MEF

MEF 3.0 PoC (138) White Paper

Extreme Services over End-to-End Multi-Layer Slices with LSO

Feb 2021

MEF-POC-138-WP2/2021

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 Abstract

MEF 3.0 PoC (138) provides a detailed demonstration of how 5G, network slicing and LSO orchestration can be used together by a service provider to support high-value real-time image analysis, RAN mobile transport, Time-Sensitive Network (TSN) computing and disaggregated computing. NTT and Okinawa Open Lab list the components used and their purpose, including MLTPE (Multi-Layer Transport Provisioning Engine), ACPE (Automated CNF Provisioning Engine), Slice Gateway, HANMOC (High Accuracy Network Monitoring and Control) and OpenNaEF orchestrator. This White Paper elaborates on the demonstration and explains the benefits for Communications Service Providers and enterprise customers.

2 Introduction

A private or non-public 5G network is a dedicated Local Area Network (LAN) that delivers enhanced network connectivity to industries, enterprises, and other customers. The evolving next-generation network services are expected to address the need for critical wireless communication for industrial operations, public safety, and critical infrastructure connectivity. Moreover, the global market is primarily driven by the growing need for ultra-reliable, low-latency connectivity for Industrial Internet of Things (IIoT) applications, including collaborative robots, industrial cameras, and industrial sensors.

Private 5G networks play a vital role in providing seamless and secure Internet connectivity to the above-mentioned IIoT devices. Moreover, the demand for Time-Sensitive Networking (TSN) and real-time-based networking is rapidly increasing for several mission-critical applications across many industries, including manufacturing, oil and gas, aerospace, and transportation. Thus, to provide high-speed bandwidth connectivity during TSN and real-time networking across industries, the demand for private 5G networks is expected to grow significantly over the forecast period.

Figure 1 shows the share and growth rate for service providers from 5G business-to-business (B2B) opportunities and shows that most industries are expected to have larger markets in 2030. Healthcare and manufacturing are the two biggest segments, and are expected to be significantly larger in 2030.

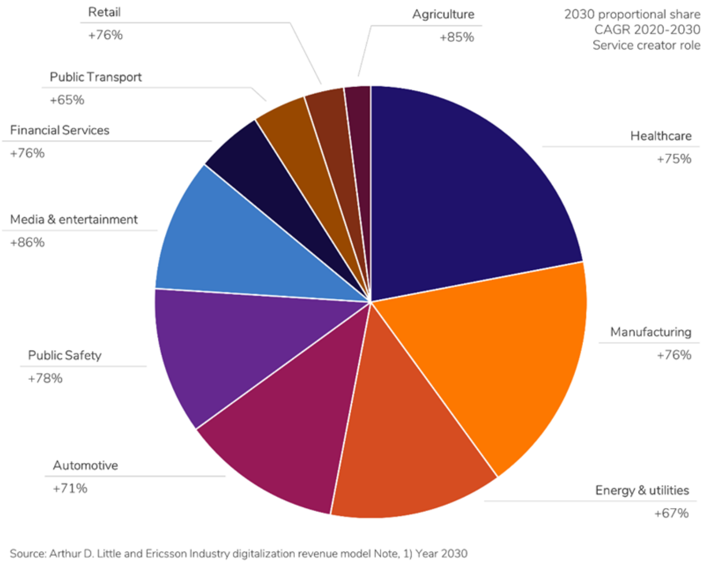

Figure 1: Share and growth rate from global total 5G-enabled B2B opportunities for service providers.[1]

Enterprise users are looking for an end-to-end (E2E) experience with guaranteed performance from their devices or machines to their applications/services, regardless of the user's location. Network slicing is fast becoming a hot topic in the context of 5G, as it provides the ability to allocate wireless performance resources for a given subscriber's slice from the Radio Access Network (RAN) side to 5G cores (5GC). While 5G promises many advanced capabilities, it provides only wireless access and mobile technology. Without E2E slicing, we will not be able to effectively manage 5G-enabled services from the RAN to 5GC and many existing wired network infrastructures.

3 Challenges in Delivering New Network-Slicing Services

Conventional per-domain individual configuration and provisioning of network slices is highly complex and time-consuming, requiring manual network provisioning of both 5G and WAN domains by specialist engineers. Per-domain provisioning and control of network slices have been addressed in some SDOs and communities, such as TIP OOPT CANDI, as well as MEF LSO Legato/Presto APIs.[2] Furthermore, owing to higher performance requirements than conventional VPN and Ethernet services in terms of latency and jitter of mission-critical applications, this complexity increases exponentially to assure network-slice service performance resulting in the following challenges and requirements:

- Per-domain provisioning of network slicing in private 5G and WAN increases provisioning time and lack of service flexibility, such as optimization.

[1] Source: Ericsson research https://www.ericsson.com/en/5g/5g-for-business/5g-for-business-a-2030-market-compass

[2] The typical development work covers the inventory for overlay services in conjunction with underlay resources, functional architecture of overlay, and underlay control and management, and coordination mechanisms between packet and optical transport networks.

- Per-domain slicing of private 5G and WAN makes traffic management in the E2E network slice difficult owing to the lack of identification of network slices.
- Manual service provisioning and configuration makes ultra-high performance difficult and time-consuming to assure and scale.
- Network provisioning for different performance requirements requires manual verification per application or customer in advance to assure the required performance.
- Mission-critical applications' performance requirements become much more demanding, especially in terms of latency and jitter, and service performance needs to be monitored more accurately than ever.

4 Goals of the Proof of Concept

To address these challenges, the MEF 3.0 Proof of Concept (PoC) (138) has been established to demonstrate the next level of innovation in network slicing over private 5G and WAN with the following goals:

1. Reduce the complexity of offering an E2E service over multiple network slices by:

- Automating service provisioning, lifecycle management, network-slice provisioning, management and E2E assurance across private 5G domains and transport network domains.
- Showing how network-slice performance can be measured and monitored with an E2E view.
- Realizing network interconnections between operators involved in a 5G service by using the MEF LSO APIs.

2. Advance MEF Lifecycle Service Orchestration (LSO) APIs and service definitions to expand network slice applicability:

- MEF 3.0 PoC (138) uses established MEF LSO APIs, including LSO Sonata APIs for inter-service provider service/slice ordering, as well as LSO Presto APIs between Service Orchestration Functionality (SOF) and Infrastructure Control and Management (ICM) for transport and 5GC. The aim is to provide MEF with valuable feedback to advance and mature the LSO Presto APIs to enable network slicing.
- The PoC also demonstrates the network service over an E2E network slice and its performance assurance to support mission-critical use cases. The aim is to show enhancements of MEF service definitions for network slicing.

5 Participants

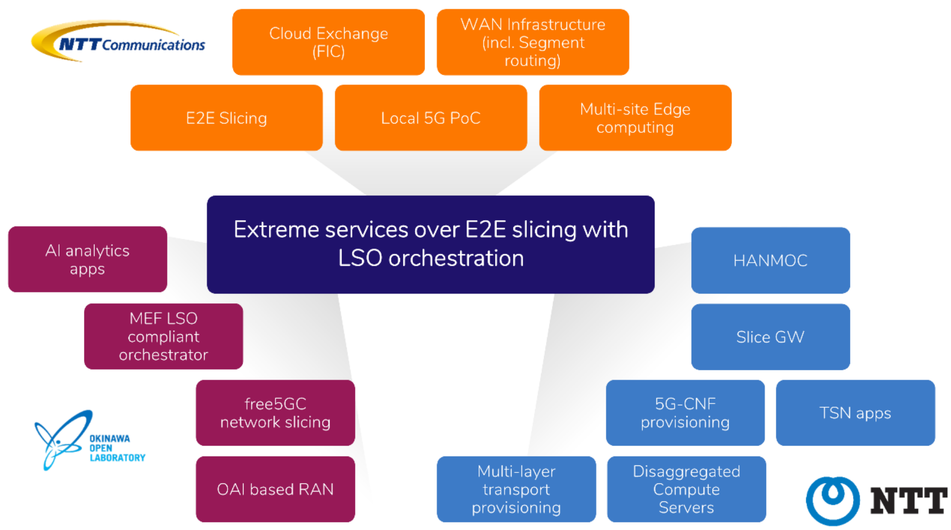

Figure 2: Participants and contributed components.

COMPANY	ROLE	PRODUCT NAMES
NTT Communications	• Private 5G and E2E network-slicing integrated solution service, including AI analytics apps. • Conducting multi-site edge computing PoC, E2E slicing over WAN with segment routing. These components are also proven with these activities.	Flexible Interconnect
NTT	• PoC lead, including implementation of architecture design, integration, and test. • Transport and container infrastructure controller. • Per slice traffic isolation and policing. • Extreme service applications.	Multi-layer transport and 5G slicing provisioning engine (E2E slice controller) Slice gateways HANMOC
Okinawa Open Laboratory	• OSS based 5G core slicing • Service orchestrator • AI analytics apps	

Table 1: Participants and Products

6 PoC Overview

MEF 3.0 PoC (138) demonstrates the next level of innovation and directions in E2E network slicing, which can accelerate implementation of an orchestrated and assured multi-domain network slice and extreme services through automated resource provisioning and network activation. Powered by a service orchestrator, together with a multi-layer transport and 5G slice controller, Slice Gateways (SLG), and HANMOC (High-Accuracy Network Monitoring and Control), E2E network slicing across different domains can be instantiated to enable a high-performance network service for mission-critical applications. Functional, specific, dedicated networks over multi-layer and the multi-domain network are orchestrated, activated, assured, and managed using simple Service Level Agreement (SLA)-based service ordering and control using service automation in alignment with MEF's LSO framework.

7 PoC Demonstration

The PoC shows the journey of a business customer that wants to cover its local site with private 5G and connect to the edge cloud through WAN for its mission-critical applications. 5G services that meet the various performance requirements provided by service providers are realized by the interconnection of networks between the several operators that make up the service. It is expected that private or non-public 5G providers have to collaborate with various operators to implement services, so we consider this to be a good use case for service configuration using MEF LSO APIs.

The demonstration scenario starts with the ordering of a 5GC slice from the self-service portal and its connection to the customer's private 5G base stations. The 5G slices are dynamically designed, optimized, and instantiated, based on the customer's Service Level Specification (SLS), using an automated Cloud-Native Network Function (CNF) provisioning engine. The information on the 5GC slices, subscriber information, and data-network-address range and allowed destinations are configured into 5GC containers.

The demonstration shows transport slice instantiation and activation after the inter-provider service ordering between 5GC and transport service providers via an LSO Sonata reference point. The transport network slices using multiple underlay connectivity such as L1 optical transport, Carrier Ethernet, and IP services are created in accordance with the SLS information in the LSO Sonata ordering process.

E2E slice connecting terminals to edge/central clouds are continuously monitored and visualized in real-time, including connecting sites, bandwidth utilization, latency, and jitter per slice.

In the demonstration, typical applications that enterprises are expecting over E2E slices over private 5G networks are shown to prove that mission-critical applications work in the demo system over the implemented slices. The mission-critical applications being demonstrated include:

7.1 Real-Time Image Analysis

The example requirements of real-time Artificial Intelligence (AI)/Machine Learning (ML)-based image analysis applications using 4K or 8K video stream transmission over a network can be more than 200 Mbps per stream and its latency should be in the range of 10-100 ms to enable immediate feedback to the camera or send alerts on the basis of the resulting analytics.

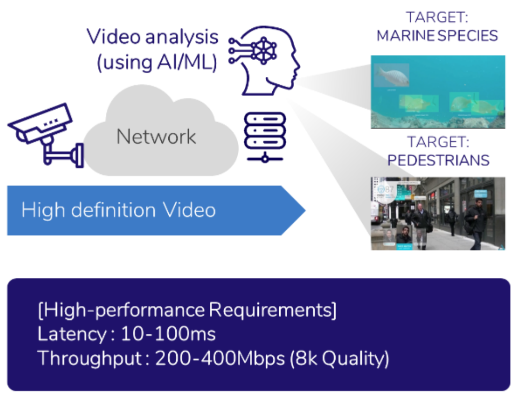

Figure 3: Real-time video analysis using AI/ML.

7.2 RAN Mobile Transport (xHaul)

Performance requirements for RAN are very strict: latency should be less than 250 µs in fronthaul, and throughput is more than 22.2 Gbps per Radio Unit (RU) for wireless frequency-synchronized control and eCPRI protocol characteristics. For mobile midhaul, the latency should be less than 4 ms, and the throughput requirement is more than 4.02 Gbps per RU. For mobile backhaul, latency should be less than 10 ms, and the throughput requirement is more than 4 Gbps/cell per RU. [2]

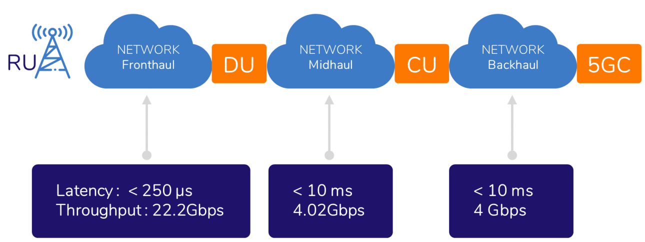

RU: Radio Unit; DU: Distributed Unit; CU: Central Unit; 5GC: 5G Core network

Figure 4: Mobile xHaul (transport).

7.3 Time-Sensitive Network (TSN) Communication

To remotely enable precise robot or Automated Guided Vehicle (AGV) remote motion control in a smart factory via a private 5G network, TSN provides exact packet transmission scheduling at µs granularity in the local network, so that normal delay variations are in the order of nanoseconds. By using network slices, TSN communications are enabled at the same delay variation as the local network, even if it is through a WAN. [3][4]

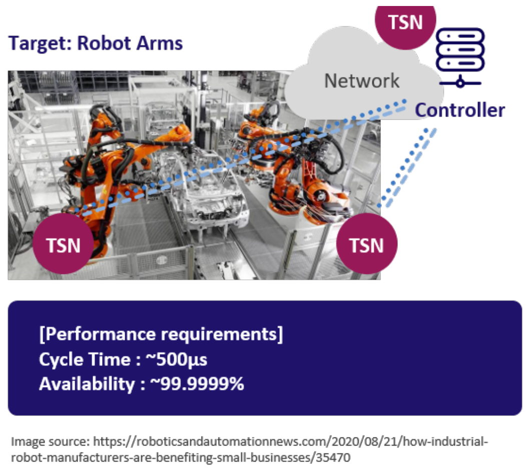

Figure 5: Remote motion control with high accuracy and low jitter.

7.4 Disaggregated Computing

As disaggregation of server hardware in the data center increases, server hardware like CPU, GPU, Storage, and NIC will be connected over the network using technologies such as Remote Direct Memory Access (RDMA)—equivalent to a computer internal bus over the network, such as PCI-express. RDMA can transfer data from the memory of a local computer to the memory of a remote computer. An example of network performance requirements for disaggregated computing is more than 4 GB/s throughput for data transfers between memory and storage, with time synchronicity of be less than 10 µs for parallel processing between CPUs.

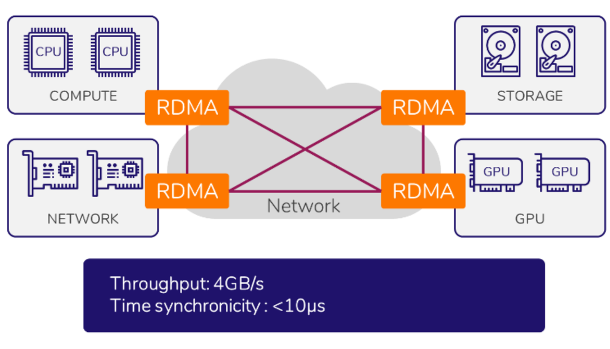

Figure 6: Remote Direct Memory Access (RDMA).

As described above, these performance characteristics need E2E network slicing to assure performance across the 5G network domain and transport service domain. The key technological innovations enabling E2E network slicing that are being demonstrated include:

- Segment Routing (SR)-based multi-layer path selection and connectivity control, based on multi-layer transport-resource management and automated optimal resource allocation for high network performance assurance.
- 5G container resource management and automated optimal resource allocation for high network performance assurance.
- Per-domain network slice selection and stitching based on E2E slice identification.
- High-accuracy network performance (latency and jitter) monitoring in μs and automated and immediate SR-path optimization.
- Demonstration of mission-critical applications working over E2E network slicing from AI-video analytics, TSN communication, mobile xHaul, and distributed computing.

8. PoC Architecture

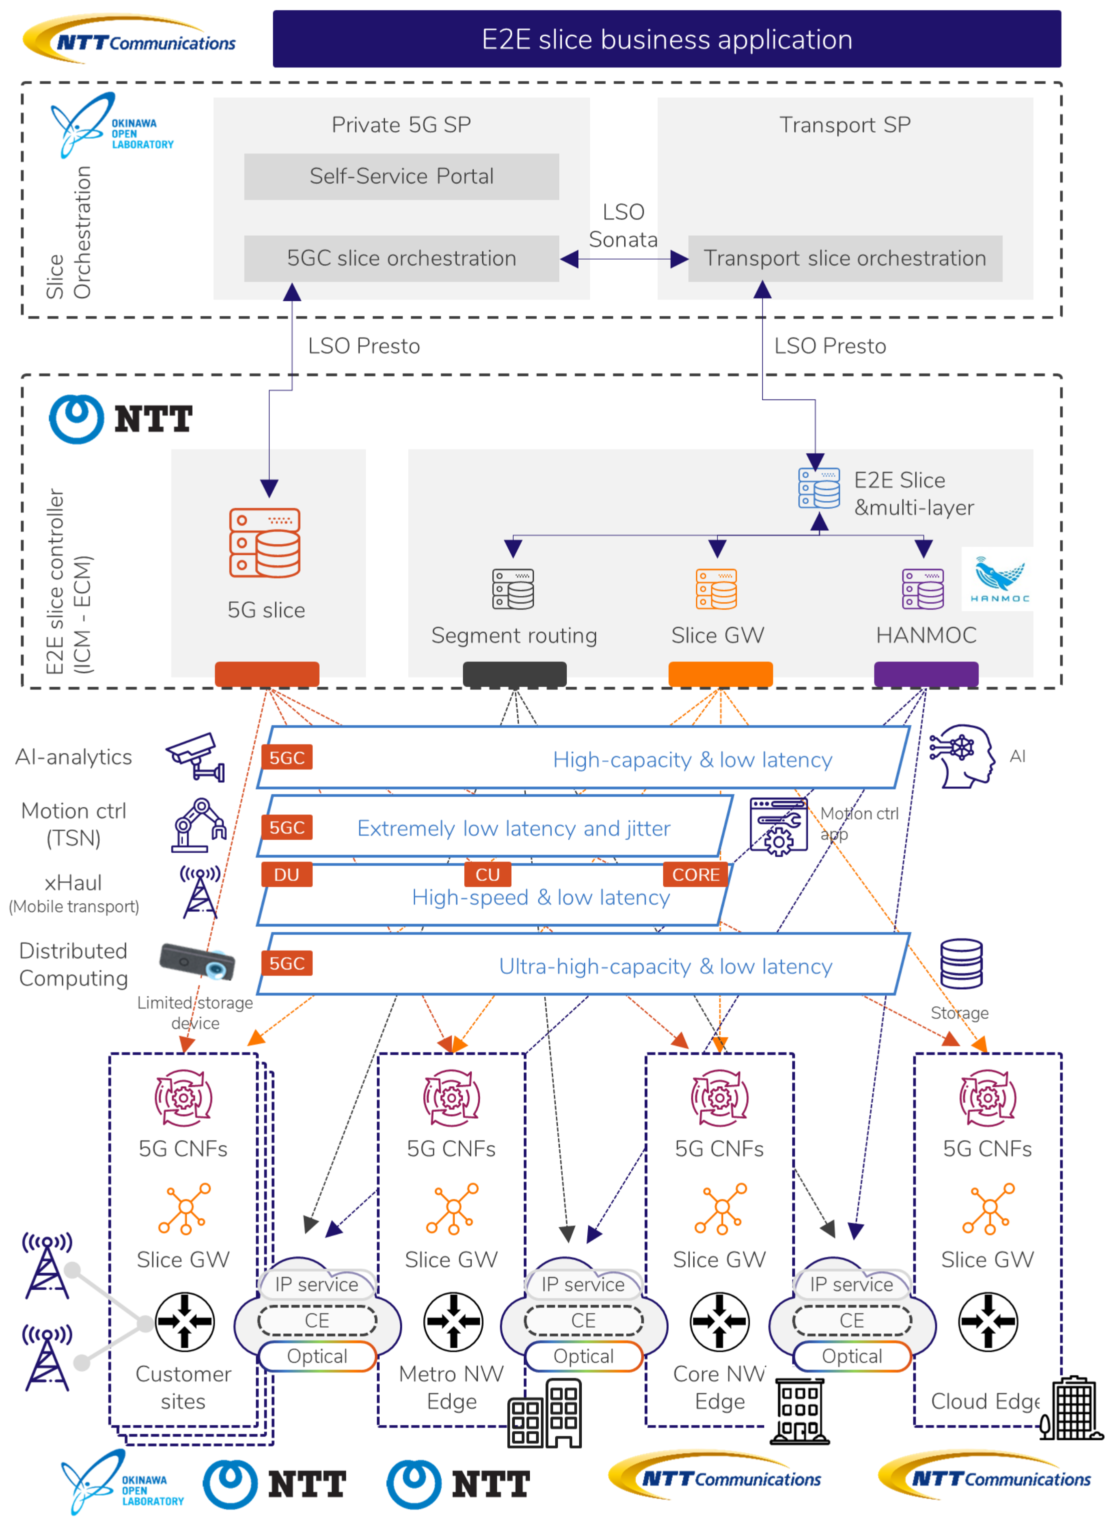

Figure 7: PoC architecture.

9 Benefits

For Communication Service Providers (CSPs):

- Enables CSPs to offer and manage services using E2E network slice with high flexibility.
- Enables accurate network service-performance monitoring.
- Lowers operation costs by streamlining deployment through systematic resource provisioning, performance validation, and testing through automated resource provisioning.

For enterprise customers:

- Complete visibility and control of services procured through a CSP as a managed service.
- Confidence that services and changes are validated before becoming operationalized especially in terms of mission-critical performance.
- Ability to set up or change network slices at any time without incurring wait time for validation and testing at the CSP.

10 Supported MEF Standards

10.1 MEF 3.0 LSO Architecture

The system design and implementation conform to the MEF 3.0 LSO Reference Architecture Framework as follows:

Service Orchestration Functionality (SOF)

- Open Okinawa Labs' (OOL) Open Network and Engineering Framework (NaEF) enables on-demand and customer-initiated network inter-provider E2E network-slice service orders.

Infrastructure Control & Management (ICM)

- This PoC involves multiple ICMs for different domains or Service Providers (SPs) of an E2E slice, including 5G and multi-layer transport controllers from NTT, such as Multi-Layer Transport Provisioning Engine (MLTPE) for layer 1 to 3 transport, Automated CNF Provisioning Engine (ACPE) for 5G cores, and High Accuracy Network Monitoring and Control (HANMOC).

10.2 MEF 3.0 Services

To guarantee the ultra-high-performance and automated service provisioning over multi-domain network slices, this PoC leverages and demonstrates MEF underlay connectivity service specifications, and an early implementation of the MEF 84 Network Slicing draft specification.

10.3 MEF 3.0 LSO APIs

LSO Presto APIs: This PoC utilizes MEF API standards and SDKs, as well as facilitating the pre-standard work of the MEF Network Slicing projects and Slice Profile ad-hoc currently underway in MEF.

11 Overview of PoC Components

11.1 Multi-Layer Transport Provisioning Engine (MLTPE)

MLTPE enables optimal transport-resource provisioning for both packet and optical networks, in accordance with an SLA. MLTPE is implemented to work with the ONF TAPI Reference Implementation and MEF LSO Presto SDK to control connectivity services through Element Control and Management (ECM), such as packet domain controllers and optical domain controllers.

The MLTPE is utilized to calculate the optimal resources and paths for transport slices to meet the ultra-high-performance requirements of mission-critical applications.

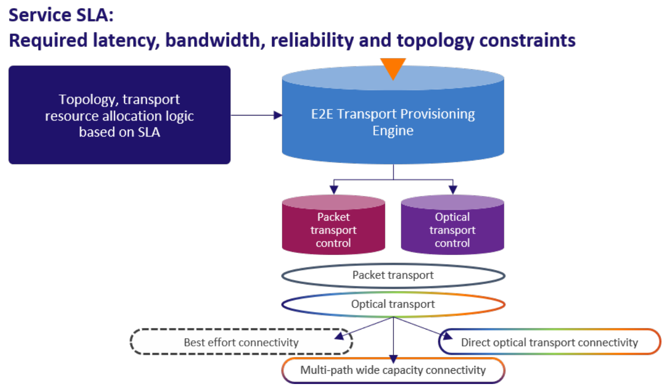

Figure 8: MLTPE

11.2 Automated CNF Provisioning Engine (ACPE)

ACPE enables automated optimal container resource allocation for CNFs. It designs the optimal allocation of server resources for a specific SLS. The optimized set of virtualized resources, such as CPUs, memory, NIC, and GPUs, are allocated to assure the ultra-high performance for 5G network slices. The performance for CNFs is predicted by an ML-based model and performs optimal resource calculations and continuously updates itself in line with traffic changes. It also supports the multi-cluster deployment of CNFs for low-latency and reliability, based the calculations of the ACPE. ACPE works with many container control-and-operation software packages, such as Helm and Kubernetes operator for instantiation, modification, scaling, and deletion of 5G network slices.

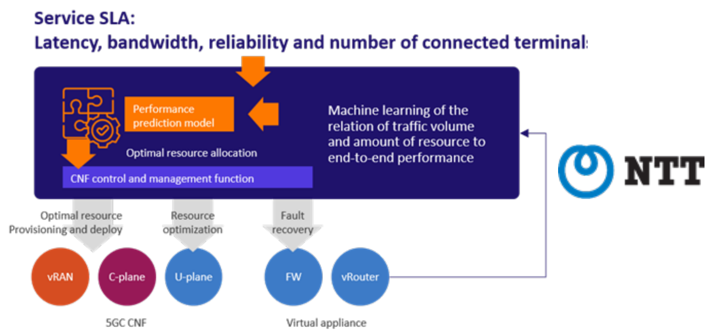

Figure 9: ACPE

11.3 Slice Gateway (SLG)

To enable E2E network slices across multiple domains, network resources of each domain need to be mapped to the slice as per its requirements. Here, a network resource is a logical path that has a specific forwarding performance, configured by SR, using forwarding routes and methods (SD-WAN overlay, L2 path, optical direct path) that each domain can provide. SLG[5] provides the functionality to map resources and the slices at the edge points of each domain. The SLG has the following functions:

1. Identifying the slice to which each packet belongs at 5-tuple granularity.
2. Selecting a logical path (SR-MPLS path in this PoC) that meets the requirements of the identified slice.
3. Policy control, based on the requirements of each slice, such as QoS control and rate limit in SLG.
4. SLG Controller provides APIs to MLTPE, or an orchestrator, to automatically create slices triggered by customer orders.

SLGs and an SLG Controller are currently installed using open source software (OSS) and trials are underway.

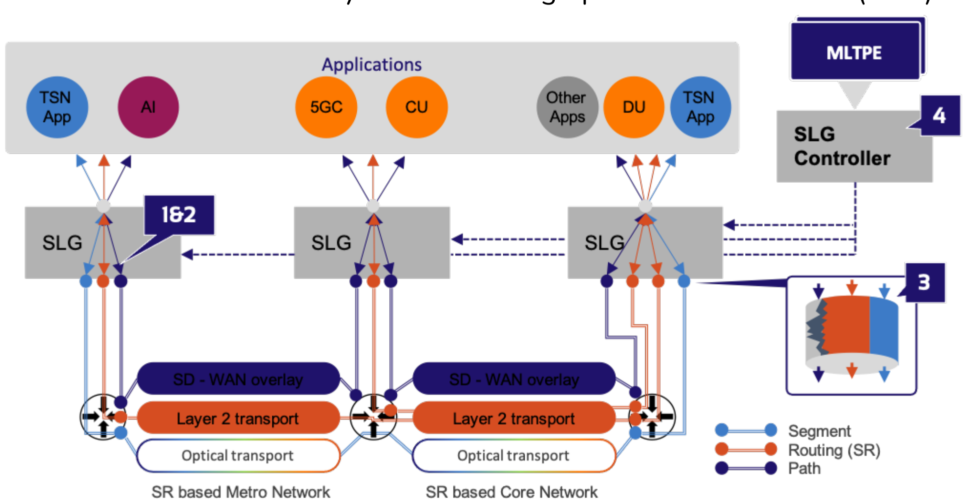

Figure 10: Slice Gateway

11.4 HANMOC (High Accuracy Network Monitoring and Control)

HANMOC[6] measures transport delay and jitter in the network with high accuracy by utilizing SR technology and provides flexible path control on the basis of its results.

Features:

1. Lowers CAPEX/OPEX by using probe-loopback method (adding on just one system to the target network).
2. Rapid identification of delay section and direction by proprietary algorithm.
3. Adds timestamps for delay/jitter measurement to packets with high accuracy using generic hardware. (Possible to monitor in high accuracy in μs units.)
4. Applicable if the network supports SR.

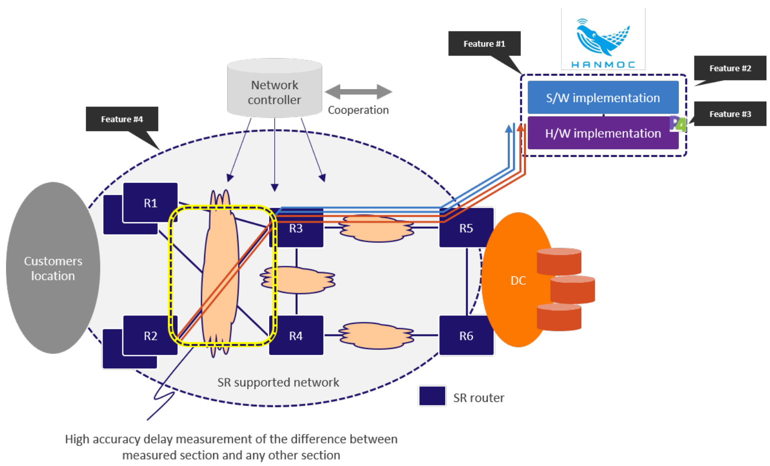

Figure 11: HANMOC

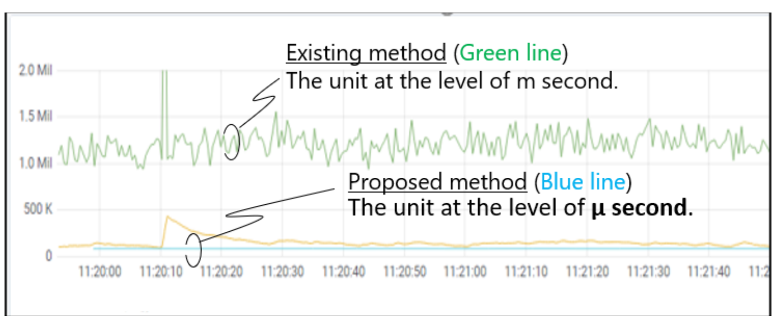

Figure 12: Results of the comparison between H/W processing and S/W processing.

11.5 OpenNaEF Network Topology Lifecycle Data-Store Orchestrator

OpenNaEF is a framework that implements lifecycle orchestration based on an on-memory graph DB—a network topology lifecycle data-store. OpenNaEF has the following features: Flexible and scalable object model support, on-memory DB, Schedule Management, History Management.

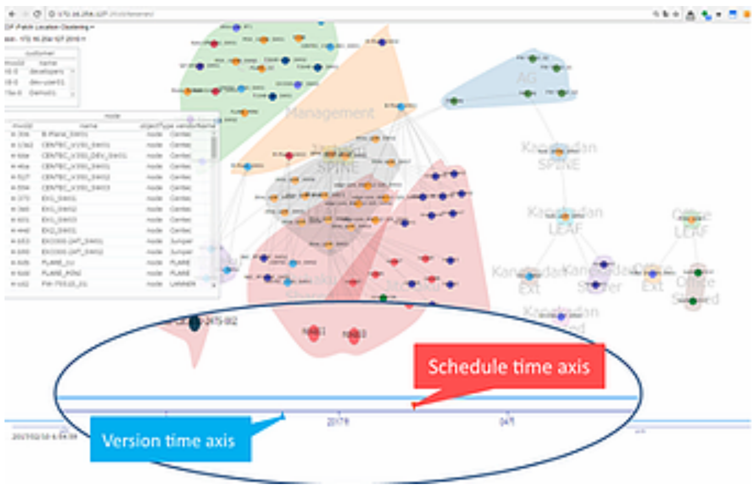

Figure 13: OpenNaEF portal example.

1. It has been used in various systems for over 20 years in many use cases.
2. Multi-vendor, multi-protocol and multi-service support.
3. All network model objects have a two-dimensional time axis and can manage the lifecycle of the network topology historically and predictively.
4. OpenNaEF natively possesses conventional network model configuration elements. By combining these configuration elements, almost all network topologies can be modelled.
5. In addition to network configuration operations that are reflected immediately like SDN, all model information can be converted to settings for network devices, servers, etc., and settings can be set for each job.

12 Summary

In this PoC, MEF digital service architecture and LSO are combined and extended to demonstrate the enterprise-dedicated network-slice service over private 5G, metro, and core transport to edge/cloud infrastructure. Extreme services, such as low latency real-time video analysis, jitter-less robotic communication, ultra-capacity storage data transfer, and xHaul transport are enabled over the end-to-end slices with LSO orchestration across multi-layer underlay slices and 5GC slices.

Service providers can deliver high-margin services using a combination of AI/ML in the Edge, 5G, and network slicing in a way not possible until now. 5G/private 5G is accelerating automation and digital transformation in various industries. This PoC demonstrates E2E-dedicated slice instantiation between 5G and transport SPs, and also how it helps enterprise customers' business in terms of service order and management.

Standardized service models for E2E-dedicated network service, including multi-layer underlay slices, as well as interconnection and performance implementation agreements, in addition to standards governing interoperability between metro and core transport slices, and container-based 5GC slices are required. The lack of an organizing feature for E2E slice interconnection, control, and management causes individual provisioning and monitoring of service providers, leading to frequent, complex inter-operator operational procedures, especially for SLS monitoring.

13 About MEF

An industry association of 200+ member companies, MEF has introduced the MEF 3.0 transformational global services framework for defining, delivering, and certifying assured services orchestrated across a global ecosystem of automated networks. MEF 3.0 services are designed to provide an on-demand, cloud-centric experience with user- and application-directed control over network resources and service capabilities. MEF 3.0 services are delivered over automated, virtualized, and interconnected networks powered by LSO, SDN, and NFV. MEF produces service specifications, LSO frameworks, open LSO APIs, software-driven reference implementations, and certification programs. MEF 3.0 work will enable automated delivery of standardized Carrier Ethernet, Optical Transport, IP, SD-WAN, Security-as-a-Service, and other Layer 4-7 services across multiple provider networks. For more information, visit https://www.mef.net/ and follow us on LinkedIn and Twitter @MEF_Forum.

14 References

[1] *Ericsson research, 5G for business a 2030 market compass*, https://www.ericsson.com/en/5g/5g-for-business/5g-for-business-a-2030-market-compass

[2] *Cisco whitepaper, Reimagining the End-to-End Mobile Network in the 5G Era* https://www.cisco.com/c/dam/en/us/products/collateral/cloud-systems-management/elastic-services-controller-esc/reimagining-mobile-network-white-paper.pdf

[3] *Ericsson Technology Review, 5G TSN - integrating for industrial automation – Ericsson, Aug. 2019*, https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/5g-tsn-integration-for-industrial-automation

[4] *5G-ACIA, 5G for Connected Industries and Automation, Feb. 2019*, https://www.5g-acia.org/publications/5g-for-connected-industries-and-automation-white-paper/

[5] *NTT Technical Review, Service-partner-oriented Network Slicing, Jun. 2019*, https://www.ntt-review.jp/archive/ntttechnical.php?contents=ntr201906fa5.pdf&mode=show_pdf

[6] *NTT R&D, "HANMOC"*, https://www.rd.ntt/e/research/NS0030.html

MEF

MEF-POC-138-WP2/2021