rise to Network-as-a-Service (NaaS) for AI, a new operating model built around four strategic themes:

Connectivity for AI: Carrier Ethernet, wavelengths, Carrier Ethernet over broadband, IP Broadband/ DIA and certified Carrier Ethernet for AI performance profiles.

Automation for AI: Lifecycle Service Orchestration (LSO)-enabled NaaS payloads, Model Context Protocol (MCP) and Agent-to-Agent (A2A) communication, federation frameworks, and LSO automation certification.

Cybersecurity for AI: Certified Secure Access Service Edge (SASE), Security Service Edge (SSE), Zero Trust, and quantum-safe Carrier Ethernet and wavelength services.

Revenue Services for AI: GPU-as-a-Service and AI Model-as-a-Service.

NaaS unites on-demand connectivity, application assurance, cybersecurity, and multi-cloud networking within a standards-based automated ecosystem. Through Mplify certification programs, service providers validate the performance, automation, and security capabilities agentic AI demands, positioning certified providers at the forefront of the multi-trillion-dollar infrastructure buildout.

Market Context – From Generative to Agentic AI: The Infrastructure Inflection

The Global AI Infrastructure Buildout

The reasoning phase of AI is reshaping digital infrastructure at unprecedented scale. In October 2025, NVIDIA CEO Jensen Huang projected that inference demand could increase a billionfold as AI systems evolve from single-response transactions to continuous chains of reasoning—what he calls "thinking AI." He describes this transition as an "extraordinary industrial revolution," a global race to construct AI factories unifying compute, power, and connectivity.

Intelligence Production Becomes Infrastructure-Dependent

Intelligence has become infrastructure-anchored, no longer confined to analyzing data but dependent on compute, energy, and networks to operate and act in real time. **Agentic AI makes infrastructure an active participant in intelligence production**.

In traditional cloud computing and generative AI, intelligence resides in models. Infrastructure transports inputs and outputs, but reasoning happens within bounded computational environments. Agentic AI fundamentally changes this: reasoning chains now span multiple data sources, model providers, and execution environments simultaneously. Networks don't just deliver results; they coordinate the continuous flows that enable reasoning. Without programmable connectivity, latency control, and cross-domain orchestration, agentic systems cannot maintain the persistent context and real-time coordination that defines their capability.

The dependency creates a strategic opportunity for service providers. When intelligence production depends not just on bandwidth but on programmability, automation, and federated orchestration, connectivity becomes a differentiating input to advantage. **Enterprises with superior infrastructure deploy agentic systems faster, scale them more reliably, and orchestrate them more effectively than competitors**. This positions providers not as utilities but as strategic partners in intelligence generation, enabling premium positioning and long-term relationships tied to business outcomes.

The infrastructure buildout reflects this imperative. McKinsey estimates $6.7 trillion in data center investment by 2030, with $5.2 trillion tied to AI workloads, as global compute demand nearly triples. The bottleneck is shifting decisively from compute to network programmability, security, and automation. Omdia projects AI-related network traffic will rise from 79 exabytes per month in 2025 to 820 exabytes per month by 2030—a tenfold increase. **This traffic growth reflects infrastructure's new role: coordinating intelligence flows, not just transporting data**. A veteran US Tier 1 service provider executive told Mplify, "I have never seen this level of explosion in bandwidth requests."

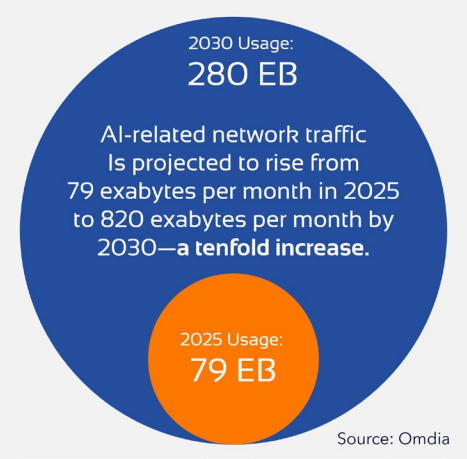

The CEO agenda mirrors this urgency. IoT Analytics found AI to be the leading technology theme in Q3 2025, mentioned in 45% of earnings calls from 4,500 companies. As reasoning workloads scale, networking—not power or compute—is emerging as the decisive constraint on AI growth.

Matching Network Architecture to AI Development Stages

This infrastructure expansion isn't just about capacity. It's about aligning network architecture with workload requirements. **Each stage of AI development demands a distinct connectivity approach optimized for its data flow, latency, and orchestration needs**.

AI STAGE	WORKLOAD CHARACTERISTICS	CONNECTIVITY	KEY CAPABILITIES
Training	• Large-scale transfers • Predictable bandwidth • Latency-sensitive synchronization • Factory-to-factory or distributed training connectivity	**Wavelengths** 400G-1.6T+ **Carrier Ethernet** 10G-400G+	• Ultra-high capacity for model-scale training • Lowest possible latency • API-enabled automation • Deterministic Carrier Ethernet for regional training
Inference	• Distributed operations • East-west traffic dominance • Variable bandwidth needs • Multi-destination traffic • Continuous deployment	**Carrier Ethernet** 10G-400G+, Carrier Ethernet over broadband	• Multi-cloud fabric • Carrier Ethernet for AI SLAs (<3ms FD, 99.999% uptime, 0.001% FLR) • Symmetrical bandwidth for bidirectional AI data flows • API-enabled automation • Federated assurance and policy control
Agentic AI / Reasoning	• Continuous reasoning loops • Agent-to-agent and agent-to-model flows • Ultra-low and consistent latency • Real-time orchestration • Multi-modal, cross-domain data exchange	**Hybrid** Wavelengths, Carrier Ethernet, and Carrier Ethernet over broadband	• Cross-provider automation • Wavelengths for retraining • Carrier Ethernet for agent-to-agent coordination and edge inference • Carrier Ethernet over broadband for peripherals and distributed agents • Dynamic SLA management • Federated security

Automation-ready wavelengths and Carrier Ethernet provide the on-demand provisioning and dynamic management AI workloads require. Training workloads vary by scale: exabyte-scale transfers between AI factories depend on the dedicated capacity and ultra-low latency of wavelengths, while regional and distributed training leverages deterministic Carrier Ethernet. Inference workloads need advanced Carrier Ethernet for AI to support distributed multi-cloud operations and edge inference. Emerging agentic AI, with continuous reasoning loops spanning multiple providers and trust domains, demands hybrid architectures that combine wavelengths for retraining, Carrier Ethernet for agent-to-agent coordination, and Carrier Ethernet over broadband for peripheral and distributed agent connectivity. For development, testing, and administrative traffic across all AI stages, as well as peripheral connectivity in distributed inference deployments, Mplify-standardized DIA, broadband Internet access, and IP VPN provide cost-effective connectivity with consistent LSO automation.

Wavelengths serve dual roles: as dedicated point-to-point optical paths for training, and as the high-capacity physical layer over which high-speed Carrier Ethernet services deliver Layer 2

automation. Automation through LSO APIs unifies these services across federated infrastructure, forming the foundation of the NaaS connectivity layer for AI.

The Rise of Enterprise Agentic Systems

Agentic AI has moved from boardroom strategy to operational reality. Mentions of "agentic AI" on corporate earnings calls rose 40% quarter-over-quarter in Q3 2025, according to IoT Analytics. Goldman Sachs CIO Marco Argenti framed the stakes: "The advancements in agentic AI have been incredible. Having AI that can not only answer questions but reason and plan to complete tasks and collaborate with other agents is a game changer across every industry." He is leading 12,000 engineers in building a hybrid workforce of humans and AI agents.

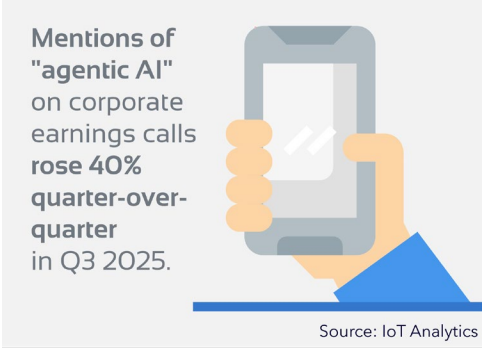

Source: IoT Analytics

The shift is visible in traffic patterns. Data center operators report unprecedented east-west traffic growth as AI agents communicate directly, a fundamental change from human-initiated requests to continuous agent-to-agent reasoning flows. The race is no longer to deploy AI models but to orchestrate them, enabling reasoning agents to collaborate across clouds, data sources, and enterprise boundaries without friction. Service providers must adapt networks to support continuous agent-driven operations rather than human-initiated transactions.

Leading service providers are pivoting from connectivity suppliers to AI infrastructure providers. This reflects recognition that enterprises increasingly evaluate network providers based on their ability to support AI workloads, not just deliver bandwidth. Service providers achieving this through standards-based certification and federated orchestration gain ecosystem partnerships and multi-provider interoperability that proprietary approaches cannot deliver.

Why Traditional Networks Fall Short

Agentic AI demands a new class of infrastructure. As Lumen CEO Kate Johnson put it, "AI needs data, data needs data centers, and data centers need to be connected." Digital Realty's Travis Ewert adds that this connection must deliver "scaled-out, private, performant, on-demand" capabilities: what the industry calls NaaS for AI.

Traditional networks excel at fixed connectivity but falter with the dynamic, distributed nature of agentic AI. They fall short in four key areas:

- Manual processes in a dynamic world. Legacy networks depend on manual workflows taking weeks or months, from qualification and ordering to provisioning and settlement. Performance monitoring remains siloed within providers, lacking automated coordination across boundaries. Agentic AI demands real-time adaptability as reasoning chains span multiple execution environments; a responsiveness manual processes cannot deliver.
- Compute-centric thinking. Infrastructure investment has centered on GPUs and data-center capacity, assuming connectivity works. This mindset leaves networks as an afterthought, provisioned reactively rather than strategically. Without high-capacity, low-latency, programmable connectivity linking AI factories, data hubs, and enterprise edges, expensive compute resources at times sit idle.

- Fragmented visibility and assurance. Traditional monitoring stops at domain boundaries, offering incomplete insight into distributed AI workflows. Agentic AI requires end-to-end observability and automated remediation across providers, clouds, and execution contexts, capabilities current tools lack.
- Security built for static perimeters. Agentic AI operates across changing, multi-domain environments requiring continuous authentication, dynamic policy enforcement, and federated zero-trust architectures. As AI agents collaborate across systems, the attack surface expands rapidly, exposing gaps that traditional security models cannot close.

To meet agentic AI demands, **networks must evolve into programmable, automated, and federated systems that deliver cloud-like flexibility and policy-driven orchestration**.

Sovereign AI and the Network Infrastructure Imperative

Nations increasingly view AI as a strategic asset requiring domestic control over training, inference, and data security. This sovereign AI imperative—projected by Oppenheimer as a **$1.5 trillion market opportunity**—extends beyond models and compute to the full infrastructure stack. Agentic AI amplifies both the complexity and urgency of achieving network sovereignty. Mplify has identified nearly 30 sovereign-AI deployments now underway with telecom providers worldwide.

Networking infrastructure is foundational to AI sovereignty. Without control over connectivity, data sovereignty remains incomplete regardless of domestic compute capacity. Nations require three critical capabilities:

- Infrastructure and Performance Control: Networking assets residing within borders under domestic management, delivering high-speed, low-latency fabrics capable of handling exabyte-scale training transfers and real-time inference. Without this foundation, data residency requirements cannot be met and AI capabilities face performance bottlenecks.
- Security and Operational Autonomy: Traffic segmentation, encrypted communications, and auditable routing protecting sensitive AI data from unauthorized access. Independent network management and troubleshooting capabilities ensure core AI services remain operational during geopolitical disruptions without dependency on external providers.
- Regulatory Compliance: Network control ensuring adherence to jurisdiction-specific sovereignty regulations while enabling authorized cross-border collaboration where beneficial.

Service providers with domestic infrastructure, standards-based connectivity, and certified performance profiles are positioned to capture sovereign AI opportunities.

The NaaS Solution for AI

NaaS for AI transforms networks from bottleneck to amplifier. *Mplify's 2025 NaaS Industry Blueprint* defines NaaS as combining on-demand connectivity, application assurance, cybersecurity, and multi-cloud networking within a standards-based, automated ecosystem, providing the programmable infrastructure agentic systems require at the speed and scale of AI reasoning. NaaS delivers lifecycle automation, end-to-end observability, and dynamic trust frameworks that adapt to distributed AI operations.

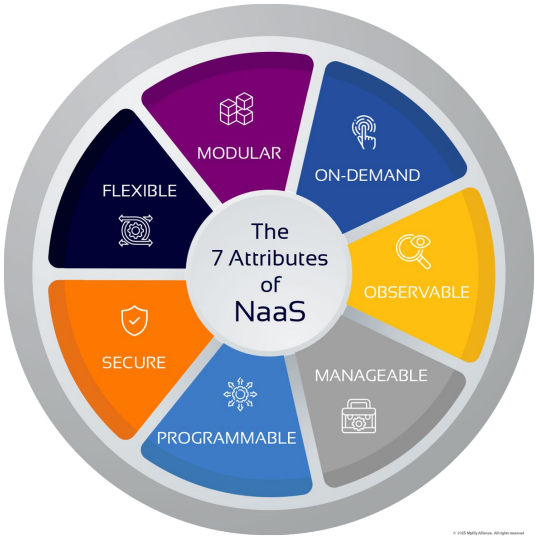

NaaS services are characterized by seven core attributes (on demand, observable, manageable, programmable, secure, flexible, and modular) that transform delivery and consumption models. These attributes and key features are detailed in *Mplify's NaaS Customer Experience Whitepaper*.

For AI workloads, NaaS addresses agentic infrastructure requirements across four strategic themes. Connectivity for AI establishes deterministic connectivity foundations. Automation for AI enables real-time multi-provider coordination. Cybersecurity for AI protects distributed reasoning across trust domains. Revenue Services for AI unlocks new business models. Each theme is detailed below.

Network-as-a-Service (NaaS) for AI, a new operating model built around four strategic themes:

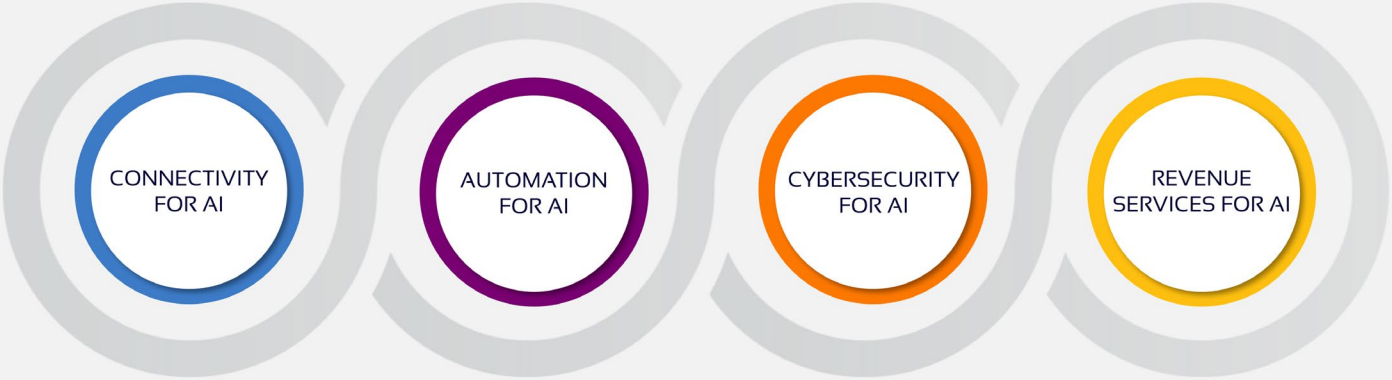

Connectivity for AI

AI's evolution from training to reasoning has elevated network connectivity from a background utility to a critical coordination layer. Meeting the scale, latency, and assurance demands of continuous reasoning requires a deterministic, programmable, and globally available foundation. For most providers, the NaaS journey begins with connectivity. Mplify members are evolving Carrier Ethernet and wavelength networks into dynamic, automated systems that form the foundation of AI-ready infrastructure. Mplify's strategy is organized around three connectivity domains—Carrier Ethernet, wavelengths, Carrier Ethernet over broadband—each optimized for

distinct roles across AI operations, from large-scale training and model synchronization to inference, reasoning, and edge coordination.

Carrier Ethernet

Carrier Ethernet provides the scalable, programmable backbone of AI infrastructure connectivity, delivering the performance assurance, reach, and ecosystem interoperability essential to the AI economy. Carrier Ethernet's global ENNI fabric provides an immediately available foundation for linking data centers, edges, and enterprise sites, combining predictable latency, guaranteed throughput, and class-of-service control critical to continuous inference and coordination loops. With interconnections ranging from 10G for edge and regional training sites to widespread 100G and accelerating transition to 400G for high-performance deployments, Carrier Ethernet is well-positioned to meet AI demands.

Reaching its full potential requires lifecycle automation enabling on-demand provisioning, dynamic scaling, and federated assurance across providers. Enterprises that can acquire AI-optimized Carrier Ethernet circuits in days, not months, gain a competitive advantage.

Core Use Cases

- **Data Center to Data Center:** High-performance, cost-efficient interconnection leveraging globally deployed Carrier Ethernet ENNIs.
- **Edge to Data Center:** Optimal balance between AI performance and scalability, making Carrier Ethernet the preferred transport for edge-to-core connectivity.
- **Subscriber Premises to Edge:** Guaranteed performance for latency-sensitive AI applications such as industrial automation, financial trading, and real-time inference.

Mplify is developing a certification to validate Carrier Ethernet for AI. The Carrier Ethernet for AI profile will include specifics for symmetrical, committed bandwidth with ultra-low latency (≤3ms frame delay, ≤1ms frame delay variation), frame loss ≤0.001%, and uptime ≥99.999%. These parameters define a new assurance tier for certified Carrier Ethernet services capable of supporting continuous AI workflows and scaling to meet the evolving demands of agentic AI.

Mplify certification becomes the tangible proof of readiness. Providers that achieve Carrier Ethernet for AI certification demonstrate measurable differentiation, facilitating interoperability, automation readiness, and deterministic performance in the federated ecosystem.

Wavelengths

While Carrier Ethernet underpins inference and production connectivity, AI training and large-scale synchronization demand dedicated, lossless transport at extreme bandwidths. Wavelength services deliver this capacity with the lowest achievable latency, connecting AI factories engaged in model training and retraining.

Today, 100G and 400G wavelengths are widely deployed, with providers introducing 800G and 1.6T interfaces to accommodate model-scale transfers. However, wavelength interconnection

remains constrained by limited standardization and automation across provider boundaries, often requiring manual cross-provider optical turn-up.

Mplify is working with members to define wavelength payloads for Layer 1 services, enabling dynamic capacity management, pre-provisioned activation, and federated lifecycle orchestration. Along with published standards, this work will establish the foundation for certified wavelength services that deliver deterministic, high-capacity transport between hyperscaler AI data centers and distributed model training sites.

Complementary Technologies

Together, Carrier Ethernet and wavelength services form a complementary connectivity architecture. Wavelengths deliver the ultra-high-capacity, low-latency paths required for large-scale model training and synchronization between AI factories, while Carrier Ethernet provides deterministic, scalable connectivity for regional and distributed training workloads, distributed inference, and agent-to-agent coordination in continuous reasoning.

As operations shift from centralized training to distributed, inference-driven architectures, Carrier Ethernet (including Carrier Ethernet over broadband at the edge) becomes increasingly central, while wavelengths remain indispensable for high-volume model synchronization across global AI factories. Beyond these three AI-optimized domains, Mplify-standardized DIA, broadband Internet access, and IP VPN provide additional connectivity options with consistent LSO automation for supporting workloads.

Carrier Ethernet Over Broadband

As AI models extend toward the edge, connectivity must reach beyond core data centers to include sensors, controllers, and edge inference devices. Carrier Ethernet over broadband brings Layer-2 connectivity to mass-market access networks, enabling deterministic performance over PON, DOCSIS, satellite, and 5G fixed wireless.

These environments can be configured to support Carrier Ethernet with performance guarantees and committed information rates. Parameters (maximum frame size, information rate, and latency) depend on the access technology. Many broadband providers already operate ENNI interconnections, creating a ready foundation for integrating Carrier Ethernet over broadband into AI ecosystems.

Access E-Line over broadband extends the reach of Carrier Ethernet for AI, linking peripherals to regional and cloud data centers through a consistent operational model. This preserves Carrier Ethernet's simplicity, security, and performance while enabling uniform management and assurance across diverse access technologies.

Mplify is planning to introduce certification for Carrier Ethernet over Broadband with an AI profile.

Automation for AI

The promise of NaaS for AI depends on more than connectivity performance; it requires an orchestration layer that turns static capacity into programmable, federated resources. Many

providers automate internally or through proprietary interfaces, but these cannot scale to the multi-provider ecosystem agentic AI demands. Connectivity infrastructure cannot meet agentic AI requirements if automation depends on bilateral agreements rather than open standards. Agentic AI demands networks that function as unified supply chains where services can be composed, activated, and managed seamlessly across domains.

Standards and LSO-Enabled Payloads

This orchestration layer is powered by standards-based LSO APIs. Mplify research indicates **nearly 100 providers worldwide have implemented or are committed to implementing LSO APIs**, reflecting strong industry alignment around a common automation framework. Ecosystem participants apply LSO-defined processes—either directly or via portals/adaptors—using common data models and payloads across business and operational functions such as address validation, quoting, ordering, provisioning, performance monitoring, and settlement.

Certification programs are expanding to validate conformance to Mplify standards. Existing certifications cover core LSO API functions, including address validation, product offering qualification, quote, order, and inventory. In 2026, validation expands to product payloads (Carrier Ethernet, Internet access, SD-WAN, cross-connect, and wavelength) to standardize automation across high-value NaaS solutions.

Mplify is developing Domain Specific Languages (DSLs) from these payloads to create AI-native schemas that agents and models can interpret directly. These DSLs define the ontology, semantics, constraints, and context for network services, **enabling AI systems to understand service attributes, valid configurations, and relationships without human translation**. Reference MCP server implementations will demonstrate how to wrap LSO APIs while preserving DSL semantics, providing prescriptive guidance for production deployments.

Future phases will extend to AI API Wrappers (MCP) and Agent-and-Model frameworks, enabling AI systems to interact directly with LSO APIs and apply contextual reasoning to ordering, provisioning, and assurance. Together, these initiatives establish the foundation for AI-driven orchestration, enabling agents to autonomously consume and manage network services.

NaaS payloads span connectivity, application assurance, security, and multi-cloud networking—from Carrier Ethernet and wavelength to SD-WAN, SASE, and AI clouds. These standardized interfaces create a shared language for federation that allows connectivity, performance, and policy data to be exchanged consistently between buyers, sellers, and orchestrators in real time. The result is automation that scales from enterprise edge to cloud, reducing delivery times from months to minutes while ensuring interoperability across domains.

Intelligent Federation and Agent Interaction

Standardization also enables new modes of consumption. NaaS marketplaces are emerging where providers syndicate services on-demand, allowing enterprises to assemble multi-provider connectivity as easily as they configure cloud workloads.

For agentic AI, this shift is transformative. Through protocols such as the MCP and A2A communication, AI agents can request, activate, and optimize network resources programmatically, without relying on human operators to modify connectivity. Federated trust frameworks—built on identity providers and AI-specific authentication, authorization, policy enforcement, fingerprinting, and auditing—enable secure interactions across enterprises, service providers, and partner ecosystems. Automation becomes the bridge between the physical and cognitive layers of AI infrastructure, allowing reasoning systems to interact directly with the network itself.

Cybersecurity for AI

Agentic AI requires continuous, federated security beyond perimeter defense. As reasoning spans domains and agents exchange data and policies, new attack surfaces emerge. Security must be standardized, certified, and automated within the same NaaS framework that governs connectivity and orchestration.

Certified SD-WAN, SSE, and Zero Trust: Building Blocks of SASE

Mplify certifies SD-WAN, SSE, and Zero Trust technology solutions through independent testing that verifies policy consistency, identity enforcement, and threat protection at scale. Service providers deploying certified vendor solutions inherit these certifications, enabling them to offer validated security capabilities without developing proprietary solutions. When all three components are certified, they form a complete SASE architecture validated at each layer.

The framework encompasses three independently certified domains:

- SD-WAN: Delivers standards-based, application-aware connectivity with policy-driven routing across multiple underlays; certified solutions validate performance for AI workloads.
- SSE: Validates threat protection, data security, and performance in real-world conditions.
- ZT: Enforces continuous verification across identity, device, and application layers for policy-driven access in multi-domain AI environments.

Together, these certified components enable consistent policy enforcement and create clear integration paths to future LSO-enabled delivery. As standardization advances, they will support secure multi-provider orchestration and readiness for AI-driven detection and response.

Quantum-Safe Connectivity Services

The longevity and sensitivity of AI data and models demand encryption resilient to quantum-era attacks. Quantum-safe Carrier Ethernet and wavelength services—part of Mplify's certification roadmap—apply post-quantum cryptography and quantum key distribution to secure inter-domain traffic and federated AI workflows. These capabilities ensure that sensitive AI data transfers,

including model weights, inference telemetry, and policy exchanges, remain protected across multi-provider environments well into the post-quantum era.

New Revenue Services for AI

Service providers can monetize their role in sovereign and enterprise AI infrastructure through AI-centric services that unify compute, connectivity, and intelligence. Built on Mplify's LSO API framework and edge computing IaaS standard, these services turn networks into revenue-generating platforms.

Mplify is developing standardized payloads and emerging certification frameworks that automate AI-oriented service delivery of GPU-as-a-Service (GPUaaS) and AI Model-as-a-Service (MaaS). These services turn the network into an intelligent supply chain for AI training, inference, and reasoning across clouds and providers. Together with Mplify's connectivity, automation, and cybersecurity capabilities, they form a unified NaaS experience.

GPU-as-a-Service

GPUaaS delivers high-performance GPU resources at the edge for scalable, efficient inference and training. LSO-enabled GPUaaS supports automated ordering, provisioning, and assurance of GPU resources across service providers, reducing latency and enhancing performance for workloads such as real-time video analytics, generative design, and autonomous systems.

AI Model-as-a-Service

AI MaaS allows enterprises to consume pre-trained, domain-specific AI models on demand without managing infrastructure or lifecycle operations. Delivered at the network edge or within regional AI clouds, MaaS leverages Mplify-standardized APIs and partnerships among service providers, ISVs, and GSIs to deliver turnkey capabilities such as computer vision, predictive analytics, and intelligent automation with the federated deployment flexibility that complements hyperscaler models.

Standards and Certification

Mplify is developing standardized payloads and certifications for GPUaaS and AI MaaS. Certifications will validate performance, API interoperability, and assurance so providers can deliver AI services with the same consistency and federation as certified connectivity and security.

Together, these services position providers to capture a share of the sovereign AI market while expanding margins beyond traditional connectivity. As enterprises and governments value integrated AI capabilities over standalone transport, **standardized and certified AI service delivery transforms service providers from connectivity vendors into strategic AI infrastructure partners**.

Next Steps

The transformation to NaaS for AI is underway. Service providers ready to position as strategic AI infrastructure partners should:

- Validate connectivity readiness with Carrier Ethernet for AI certification, demonstrating deterministic performance for continuous reasoning workloads.
- Achieve automation certification for LSO APIs and payloads, enabling federated orchestration across provider boundaries.
- Implement certified security through SASE and quantum-safe profiles protecting distributed AI operations.
- Explore revenue services, including GPUaaS and AI MaaS, to capture margin expansion opportunities.

As the $6.7 trillion AI infrastructure buildout accelerates, certified providers positioned within standards-based ecosystems gain the competitive advantage that proprietary approaches cannot deliver.

Acknowledgements

Stan Hubbard (Principal Analyst, Mplify) – Author

Pascal Menezes (Chief Technology Officer, Mplify) – Contributor

About Mplify

Mplify is a global alliance of network, cloud, cybersecurity, and enterprise organizations working together to accelerate the AI-powered digital economy through standardization, automation, certification, and collaboration. As the defining authority behind Carrier Ethernet, Lifecycle Service Orchestration (LSO) APIs, and certified SASE and SD-WAN, Mplify has developed the global blueprint for Network-as-a-Service (NaaS) that is empowering the industry to innovate, interoperate, and scale trusted network services across a global ecosystem. For more information about Mplify Alliance, visit Mplify.net.

Disclaimer

NaaS: The Automated Network Supply Chain for Agentic AI