MEF White Paper

NaaS Customer Experience

Articulating the Essential Features of a NaaS Offering

September 2024

Disclaimer

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and MEF Forum (MEF) is not responsible for any errors. MEF does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by MEF concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by MEF as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this document. MEF is not responsible or liable for any modifications to this document made by any other party.

The receipt or any use of this document or its contents does not in any way create, by implication or otherwise:

a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor

b) any warranty or representation that any MEF members will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor

c) any form of relationship between any MEF member and the recipient or user of this document.

Implementation or use of specific MEF standards, specifications, or recommendations will be voluntary, and no Member shall be obliged to implement them by virtue of participation in MEF Forum. MEF is a non-profit international organization to enable the development and worldwide adoption of agile, assured and orchestrated network services. MEF does not, expressly or otherwise, endorse or promote any specific products or services.

Contents

1. Executive Summary

This white paper provides both users and providers of NaaS (Network-as-a-Service) offerings with a neutral baseline description of service features. It is based on expertise provided by the world's leading NaaS providers in the MEF Forum membership. Background information is also provided on the challenges faced by enterprises that are served by NaaS providers and the transformation that telecom service providers are undergoing to address those enterprise challenges. Also included are summaries of the stakeholder landscape in the NaaS industry as well as available commercial models for NaaS offerings.

The document is expected to evolve over time as MEF Forum members increase their offerings in the NaaS market segment. Additional complementary work in MEF to this NaaS Customer Experience White Paper includes the development of a NaaS federation of MEF members offering Network-as-a-Service and a MEF NaaS Blueprint, which details MEF's extensive offering of Network-as-a-Service building blocks for NaaS implementers, including MEF Standards and MEF LSO APIs.

2. Introduction

Enterprises are facing challenges with their existing network architectures due to huge shifts in their business models and requirements. These shifts are driven by evolving needs and expectations from their employees, consumers, and ecosystem partners, who are looking for enhanced value, convenience, and experience in conducting their business operations in a mix of hybrid cloud-centric as well as edge-based ecosystems.

Enterprises are looking for network-architecture transformations to manage these business-requirement shifts and overcome existing challenges within current networks. In turn, to meet these enterprise-network demands, service providers also need to transform their networks to deliver on-demand, programmable, and intent-based, connectivity solutions.

2.1 Enterprise Transformation Drivers

Enterprises are increasingly moving their workloads to cloud environments, which offer major benefits like scalability and agility, allowing enterprises to spin up applications, and scale them up and down, in near real time. However, the networks required to connect to these workloads are still often sold and managed in a legacy format, with long, fixed-term contracts and lengthy delivery timelines. Whereas enterprises can use services from cloud providers very quickly and for short durations, they cannot necessarily achieve that with the legacy network services that connect them to those clouds. For example, if an enterprise can buy cloud compute for a month-long project within ten minutes, but they must buy the network service that connects to that cloud for one year and it takes two weeks to put in place, then, from the enterprise's point of view, the network service is irrelevant. NaaS offerings need to match cloud offerings in terms of time to implement and length of contract. As a result, enterprises are now seeking connectivity services within a cloud-like consumption model that can be provisioned in near real time, scaled or adapted quickly, and are available under flexible commercial models. This paradigm empowers enterprises to react quickly to changing market conditions, for example scaling bandwidth up when demand increases or rerouting a circuit when an outage occurs. They increasingly expect immediate delivery of services, including delivering multiple services to the same access point in particularly high-density locations like major datacenters. Enterprises are looking for platforms on which the service provisioned is fully automated, allowing real-time changes to the network, and provisioning multiple services over the same access element. Although service providers may need to compromise the experience in some situations, like when manual work is required for physical access such as installing new fiber or equipment, they need to minimize those scenarios through strategies like pre-provisioning.

Enterprises also require the ability to mix and match short- and long-term agreements to accommodate a combination of network services. While they require some network services on a more permanent basis, they need the flexibility to order others on a short term, pay-as-you-go basis. Traditional billing models of fixed, monthly rental fees that align with fixed, static agreements no longer suit their needs. To ensure that enterprises can adjust their demand (and charges) to their consumption, they require the ability to change and flex the network in real time, multiple times a day or even per hour, with more granular pricing models like per-hour rates.

Aligned with cloud-like consumption, enterprises don't want to depend on a service provider salesperson and human interactions to change and adapt their network—they increasingly want to control their own network through self-service capabilities like real-time quoting, ordering, and where possible provisioning of the network. Initially, they will access self-service through web portals, but as the number of partner web portals that an enterprise needs to use continues to grow, they will become difficult to manage. As enterprises advance in their digital-transformation journey, they are increasingly looking for API support to ensure they can consume network services in a programmatic manner through their own internal applications.

API integration with network service providers will allow enterprises to integrate network orchestration into their own workflows, increasing efficiency and, most importantly, enabling new consumption models that foster innovation. As the use of APIs to consume network services increases, the availability of a standardized API definition that can be used across the industry becomes critical to ensure enterprises can integrate with new service providers quickly and efficiently and avoid vendor lock-in due to the high cost of integration.

2.2 Service Provider Transformation Drivers

To meet the requirements of connected enterprises operating in complex digital ecosystems with diverse applications, services, users, and devices, the service provider networks have in turn also been transforming to continuously adapt to changing business requirements.

Examples of Service Provider transformation trends include, as described in more detail below:

- From human resource-centric processes to automated service lifecycle management.
- From fixed network connections to agile, on-demand dynamic services.
- From traditional, service-based models to self-service, API-driven, usage-based models.
- From network availability as SLA to user experience as SLA.
- From reactive incident handling to automated, proactive service assurance management.
- From disaggregated network & security to converged network & security.

From Human Resource-centric Processes to Automated Service Lifecycle Management

To achieve operational efficiencies, the traditional human-resource-centric business and operational processes, which were designed with longer planning, upgrade, and delivery cycles in mind, are being migrated towards leaner, fully automated and highly scalable, digital IT-systems-centric workflows.

From Fixed Network Connections to Agile, On-Demand Dynamic Services

The traditional fixed service architectures have offered stable connectivity and predictable performance. However, as enterprises transform towards rapid cloud integration with expectations focused on application performance, there is a need for flexible, scalable, on-demand-network service provider solutions that customers can self-manage based on their application requirements, service quality requirements, and commercial considerations, such as 'On-Demand' versions of the following:

- Bandwidth: to scale capacity up (or down) from a base of committed connection bandwidth, in near-real time.
- Connectivity: to order and activate connectivity services in minutes rather than months.
- Route: to dynamically select a path of choice to meet specific performance parameters.
- QoS Management: to adjust the quality of service (QoS) parameters for customer traffic or application requirement.
- Security: to dynamically activate and change different levels of security policies based on customer applications.
- Virtual CPEs and VNFs: to create virtual CPEs and virtual network functions (VNFs) for their network in cloud-like, on-demand models.

From Traditional, Service-Based Models to Self-Service, API-Driven, Usage-Based Models

As enterprise businesses evolve into a cloud-first approach, their network requirements also need changes from traditional sales-organization-driven service models to self-service, cloud-like, usage-based models with new technology introductions. Manual service-provider business and operational workflows face challenges to adapt and scale due to a fragmented network architecture with multiple technologies, thus hampering their inability to align with future needs and network cost with business demands. These challenges have led service providers to offer customer-managed, self-service workflows and enable API-based integration into customers' network operations (net-ops) systems.

From Network Availability as SLA to User Experience as SLA

Traditional, disaggregated networks provide Service Level Agreements (SLAs) on specific aspects of service delivery, while enterprise network architecture and businesses rely heavily on applications and software, making application-based SLAs essential for ensuring optimal performance and consistent user experience. To that end, offering customers the ability to define business intents as expectations or objectives, which the platform automatically translates into low-level configurations and actions, is helping enterprises to simplify network management, reduce human errors, and ensure user-experience objectives are met.

From Reactive Incident Handling to Automated, Proactive Service Assurance Management

As the complexity of network and services grows, there's an urgent need to shift from traditional, manual incident handling to more sophisticated, data-driven, automated, and ideally AI-empowered, network and services incident and problem management. Modern service assurance solutions continuously track ever-changing configurations of enterprise customer networks and services to react to occurring incidents and problems as well as predict potential problems (e.g. capacity or hardware issues) and even recommend proactive resolutions. Broad automation and proactive management in this area is required to ensure customer experience requirements are met.

From Disaggregated Network & Security to Converged Network & Security

Internet and cloud-centric service architectures bring in new security challenges for enterprise networks. The current distributed nature of network and security architectures makes it difficult to implement end-to-end security-management practices, leaving businesses more vulnerable to cyberattacks. To address this challenge, a shift towards converged network and security solutions is a key requirement from enterprises, by leveraging an orchestration platform that enables the on-demand management of the network and, at the same time, the security functionalities.

3. NaaS Defined

MEF defines NaaS—Network-as-a-Service—as the combination of one or more of following: on-demand connectivity; application assurance; cybersecurity; and multi-cloud-based services, across a standards-based automated ecosystem.

4. Essential NaaS Features

The following are considered essential features of any Network-as-a-Service offering as described in more detail in this section:

1. Always-On Platform Access
2. Management Console
3. Standardized API Access
4. Identity Management and Role-Based Access Control (RBAC)
5. Product Catalog
6. Location Discovery
7. Multi-Service Dynamic Access
8. Real-Time Serviceability
9. Real-Time Firm Pricing
10. Granular Ordering
11. Bundle Ordering
12. Short or Zero-Commitment Contracts
13. Online Ordering, Including Signature and Acceptance of a Contract
14. Self-Service Contractual Changes
15. On-Demand Configuration Changes
16. Dynamic SLA Management
17. Automated Order Fulfillment
18. Usage Tracking
19. Pay-As-You-Use Billing
20. Regulatory Compliance

Always-On Platform Access

NaaS platforms have a high level of availability, typically close to 24x7x365, recognizing that customers use them outside business hours and expect the same platform service at all times.

Management Console

A management console, typically accessed via an Internet connection, is provided for the customer to access the NaaS platform, such that the customer can:

- Manage their account, optionally including any payment details.
- Discover, purchase, and manage services.
- Understand their consumption.

The console must have a modern security approach, including multi-factor user authentication.

Standardized API Access

A standardized API is provided for the customer to access the NaaS platform's functionality. The API is defined by the MEF standards definition process. All business and operational functionality is available via the APIs.

Identity Management and Role-Based Access Control (RBAC)

Identity Management refers to the policies, processes, and technologies used to manage and secure the digital identities of users within an organization. It includes the creation, maintenance, and deletion of user accounts, as well as the management of access rights and permissions. The goal of identity management is to ensure that only authorized NaaS users have access to the resources they need while protecting sensitive data from unauthorized access. Identity management typically encompasses user authentication, authorization, and the provisioning of user credentials.

RBAC allows the customer to define a set of roles for their NaaS platform users, which can include:

- Specific capabilities the role can access, such as purchasing.
- Functions the role can perform on specific areas or components of the deployed network.
- Cost centers to which a role can allocate spend.
- Management rights the role has on the account.

Product Catalog

The customer can consult a catalog to see the product offerings available for them to order on the NaaS platform.

Location Discovery

The customer can find locations, physical or virtual, at which products can be delivered with the appropriate level of granularity. These are typically locations for endpoints and cloud on-ramps, geographic boundaries for cellular services, etc., and may be delivered by the NaaS provider or one of their partners.

Multi-Service Dynamic Access

NaaS platform-based services are differentiated from traditional connectivity product offerings in that they enable access to multiple network services and service types from different providers or platforms, using a single access point or device, such as universal CPE or Virtual Network Functions (VNFs).

Real-Time Serviceability

The NaaS platform provides accurate serviceability to the customer in real time for the provisioning of new products or making changes to an existing product. Options for feasible technical configurations of a product offering are usually presented to the customer as a menu of options, such that the customer can select one to get pricing.

Serviceability for a product that will be delivered through zero-touch automation by the NaaS provider, or in collaboration with their ecosystem partners, must be firm and committed. Where manual processes will be used to extend the NaaS footprint to deliver the product, serviceability is provided on a best-efforts basis without firm commitment that the product can be delivered.

Real-Time Firm Pricing

The NaaS platform provides firm pricing to the customer for technical configurations for a product offering with commercial terms within a specific context for delivery in real time. Options for prices and contract terms are often presented to the customer as a menu so the customer can know what is available to them and select one for their product order.

Like real-time serviceability, products that will be delivered through zero-touch automation by the NaaS provider or in collaboration with their ecosystem partners must be firm and committed. Where manual processes will be used to extend the NaaS footprint to deliver the product, pricing can optionally be provided without firm commitment.

Granular Ordering

The NaaS platform is differentiated from the traditional approach by following a building-block approach where customers can choose atomic product offerings to construct and adapt their own end-to-end service over time. Each product offered is contracted individually and has its own commercial and operational lifecycle.

Bundle Ordering

The NaaS platform also allows the ordering of bundled product offerings, in case the customer wants to choose more than one atomic product offering to construct the end-to-end service (for example, bundling core underlay connectivity with SD-WAN overlay and customer web services), to achieve an optimized bundle price.

Short Or Zero-Commitment Contracts

Products can be contracted on a NaaS platform on a short-term basis, typically measured from an hour to a few days. This short duration means that these products can be disconnected or decommissioned at any time.

The exception to this rule is when the NaaS provider has expended significant capital expenditure to deploy infrastructure dedicated to a single customer and may need to force a longer-term commitment to recover the outlay. This exception must be clearly conveyed to the customer before signature and acceptance of the contract.

Online Ordering, Including Signature and Acceptance of a Contract

The NaaS platform allows the customer to submit a product order for a specific configuration and a contract term of a product offering using an online mechanism, such as the management console or the API, which represents a legally binding acceptance of a contract for the product that results from the product order's fulfilment.

Self-Service Contractual Changes

The NaaS platform allows customers to change the contractual conditions for a product that has already been fulfilled. This could be to enhance or extend the duration of a contract to improve the commitment of service availability or quality. Changes in contracts often include changes in pricing.

The customer can accept the terms (and pricing) of the updated contract using the management console or API, which results in a legally binding commitment.

On-Demand Configuration Changes

Customers can change the configuration on a temporary basis to get more capacity, or to change technical settings of a product, when these products have been contracted on a longer-term fixed basis. These configurations are charged at small increments of time, such as hours, often at a premium rate. The customer can return to the original configuration, or may change the configuration again, whenever they want and without any penalty.

Dynamic SLA Management

Dynamic SLA management involves modernizing service level agreements and contractual requirements to align with the evolving digital landscape. By utilizing agile metrics and flexible reporting, organizations can support a wide range of advanced services—including cloud, IoT, and beyond—ensuring that SLAs adapt to changing environments and service transformations. This approach enables responsive and scalable management of connectivity and digital services, maintaining high service quality and performance in a flexible, dynamic setting.

Automated Order Fulfilment

Product and service orders are almost always fulfilled in near real-time through end-to-end automation. The NaaS platform must provide detailed visibility of order fulfilment progress. Product orders are guaranteed to be fulfilled and not to fail due to lack of available capacity.

The only exception for product orders is when an initial physical deployment is needed to enable a new, typically remote location, or install hardware or cabling.

Usage Tracking

The NaaS platform presents information, in the management console and via the API, to enable a customer to track their consumption. This information includes capacity utilization and spend; it allows the customer to slice data based on several dimensions, including cost centers, locations, type, or who actioned the order. Data is kept up-to-date in near-real time.

Pay-As-You-Use Billing

Products are charged based on consumption (or utilization) or time-based subscription billing. Dimensions for consumption could include the amount of data transferred or the number of executions of an action or rule or number of users. Time-based subscription billing is done in a range of intervals that may be as small as a few hours. A customer can get an up-to-the-minute view of charges at any time.

Regulatory Compliance

Regulatory compliance in Network-as-a-Service (NaaS) involves ensuring that services meet legal and industry standards (e.g., GDPR for data privacy). NaaS service providers must handle these regulations to maintain secure and compliant network operations, adapting to the specific requirements of different regions and industries.

5. NaaS Customer Lifecycle

Traditionally, the telecom customer's lifecycle focuses around determining availability and pricing; only after finalizing the order (typically connectivity services), will the customer expect to receive support from the service provider.

In contrast, in the NaaS customer lifecycle, the service provider first onboards the customer to the NaaS platform and then engages with the customer to determine availability and pricing; the customer can expect support from the service provider before and during the customer's operation of the network. This change in the timing of the customer's commitment to use specific, chargeable items means that the customer can expect to engage with the NaaS provider extensively, even before committing to buy one or more components to fulfill their NaaS requirements.

The NaaS customer lifecycle can be broken down into three major areas:

1. Customer Onboarding by the NaaS provider: including working with the customer to jointly assess the networking, security, and application requirements and capabilities.
2. Network Operation: Continuous operation of the network as a service by the customer on the NaaS platform for as long as required, including:
    a. Discovery and ordering of NaaS components, including accepting incremental charges for every component ordered.
    b. Performing acceptance testing and delivery of individual components.
    c. Performance monitoring of the NaaS components.
    d. Creating support tickets and tracking the resolution of issues, as necessary.
    e. Terminating NaaS components.
    f. Recurring and/or one-time payments for the entire NaaS solution in effect at that time.
3. Customer Off-boarding by the NaaS provider.

6. The Application Developer's NaaS Experience

LSO APIs and programmability are key components of a NaaS offering. Application developers are provided with network-based APIs for integration into their application business logic. This capability enables developers to utilize connectivity resources alongside cloud compute and storage resources. Just as application developers can spin up virtual machines, or allocate storage space with ease, a NaaS offering provides developers with the ability to spin up connectivity.

6.1 NaaS Application Capabilities

Specifically, application developers use NaaS APIs to provide the following capabilities, as described in more detail below:

- Dynamic Network Provisioning
- Automated Scaling
- Service Orchestration
- Security Integration
- Event-Driven Networking
- Cost Optimization

Dynamic Network Provisioning

To request network resources (such as virtual private networks, load balancers, or firewalls) based on application needs. This dynamic provisioning ensures that the network adapts to changing workloads and traffic patterns.

Automated Scaling

When an application experiences increased demand, to scale network resources to accommodate the load—e.g. adding more bandwidth, adjusting QoS parameters, adding security parameters, or spinning up additional VPN connections.

Service Orchestration

To orchestrate complex network services without manual intervention—e.g. to define rules for traffic routing, implement failover mechanisms, or set up secure communication channels.

Security Integration

To integrate security features directly into application stacks—e.g. enforcing encryption, implementing access controls, or monitoring traffic for anomalies.

Event-Driven Networking

To trigger network actions based on application events—e.g. when a new user signs up, the application can automatically create a dedicated VPN tunnel for that user; similarly, when an application component fails, the network can reroute traffic to healthy instances.

Cost Optimization

To manage network resources and optimize costs—e.g. to spin down unused resources, adjust bandwidth dynamically, and avoid overprovisioning resulting in cost savings on cloud egress charges in addition to improved application experience.

6.2 NaaS Application Examples

LSO APIs bridge the gap between application logic and network infrastructure, allowing developers to focus on building robust, scalable, and secure applications without getting bogged down by network complexities. Examples include:

1. Gaming Industry:
    a. Dynamic Game Scaling: Game developers use NaaS APIs to dynamically scale server resources based on player demand; during peak hours, additional compute capacity can be provisioned automatically, ensuring smooth gameplay.
    b. Latency Optimization: NaaS is used to enhance the gaming experience by selecting optimal routes and minimizing latency between game servers and players.
2. Banking and Financial Services:
    a. Business Automation: to enable seamless integration between banking systems and external applications; for example, APIs facilitate payment processing, fraud detection, and credit scoring.
    b. Embedded Finance: to support embedded financial solutions, allowing non-financial companies to offer financial services within their platforms; APIs handle transactions, identity verification, and risk assessment.
3. Cloud Application Connectivity:
    a. Enables the customer network to access, via public or private connectivity, any applications hosted in clouds.
4. Multi-Cloud Interoperability:
    a. Cross-Cloud APIs: to enable multiple clouds to be used by the application as if they were just one cloud.
    b. Hybrid Cloud Support: to integrate on-premises networks with cloud-based NaaS.
5. Edge Compute Integration:
    a. Edge Services: to manage network functions at the edge (e.g., CDN caching, edge security).
    b. Low Latency Routing: to optimize traffic routing for edge applications.

7. NaaS Service Building Blocks

A major differentiator between NaaS offerings and traditional connectivity offerings is the customer's ability to choose components or building blocks for their NaaS service to be assembled by the NaaS provider into a complete, working NaaS offering (the "Granular Ordering" described in the NaaS Essential Features section). In other words, customers will expect to be able to identify and select clearly defined and priced NaaS building blocks. The following are examples of such building blocks:

1. Underlay Service Components: Foundational network infrastructure with on-demand bandwidth that supports NaaS. Examples include:
    a. Carrier Ethernet
    b. MPLS, L3VPN
    c. Layer 1 Optical Transport Network (OTN), WDM/DWDM
    d. DIA (Direct Internet Access)
    e. Broadband Internet
    f. Wireless connectivity
2. Overlay Service Components: Built on top of the underlay infrastructure to provide additional functionality and flexibility beyond on-demand bandwidth, such as subscription-based payment. Examples include:
    a. SD-WAN (Software-Defined WAN) for dynamic application policy-based configuration of network resources.
    b. Network Function Virtualization (NFV) technologies for dynamic deployment of network functions.
    c. Network and Cloud Security, such as SSE (Secure Service Edge) and ZT (Zero Trust).
3. Infrastructure Service Components: Supplementing network connectivity incurring recurring charges. Examples include:
    a. Hardware
    b. Virtual instances
    c. Managed services

8. NaaS Ecosystem

The NaaS ecosystem transcends traditional communication services providers (CSPs) and brings together diverse partners to revolutionize enterprise networking. By collaborating effectively, these ecosystem players can create substantial value and deliver an exceptional end-to-end user experience (UX), as well as support enterprise customers looking to host and access SaaS (Software-as-a-Service) applications on hybrid cloud infrastructures.

Scale and geographical reach are key to a NaaS ecosystem, as a larger ecosystem delivers greater choice and unlocks more interconnection opportunities for end customers. An ecosystem that offers breadth and depth of services, as well as access to an existing customer base, is also a more attractive proposition for new partners.

The NaaS ecosystem comprises these key players, as described in more detail below:

- Communication Service Providers (CSPs)
- Internet Exchanges (IXPs)
- Datacenter Operators
- Hyperscaler Cloud Operators
- SaaS Application Providers
- Edge & 5G Service Providers
- IoT Solution Providers
- Security Solution Providers
- Network Equipment Manufacturers
- API Gateway Providers

Communication Service Providers (CSPs)

CSPs serve as the backbone of the NaaS ecosystem; their extensive network infrastructure, including fiber-optic cables, data centers, and connectivity solutions, forms the foundation for reaching other ecosystem players. Key contributions from CSPs include high-speed private connectivity, automated service orchestration, and security services.

As CSPs invest further in automation and their own NaaS platforms, the need to standardize APIs to improve interoperability between CSPs and their network infrastructure increases.

Internet Exchanges (IXPs)

IXPs act as meeting points for networks, enabling direct peering; their role in the NaaS ecosystem includes reducing latency, interconnecting CSPs, and scalability for accommodating increasing traffic demands.

NaaS opens greater opportunity for remote peering, enabling enterprise, service provider, and carrier customers to interconnect with more global IXPs without having to establish a physical presence at the exchange point.

Datacenter Operators

Datacenters are central to a NaaS ecosystem and NaaS service delivery; their contributions include hosting both physical and virtualized network infrastructure for compute, storage, and networking resources, while edge datacenters play an important role in extending the reach of NaaS services closer to end-users.

NaaS facilitates datacenter interconnection, which is the private linking of two or more datacenters to shared resources; for enterprises and datacenter operators this advantageously opens access to a wider ecosystem, such as a cloud provider hosted in an adjacent datacenter facility.

Hyperscalers

Hyperscale cloud providers contribute significantly to the NaaS ecosystem by offering a global presence, ensuring NaaS service accessibility across regions; this presence provides enterprises with scalable infrastructure for computing, storage, and content delivery—and, more importantly, seamless integration with NaaS platforms through APIs.

Increasingly, enterprises are looking for ways to privately connect their infrastructure to the cloud. In the past, their choice was limited to purchasing dedicated connections from the cloud provider—e.g. AWS Direct Connect, Google Partner Interconnect, or Microsoft Azure ExpressRoute—or from one of the cloud provider’s preferred network partners, and be locked into fixed terms and bandwidth. Today, NaaS enables enterprises to establish virtual connections on-demand from one entry point to a service fabric and onto one or more cloud providers, leading to advantages that include reduced network complexity, increased network scalability, and improved control over cloud costs. With the rise of hybrid and multi-cloud strategies, many NaaS offerings also include cloud-to-cloud connections.

SaaS Providers

Software-as-a-Service (SaaS) providers rely on robust network connectivity. Their collaboration within the NaaS ecosystem results in optimized delivery that ensures low-latency access to their applications, enhancing user productivity. NaaS also enables secure connections to SaaS platforms, thereby protecting sensitive data; API-Driven Integration with NaaS services enables programmatic interactions with critical network resources.

As applications become more mission-critical to enterprises, demand for private access to SaaS applications continues to grow; options to integrate with a SaaS provider’s network architecture are more limited at this time than the cloud provider offerings, presenting a strong market opportunity for NaaS. NaaS facilitates easier interconnection between an enterprise network and the SaaS provider’s network architecture, making it a good fit for applications that require enhanced security and/or latency. In addition, some NaaS offerings offer the ability for customers to choose their class of service so that they can prioritize traffic for specific applications.

Edge & 5G Service Providers

Edge service providers help extend the NaaS experience to the network edge and, hence, closer to end-users, minimizing latency; in a global NaaS ecosystem where enterprises operate across regions, low-latency connections are essential for real-time applications, IoT devices, and remote workers.

The rise of Generative AI is also expected to drive new demand for edge services as AI data and processing moves away from centralized datacenters and much closer to devices, reducing latency and ensuring sensitive data doesn’t leave the edge.

IoT Solution Providers

IoT solution providers offer services, products, and expertise to help businesses and individuals implement, manage, and optimize Internet of Things (IoT) systems. These providers deliver end-to-end solutions, which typically include hardware, software, connectivity, data analytics, and security components, tailored to meet specific needs in a NaaS approach.

Security Solution Providers

Security is paramount in the NaaS ecosystem. Security solution providers enable data protection for sensitive data transmission and integrate services, such as encryption, firewalls, and intrusion detection, directly into the network fabric. Enterprises trust NaaS when robust security measures are in place.

Network Equipment Manufacturers

Network equipment manufacturers contribute significantly with the supply of routers, switches, SD-WAN controllers, and other critical hardware to ensure seamless end-to-end NaaS service delivery. This equipment also includes critical edge devices like 5G, IoT, and mobile devices.

API Gateway Providers

APIs are the lifeblood of NaaS; API gateway providers offer API tools to manage, secure, and monitor APIs. Well-defined APIs enable programmability, dynamic service provisioning, and billing. APIs interconnect various ecosystem components; they allow seamless interactions between CSPs, data centers, and SaaS providers. A well-designed API gateway enhances the developer experience, encouraging ecosystem partners to build innovative NaaS solutions.

9. Conclusion

Service providers have increasing access to the capabilities required for delivering cloud-like network services and, in parallel their enterprise customers, are increasingly motivated to consume network-based services in the same way that they consume compute and storage services in the cloud today. The alignment of expectations and a common language between the users of Network-as-a-Service on the one hand, and their providers on the other, is essential for the growth of the digital economy. This white paper delineates a baseline from which users and providers can quickly and effectively agree on expectations and requirements for the evolution of the NaaS industry.

10. About MEF

MEF is a global consortium of service, cloud, cybersecurity, and technology providers collaborating to accelerate enterprise digital transformation. It delivers standards-based frameworks, services, technologies, APIs, and certification programs to enable Network-as-a-Service (NaaS) across an automated ecosystem. MEF is the defining authority for certified Lifecycle Service Orchestration (LSO) business and operational APIs and Carrier Ethernet, SASE, SD-WAN, Zero Trust, and Security Service Edge (SSE) technologies and services. MEF's Global NaaS Event (GNE) convenes industry leaders building and delivering the next generation of NaaS solutions. For more information about MEF, visit MEF.net and follow us on LinkedIn and Twitter.

11. Acknowledgements

Alex Hawkes (Console Connect) – Contributor

Antonella Sanguineti (TI Sparkle) – Contributor

Daniel Bar-Lev (MEF Forum) – Co-Editor

Divesh Gupta (Console Connect) – Co-Author

Dominik Pacewicz (Amartus) – Contributor

Ettore Pulieri (TI Sparkle) – Co-Author

Fahim Sabir (Colt) – Co-Author

Federica Maria Manini (TI Sparkle) – Contributor

Isabella Melli (TI Sparkle) – Contributor

Jack Pugaczewski (Lumen) – Contributor

Javier Lecanda (Colt) – Co-Author

Karthik Sethuraman (Tata Communications) – Co-Author and Co-Editor

Mehmet Toy (Verizon) – Contributor

NaaS Customer Experience White Paper