5G European Validation platform for Extensive trials

Deliverable D3.2
Interworking Reference Model

Project Details

Call	H2020-ICT-17-2018
Type of Action	RIA
Project start date	01/07/2018
Duration	36 months
GA No	815074

Deliverable Details

Deliverable WP:	WP3
Deliverable Task:	Task T3.2
Deliverable Identifier:	5GEVE_D3.2
Deliverable Title:	Interworking reference model
Editor(s):	Lourdes Luque Cantó
Author(s):	Lourdes Luque, Luis Miguel Contreras, Juan Rodríguez (TID); Giacomo Bernini (NXW); Marc Molla (ERI-ES); Cao Phan, Farouk Messaoudi, Rabah Guedrez, Sergio Morant (B-COM); Fabrizio Moggio, Carmen Catalano (TIM); Nicola Blefari, Mauro Femminella, Francesco Lombardo, Paolo Lungaroni, Matteo Pergolesi, Gianluca Reali, Stefano Salsano (CNIT); Ramón Pérez, Aitor Zabala (TELC); Michal Chabiera, Grzegorz Panek (ORA-PL)
Reviewer(s):	Andreas Tzoulis; Luca Stroppolo
Contractual Date of Delivery:	30/06/2019
Submission Date:	28/06/2019
Dissemination Level:	PU
Status:	Final
Version:	v1.0
File Name:	5GEVE_D3.2_v1.0

Deliverable History

Version	Date	Modification	Modified by
ToC_v0.1	*04/03/2019*	*ToC proposal*	*Lourdes Luque Cantó*
ToC_v0.2	*11/03/2019*	*Minor assignment modifications*	*Lourdes Luque Cantó*
v0.1	*25/03/2019*	*Contributions sections 1&2*	*Lourdes Luque Cantó, based on contributions by B-COM, NXW, CNIT, TIM and TID*
v0.11	*01/04/2019*	*Format & review of 2.2.4*	*Juan Rodriguez Martinez*
v0.12	*11/04/2019*	*First contributions to 4.3*	*Giacomo Bernini, based on contributions by NXW, TELC and TID*
v0.13	*22/04/2019*	*Additional contributions to sections 1&2. First contributions to section 4.1*	*Lourdes Luque Cantó, based on contributions by TELC and TID*
v0.14	*30/04/2019*	*Section 5 draft*	*Sergio Morant (b-com)*
v0.15	*07/05/2019*	*Workflows and architecture drafts in section 4.1 & 4.2*	*Luis Miguel Contreras*
v0.16	*21/05/2019*	*Multiple contributions, mainly to section 4*	*Multiple partners*
v0.17	*27/05/2019*	*Additional contributions to section 4*	*TELC*
v0.18	*30/05/2019*	*Multiple contributions to multiple sections*	*Multiple partners*
v0.19	*04/06/2019*	*Contributions to section 4*	*Multiple partners*
v0.9	*05/06/2019*	*First stable version for internal review*	*Lourdes Luque Cantó*
v0.91	*11/06/2019*	*Implemented comments from first review*	*Lourdes Luque Cantó*
v0.92	*24/06/2019*	*Second stable version for final comments*	*Lourdes Luque Cantó*
v0.93	*26/06/2019*	*Finalized version*	*Lourdes Luque Cantó*
v1.0	*28/06/2019*	*Version ready for delivery*	*Lourdes Luque Cantó*

Table of Contents

List of Acronyms and Abbreviations

Acronym	*Meaning*
3GPP	3rd Generation Partnership Project
5G-PPP	5G Infrastructure Public Private Partnership
AAA	Authentication, Authorization and Accounting
AI	Artificial Intelligence
AMF	Access and Mobility Function
API	Application Programming Interface
BGP	Border Gateway Protocol
BSS	Billing/Business Support System
BSSO	Business Service Slice Orchestrator
CESC	Cloud-Enabled Small Cell
CMD	Command (prompt)
CN	Core Network
CPU	Central Processing Unit
C-RAN	Cloud RAN
CRUD	Create Read Update Delete
CS	Communication Service
CSAR	Cloud Service Archive
CSMF	Communication Service Management System
CTRL	Control
DevOps	Development and Operations
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Server
DSSO	Domain Specific Slice Orchestrator
E2E	End to End
EGMF	Exposure Governance Management Function
eMBB	Enhanced Mobile Broadband
EMS	Element Management System
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
EVE	Evolution and Ecosystem (in ETSI NFV-EVE)
E-VPN	Ethernet Virtual Private Network
FCAPS	Fault, Configuration, Accounting, Performance and Security
GRE	Generic Routing Encapsulation
GSM	Groupe Spéciale Mobile

GSMA	GSM Association
GST	Generic Slice Template
GUI	Graphical User Interface
I/W	Interworking
ICMP	Internet Control Message Protocol
ID	Identifier
IETF	Internet Engineering Task Force
IFA	Interfaces and Architecture
IMS	IP Multi-Service
IoT	Internet of Things
IP	Internet Protocol
IPSec	IP Security
ISRB	Inter-Slice Resource Broker
IT	Information Technology
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
L2	Layer 2
L3	Layer 3
LAN	Local Area Network
LCM	Life Cycle Management
LoI	Level of Inventory
MANO	Management and Orchestration
MAPE	Monitor Adapt Plan and Execute
MD-SAL	Model-Driven Service Abstraction Layer
MdO	Multi-domain Orchestrator
MDSO	Multi Domain Slice Orchestrator
MEAO	Mobile Edge Application Orchestration
MEC	Multi-access Edge Computing
MF	Management Function
MLPOC	Multiple Logical Point of Contact
mMTC	Massive Machine Type Communications
MnS	Management Services
MTP	Mobile Transport and Computing Platform
NBI	Northbound Interface
NEST	Network Slice Type
NETCONF	Network Configuration Protocol
NMR-O	Resource and multi Network segment Orchestrator

NF	Network Function
NFV	Network Functions Virtualization
NFVI	NFV Infrastructure
NFVIaaS	NFVI as a Service
NFVO	NFV Orchestrator
NRF	Network Repository Function
NS	Network Service
NSaaS	Network Service as a Service
NSD	NS Descriptor
NSI	Network Slice Instance
NSMF	Network Slice Management Function
NSO	Network Service Orchestrator
NSSI	Network Slice Subnet Instance
NSSMF	Network Slice Subnet Management Function
NST	Network Slide Template.
NTP	Network Time Protocol
O&M	Management and Orchestration
OAM	Operations, Administration and Maintenance
ODL	OpenDaylight
ONAP	Open Networking Automation Platform
ONF	Open Networking Foundation
ONOS	Open Network Operating System
OPNFV	Open Platform for NFV
OSM	Open Source MANO
OSS	Operations Support System
OVS	Open Virtual Switch
OVSDB	Open Virtual Switch Database
P&P	Plug and Play
PBM	Policy Based Management
PCI	Peripheral Component Interconnect
PCI-PT	PCI Pass-through
PDN	Packet Data Network
P-GW	PDN Gateway
PNF	Physical Network Function
PNFD	PNF Descriptor
PNFI	PNF Instance
PoP	Point of Presence

QoE	Quality of Experience
RAM	Random Access Memory
RAN	Radio Access Network
REST	Representational State Transfer
RFC	Request For Comments
RMON	Remote Monitoring
RO	Resource Orchestrator
SA	Service and System Aspect (3GPP terminology)
SBA	Service Based Architecture
SBI	Southbound Interface
SD	Service Descriptor
SDK	Software Development Kit
SDM	Software-Defined Mobile
SDMC	SDM Controller
SDMO	SDM Orchestrator
SDM-X	SDM Coordinator
SDN	Software Defined Networks
SDO	Standards Development Organization
S-GW	Serving Gateway
SIM	Subscriber Identity Module
SLA	Service Level Agreement
SlaaS	Slice as a Service
SLPOC	Single Logical Point of Contact
SMF	Session Management Function
SNMP	Simple Network Management Protocol
SO	Service Orchestrator
SOL	Solutions
SP	Service Platform
SPOF	Single Point of Failure
SQL	Structured Query Language
SR-IOV	Single Root Input/Output Virtualization
SSH	Secure Shell
SSL	Secure Sockets Layer
SS-O	Slice Service Orchestrator
ST	Service Template
SVP	Service Virtualization Platform
T-API	Transport API

TCP	Transport Control Protocol
TOSCA	Topology and Orchestration Specification for Cloud Applications
UDP	User Datagram Protocol
UE	User Equipment
uRLLC	Ultra Reliable Low Latency Communications
V&V	Verification & Validation
VIM	Virtualized Infrastructure Manager
VL	Virtual Link
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
VNFaaS	VNF as a Service
VNFD	VNF Descriptor
VNFI	VNF Instance
VNFFG	VNF Forwarding Graph
VNFM	VNF Manager
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VS	Vertical Slicer
VxLAN	Virtual Extensible Local Area Network
WAN	Wide Area Network
XML	Extensible Markup Language
YANG	Yet Another Next Generation
YAML	Yet Another Markup Language

List of Figures

List of Tables

Executive Summary

Experimenters from Verticals willing to execute an experiment on top of the 5G EVE infrastructure will be able to interact with 5G EVE Experimentation Portal in order to completely define their interests and requirements. However, it is not the Portal the component that directly interacts with the elements at the different sites to deploy such experiment, or modify the testing conditions after different repetitions, for example. That component is the 5G EVE Interworking (I/W) Layer, which provides the required abstraction to permit the Portal not to be aware of the different underlying technologies, and to permit the parallel evolution of the different sites without affecting our single and common frontend for Verticals.

The present D3.2 document, *Interworking Reference Model*, is the specification of the Interworking Framework for 5G EVE, further developing the initial specification provided in Deliverable D3.1 (*Interworking Capability Definition and Gap Analysis*).

The foundations of the architecture have not been changed. The proposed I/W Framework is composed of five modules, with the following summarized functionalities:

- The Multi-Site Network Service Orchestrator (NSO), Multi-site Catalogue and Multi-site Inventory are responsible for the management of the lifecycle of the deployed components, jointly allowing for multi-site slices supporting Verticals' experiments.
- The Data Collection Manager collects the required performance metrics to ensure the correct operation of the infrastructure and to validate the targeted KPIs.
- The Runtime Configurator applies the required runtime configurations to the provisioned services.

In this document, a detailed description of the features required to those five modules is provided, together with the specifications of north-bound and south-bound reference points allowing for the envisioned interactions. In that sense, the Interworking API or NBI (North-Bound Interface) allows for the communication with the Experimentation Portal, while the Adaptation Layer or SBI (South-Bound Interface) is the one abstracting the heterogeneous site capabilities. Both interfaces have been defined for the Data Collection Manager, the Runtime Configurator and the Multi-site NSO. However, the Multi-site Catalogue and Inventory will only interact with the sites using the Multi-site Orchestrator as a proxy. The reason is that we can leverage on features already supported by available orchestrators to simplify the Adaptation Layer.

A very important decision, valid for the first release of the I/W Framework, is that only the local NFVO components will be exposed to the Multi-site Orchestrator. This means that local orchestrators will maintain the direct control over the local infrastructure, with other local components like the VIM or the SDN controller not exposed directly to the upper layers. This approach can be revisited in future versions of the I/W Framework, depending on new interworking requirements. In any case, we have also accounted for the possibility of any site deploying an additional SDN controller, for example for the WAN, in which case it should be possible for the I/W Framework to interact directly with it via the WIM.

The biggest complexity, however, comes from the wide range of components which could be targeted by the Runtime Configurator and the Data Collection Manager. The amount of software drivers to be implemented in the I/W Layer will depend on the number of Vertical experiments, the complexity of those experiments in terms of number of required VNFs or PNFs, the tools that are available in the different sites, the vendors providing those tools, etc. Although some level of harmonization will be sought in the project as a whole, the I/W Layer will have to be deal with a very heterogeneous environment, therefore subject to interoperability issues.

From the design point of view, it is not realistic to try and define an SBI valid for all the uses cases. Therefore, for these components, we will follow an approach based on continuous integration of new features, starting from the ones required by the Verticals participating on 5G EVE, and building on top of those as new requirements appear.

Another design decision taken to minimize integrations is to rely on standards whenever these are available, even in a very preliminary stage. That should guarantee the design being future-proof, leveraging on subsequent industry implementations, and also opens the possibility to contribute on the standards that need extensions. For

that purpose, the work on D3.2 began with a State of the Art covering the main references regarding NFV, network slicing and SDN control. That study was complemented with a detailed analysis of other European projects dealing with concepts similar to those of 5G EVE. The main goal was to understand what results were reusable from these projects, and where we needed to put the main focus for new implementations.

The first conclusion was the confirmation that, until 5G EVE, there had not been an effort to build a system allowing for multi-site Vertical experimentation across such a complex and heterogeneous environment. Therefore, all the adoptions made in 5G EVE will have to be carefully assessed, as they most probably will not be valid straight away. The second conclusion was that if there is a component that has not been sufficiently tackled yet is precisely the Runtime Configurator, for the reasons mentioned above.

Apart from a technical description of the I/W Framework itself, the document also describes the services it will offer to the Verticals, in line with the ongoing work in WP1.

In the first stage, experiments will only be deployed in a single site, so the I/W Framework will be able to support single-site Applications Deployment, Experiment Monitoring and Network Automation, the latter to build the required connectivity services (e.g., a VLAN or a VxLAN tunnel) among the deployed components. In the second phase, the I/W Framework will be upgraded with additional capabilities, like the multi-site E2E Orchestration and extended Experiment Monitoring. The E2E Orchestration feature includes the capability to deploy multi-site slices, and VNFs on top of them, across all the sites participating in a given experiment.

Even with all these considerations, experiments will only be possible in 5G EVE with the deployment of an underlying connectivity supporting the identified requirements. The first part of the activity, the requirements definition, was already included in D3.1; in D3.2, we have complemented the analysis by studying all the technical possibilities to achieve this connectivity in the different communication Planes (orchestration, control and data). Although the final implementation decision corresponds to WP2, a proposal is outlined here aiming at securing compliance with the needs of the I/W Layer in the Orchestration Plane: a star-based topology, with the I/W Framework components located in the hub, placed in the Turin site to minimize the delay. The technical implementation would be based on IPSec VPN tunnels, and an addressing scheme should have to be agreed among all the sites to avoid duplications.

We have also defined additional services or functionalities required from the 5G EVE site facilities, some of them fundamental to ensure the correct operation, and others to provide added value to the Verticals. An example of the first ones is the adoption of appropriate security measures, even more importantly, taking into account the multi-site characteristic of the 5G EVE environment. Therefore, items like hardening of the infrastructure, authentication in the Orchestration Plane, firewalling, filtering, etc., are discussed. On the other hand, the necessity of services like the remote access for Vertical to their deployed VNFs, or the availability of a ticketing tool even in scenarios as automated as 5G EVE, has also been included.

Finally, another objective of the document is to define the requirements or capabilities that WP3 is imposing to 5G EVE sites that want to participate of the interworking solution that will be developed. And more specifically, a list of requirements for the site hosting the I/W Layer has been identified, together with some considerations in case new sites wanted to join the 5G EVE ecosystem.

Beyond the obvious need to make available the required compute capacity to actually install all the software modules forming the I/W Framework, the central site will have to consider topics like the required encryption capability to ensure correct communications across the star, the operational costs (updates, troubleshooting, bug resolution, resolution of physical impairments…), the capability to monitor the status of all the Orchestration Layer components from the hub, or the potential availability of mirror platforms which can be used for preliminary testing without affecting the system in production.

1 Introduction

Deliverable D3.2 – *Interworking Reference Model* – represents the main outcome of Task 3.2 – *Interworking Reference Model Definition* –, and specifies the selected architecture meeting all the requirements for the Interworking Framework that governs the experimentation services offered 5G EVE site facilities.

1.1 Initial context and terminology

In D3.1 – Interworking Capability Definition and Gap Analysis – [1], several concepts related to the Interworking Reference Model were already outlined as part of a preliminary specification effort. The objective was to identify the minimum set of functions to be supported, permitting the initial alignment (i.e. gap analysis) with the status at that time of the different 5G EVE sites. This specification was drafted from scratch, since no project before 5G EVE had tackled the global orchestration of interworking environments with multiple, different, local orchestrators Although there was previous work on which to lean on, there are still some key differences between 5G EVE and the state of the art of the required technologies, which are explained in section 2.

Figure 1 below depicts (in blue colour) the mentioned preliminary architectural definition, with the Interworking (I/W) Framework sitting between the Experimentation Portal, frontend for experimenters, and the different sites where experiments are to be deployed.

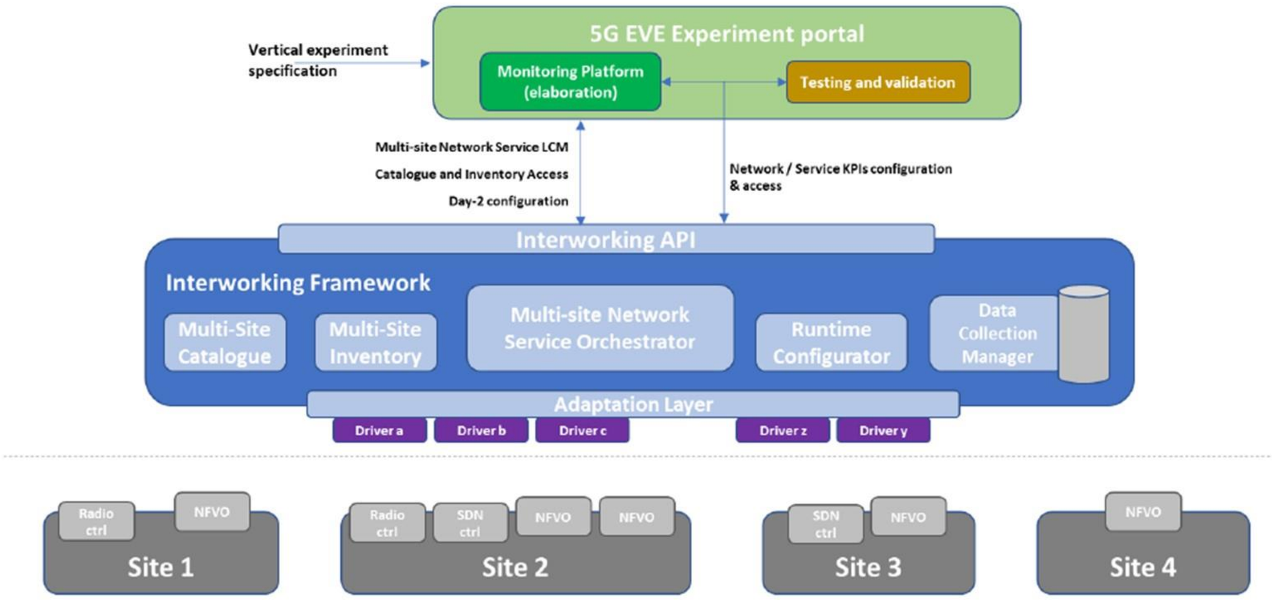

Figure 1: Interworking Framework architecture as defined in D3.1

Very briefly, the components that were already identified were:

- The **Multi-Site Network Service Orchestrator** (NSO) is responsible for coordinating the lifecycle of Network Services across the site facilities, leveraging on the local site orchestration components, and thus allowing for end-to-end network slices supporting verticals' experiments.
- The **Multi-Site Catalogue** stores Network Service Descriptors and verticals' VNFs. Descriptors can represent either multi-site network slice offers, exposed to experimenters via the Portal, or capabilities of the sites, directly collected from the facilities. Descriptors are not only referred to NFV but also to any other capability which may not be dynamically exposed.

- The **Multi-Site Inventory** keeps the information on provisioned and instantiated network slices across the 5G EVE end-to-end facility. In other words, it keeps the status of the slices which are actually up and running, helping the Multi-Site NSO decide where to deploy new slices.
- The **Data Collection Manager** collects all the performance metrics related to running Services, both to ensure the correct operation and to validate the targeted KPIs. This module also provides the performance measurements that help the Multi-Site NSO in the deployment decisions.
- Finally, the **Runtime Configurator** applies the required configurations to the provisioned Services and VNFs (and possibly other modules lacking standard interfaces) in run time.

Accompanying these modules, both northbound and southbound interfaces were also outlined. The **Interworking API** (NBI) offered the required set of services to the Experimentation Portal, including of course the provisioning of multi-site Network Services and slices. The mission of the **Adaptation Layer** (SBI) was to abstract the heterogeneous site capabilities, permitting the Interworking Framework to operate with common internal models, and allowing for the use of specific per-site drivers.

As described in Section 4 of this document, the basics of the previous architecture has been maintained now in D3.2. With the inclusion of detailed workflows, a few functionalities have been better delimited among the different components, and of course some others have been extended. However, since new interworking requirements have not appeared since the publication of D3.1, the architectural design has been maintained.

Together with the previous high-level architecture definition, additional preliminary analyses were included in D3.1, analyses which are summarized (and briefly revisited) hereafter.

As per site capabilities to be exposed to the I/W Framework, User Equipment and SIM cards, independently of the single or multi-site nature of the experiment, were declared as not to be exposed. The same was true for the subscribers' management and configuration features; these two functionalities are to be sorted out internally at each site. The rest, mainly the RAN, Distributed Cloud/MEC and SDN Controllers, and the Orchestrator, together with the monitoring and testing tools, were considered as candidates to be exposed through the Adaptation Layer, to permit the instantiation and execution of experiments. The only consideration then was that the local Controllers (RAN, Distributed Cloud/MEC and SDN) might also be hidden behind the local Orchestrator, as it actually still happens for the 5G EVE sites at the time of writing this Deliverable. Therefore, in the current design these controllers will not be exposed to the I/W Framework either. Of course, this approach may be reconsidered if needed in the future, according to new interworking needs or required capabilities.

Regarding the inter-site connectivity planes, in D3.1 it was commented that the orchestration plane communications could happen based on a star topology, centred at the site hosting the I/W Framework. However, it was said that for control and data plane multi-site communications, probably a full-mesh connectivity would be preferred. The final requirements from WP3, together with a connectivity proposal, are included in section 5 of the present document, for WP2 to consider in the finally adopted solution. The main key difference, as it will be seen, is that now it is considered that the control plane could also share the same star topology as the orchestration plane. This gives WP2 some flexibility in the early deployment stages, where the data plane is not so critical (during the first experimentation phases, experiments will not expand over multiple sites).

Finally, to end this introduction, it is worth mentioning that D3.1 also included a detailed overview of the current sites' capabilities and supported features, in line with those mentioned above (RAN, Distributed Cloud/MEC, SDN, MANO, etc.). That information was gathered mainly from WP2 Deliverables and gives an idea about the complexity of the multi-site interworking challenge at 5G EVE (see Figure 2, where the heterogeneity of the four sites is represented).

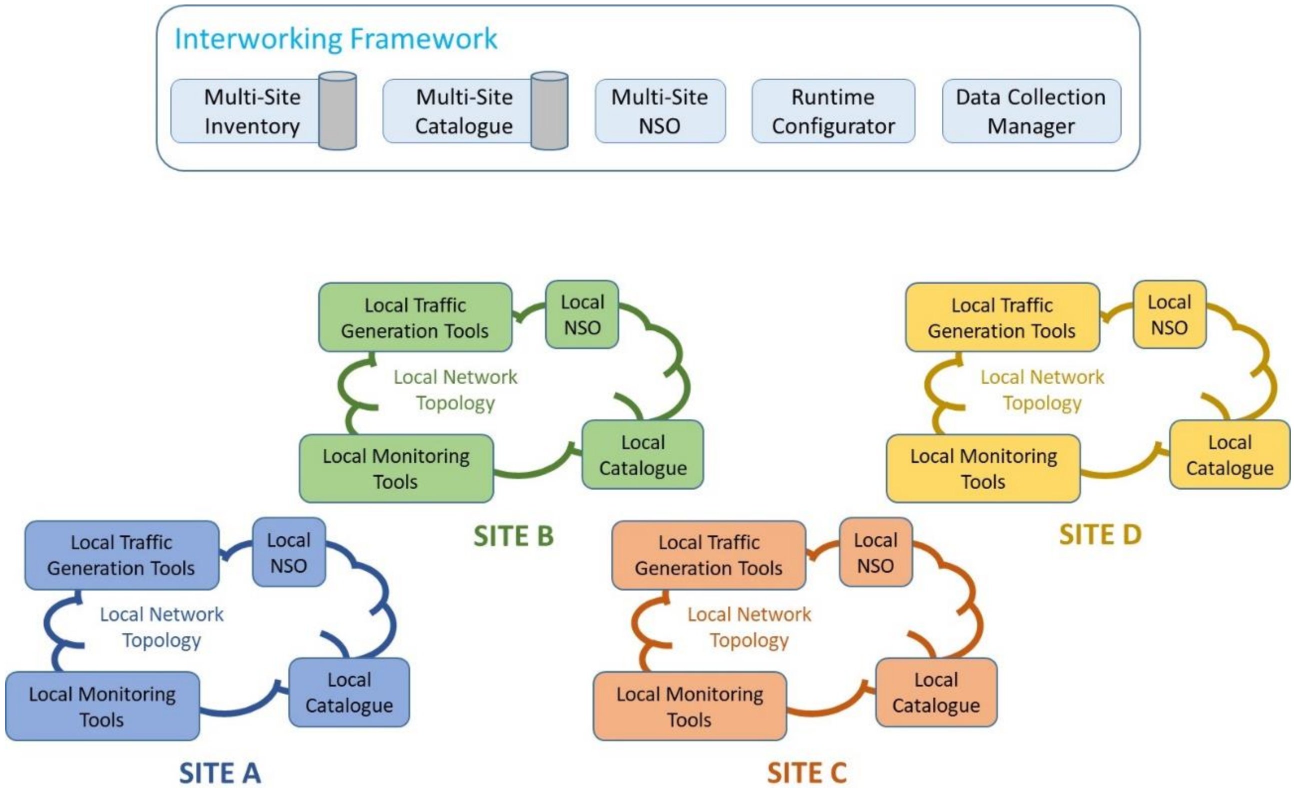

Figure 2: Orchestration challenge at 5G EVE

In the end, this is the reason why 5G EVE cannot lean only on current state of the art implementations, and has to build a new I/W Framework: such interworking control system just does not exist.

1.1.1 Deliverable D3.1 update

Due to the general scope of the project, providing an experimentation infrastructure not only for Verticals participating in 5G EVE, but also for other customer Verticals, it is foreseen that new interworking requirements appear in the future. To account for this, Task 3.1 – *Capabilities identification for interworking and gap analysis* – expands until Month 24, where the fully functional version of the Interworking Framework will be delivered. However, there are no additional Deliverables in T3.1 apart from D3.1.

Therefore, and as commented there, "*any further elaboration, improvement or update* [of D3.1] *will depend on each 5G EVE use case evolution and will be reported in future WP3 deliverables, if required and applicable*".

At the time of writing of this Deliverable, we do not have new requirements to update the list in D3.1 Therefore, the specification of the I/W Framework here contained considers only the original requirements list. In that sense, Deliverables D3.3 and D3.4 in Task 3.3 (implementation and deployment) will not only be used to report on new requirements when they appear, but also to report on their impact in the architecture (if any) and the modifications required to satisfy them (if any).

1.1.2 Terminology

Most of the terms used in this document are explained in their corresponding section for an appropriate understanding of all the described concepts. There are, however, a few of them which are used very often across the document, and that have been better considered for inclusion in a specific section, so that they can be easily consulted any time. Some of these terms have been inherited from other documents or Work Packages.

The first definitions deal with the different communication planes present in 5G EVE: orchestration, control, data and management planes, as depicted in Figure 3.

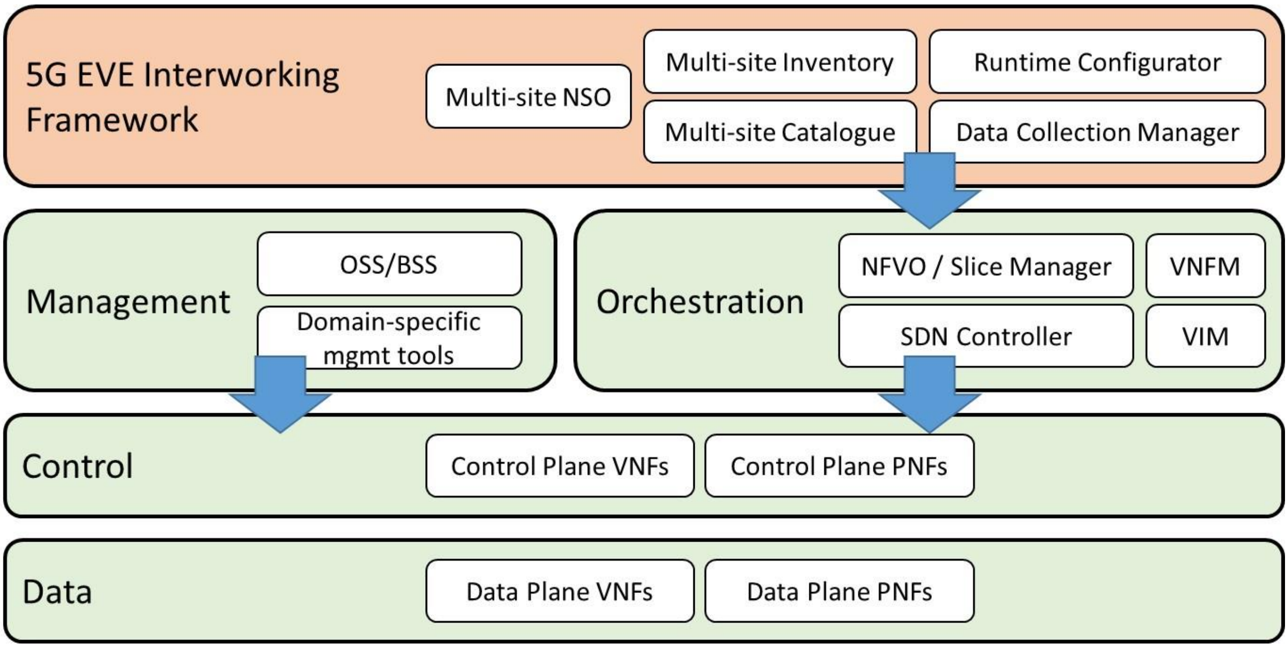

Figure 3: Communication Planes at 5G EVE

The **Orchestration Plane** normally deals with the higher modules of the 5G EVE Interworking Infrastructure. In this Plane is where the communications among components in charge of multi-site Network Services and Slices happen. Therefore, it is the plane that the Interworking Framework uses to communicate with the local orchestration-related components. Also, other processes like the data collection or the runtime configuration happen in this plane, which therefore expands to the deployed nodes (VNFs or PNFs) at the different sites.

The Orchestration Plane in 5G EVE also includes the set of local tools and solutions to provision all the segments of the 5G networks in each of the local sites. Apart from the VIM, dealing with the NFV resources, this layer hosts all the network controllers present in the facilities.

In the initial phase of experimentation, as it will be described later in this document, the components in the Orchestration layer will not be directly exposed to the I/W Framework, except for the local NFVOs. The rest (e.g., VIM or SDN Controller) will be controlled directly by those local NFVOs. This is expected to change in the future, according to new incoming interworking requirements (e.g., from ICT-19 Projects), but is still enough to force the Orchestration Plane to span across multiple sites.

The **Control Plane** is where network protocol communications happen. In 5G EVE, it is possible that network routers, for example, are deployed as part of the infrastructure supporting an experiment (equally in the form of VNFs or PNFs). It may also happen that these network routers are configured to exchange routing information, for example, using BGP. These communications are normally exchanged in-band with the Data Plane, but are considered a different plane since a different treatment can be provided via Quality of Service mechanisms.

The **Data Plane** deals with the physical infrastructure, that is, the RAN, the WAN (edge and core network nodes providing the end-to-end connectivity), the IT resources independently of where they are located (Distributed Cloud/MEC or Centralized Cloud), etc. Multi-site slices will also be built on top of the Data Plane.

Of course, both the Control and Data planes can also expand across multiple sites.

The consideration for the **Management Plane**, however, may be different. As of today, each local site has its own Management Plane used for direct access to the deployed components. This plane is critical for many internal reasons, including troubleshooting, monitoring, etc., but it is not currently envisioned that it has to be accessible to all the sites; at least, not from an I/W Framework point of view, since no interactions are currently envisioned among this component and the local systems like OSS or BSS.

Associated to these communication planes, in section 3.2 we describe the Data Plane **Connectivity Services**. These are services exclusively destined to the 5G EVE end customers, that is, the Verticals experimenting on top of the project's infrastructure. In other words, when an experimenter deploys an experiment, the mechanisms by which the VNFs, PNFs, etc., are interconnected are called Connectivity Services. They must not be mistaken by the **Interworking Framework Services**, which are the services offered by the Orchestration Layer supporting features that permit experimentation at 5G EVE: experiment deployment, experiment monitoring, KPI analysis and validation, etc.

On the other hand, the term **Interworking Connectivity** defines the connectivity solution among sites that allows for the interaction of the components in each of the Connectivity Planes. This solution is covered in section 5, and its final implementation is still being discussed with WP2.

Finally, it is worth mentioning that terms **Interworking Framework** and **Interworking Layer** have been used equally throughout the document. The same is true for **Multi-site Network Service Orchestrator**, **Multi-site NSO** or simply **Multi-site Orchestrator**.

1.2 Document objectives

The main objectives of this document (in line with those of Task 3.2) are:

1. Definition of the services supported by the 5G EVE platform, such as:
    a. Network connectivity services among VNFs
    b. Services directly offered to experimenters, like access to KPIs or automated deployment of VNFs.
2. Definition of the capabilities that sites need to deploy to jointly offer a multi-site platform for testing, capable of inter-site slicing.
3. Definition of a global orchestration framework compatible with the local orchestrators. In this sense, the material summarized in section 1.1, corresponding to Deliverable D3.1, represents a very good starting point for such final definition.
4. Definition, jointly with WP2, of the interworking connectivity solutions among sites, in terms of orchestration, control and data planes.
5. Definition of required supporting functionalities by the facilities (like hardening of all the infrastructure, a ticketing tool, access control based on authentication mechanisms, etc.), which may not be technically mandatory for the establishment of multi-site slices, but are required administratively or operationally wise.
6. Definition of requirements for the site hosting the I/W Framework.

The document starts with a high-level State of the Art analysis, which has been performed to identify clearly what components are already available, or can be leveraged on, and what gaps need to be directly addressed by the consortium.

It is also important to mention that this document is a design document, not governed by potential limitations of the technologies used for implementation. In other words, design specifications will be made as general and inclusive as possible, aligning with available standards even if these are still in their initial phases of specification, or not yet supported by the equipment deployed in sites. Later on, during the implementation phase, main efforts will be focused on those features that are strictly needed by the Verticals experimenting at 5G EVE.

The same can be said with regards to interoperability of the I/W Framework with deployed equipment. For components like the Runtime Configurator, or Data Collection Manager, it is not realistic to try and build standard interfaces, due to the great range of potential use cases. Instead, the first version of the I/W Framework will concentrate on those components that are available, and it will be extended in the future with new drivers in case they are needed.

This global design approach will help in the identification of gaps in the standards, where the project will be able to contribute, and will also permit a quicker implementation of future features when they are demanded.

1.3 Document structure

Sections of Deliverable D3.2 following this Introduction are organized as follows:

- Section 2 includes a State of the Art for the I/W Framework related topic, both in terms of standardization activities and of other European funded projects.
- In Section 3, we define the Services the Interworking Framework will support for Verticals experimentation, together with the Data Plane Connectivity Services.
- The Interworking Reference Model is fully covered in Section 4. The section is divided in four sub-sections, dealing with the internal components, the required workflows among those components, and the two needed external interfaces: the Interworking API and the Adaptation Layer.
- Section 5 outlines an Interworking Connectivity proposal among the different 5G EVE sites, already discussed with WP2, fulfilling the WP3 connectivity requirements.
- The requirements for additional supporting facility services and tools are described in section 6.
- Finally, section 7 is a summary for Site Managers, describing the capabilities that sites need to support to be able to provide the required interworking features and to host the I/W Framework.

2 State of the Art

In the previous Introduction it has been identified that a system fulfilling all the interworking requirements for 5G EVE has not yet been implemented. This does not mean, though, that the low-level design of the Interworking Framework has to be fully done from scratch. The objective of this section is, on the one hand, to present similar efforts on which we will lean in the design of components, workflows or interfaces of the I/W Framework; on the other hand, to justify why some activities which might seem similar are in fact not applicable to the 5G EVE scenario.

The resultant State of the Art provides a high-level view: it briefly explains differences and similarities with the interworking solutions adopted by other organizations, and it provides the required references to be able to expand on the details if desired.

As an introduction to position the content of the State of the Art within the 5G EVE Platform, Figure 4, not intended to be totally exhaustive, matches the SDOs and Projects described in the following sections with the I/W Framework architecture of D3.1. As seen, there are many sources that 5G EVE can use as starting point to face the complex orchestration challenge presented in Figure 2. However, many of the project's compromises will demand costly extensions to the current State of the Art, and furthermore, there has not been so far any effort to build the whole system, interoperating across all layers and among all components.

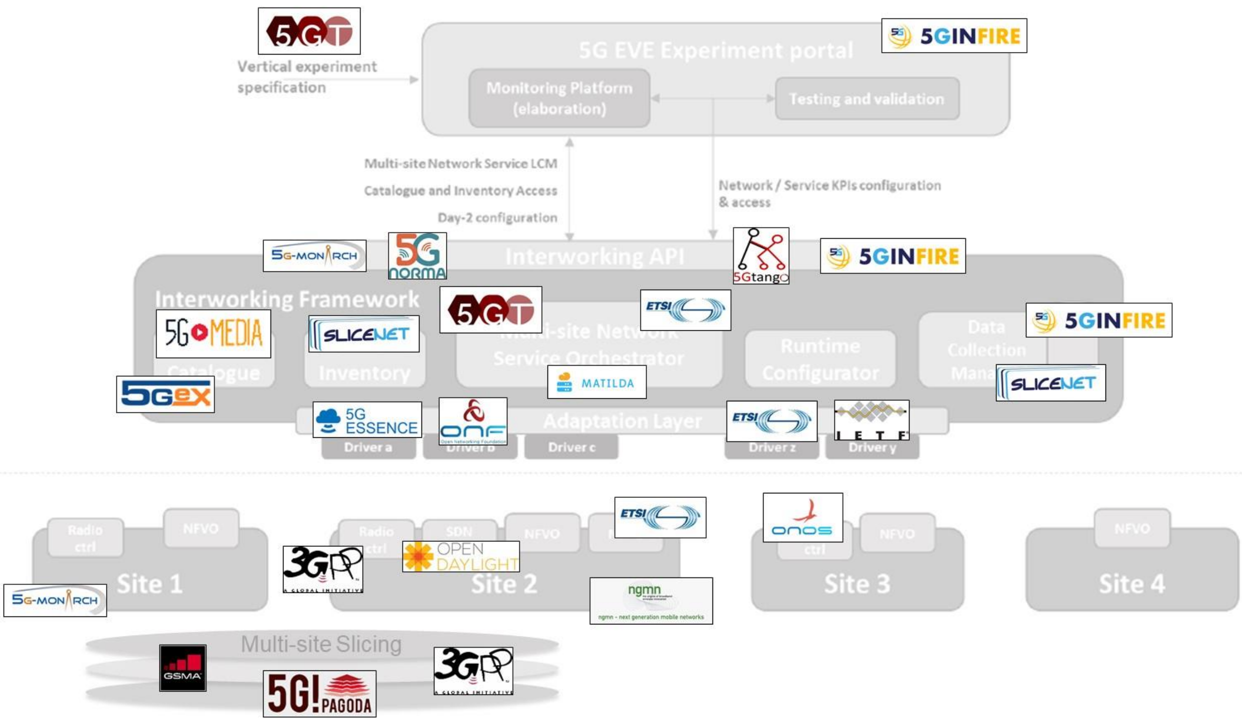

Figure 4: **SotA summary**

Another conclusion from the previous exercise is the limited work on the Runtime Configuration module. Of course, there are technologies which 5G EVE will leverage for its implementation, like for example Red Hat Ansible[1]. But one key idea obtained from Figure 4 is that no project is actually developing such an extensive and automated framework for Vertical's multi-site experimentation and KPI analysis and validation, equally

[1] https://www.ansible.com/

capable to perform technology benchmarking and measurement of 5G network KPIs. The closest component has been found in 5G-MEDIA, but it only applies to NFV.

2.1 Standardization related activities

2.1.1 Network Function Virtualization

European Telecommunications Standards Institute (ETSI) NFV (Network Function Virtualization) is an initiative to virtualize network functions and services traditionally run on proprietary, dedicated hardware. With NFV, Network Functions (NFs) like routing, load balancing or firewalling are packaged as Virtual Machines (VMs) and are instantiated on top of commodity hardware (e.g. x86-based servers). NFV has for that purpose introduced the Virtual Network Functions (VNFs) concept, software implementation of NFs, allowing the collocation of multiple instances on top of the same virtualized environment. For these reasons NFV is deemed as a critical technology to enable the transition toward 5G.

ETSI proposes a complete reference architecture that defines functional blocks and interfaces to describe the NFV framework, organized in three levels of entities and functions: VNFs, NFV Infrastructure (NFVI), and NFV Management and Orchestration (NFV MANO). Figure 5 describes this framework with the focus on the MANO level.

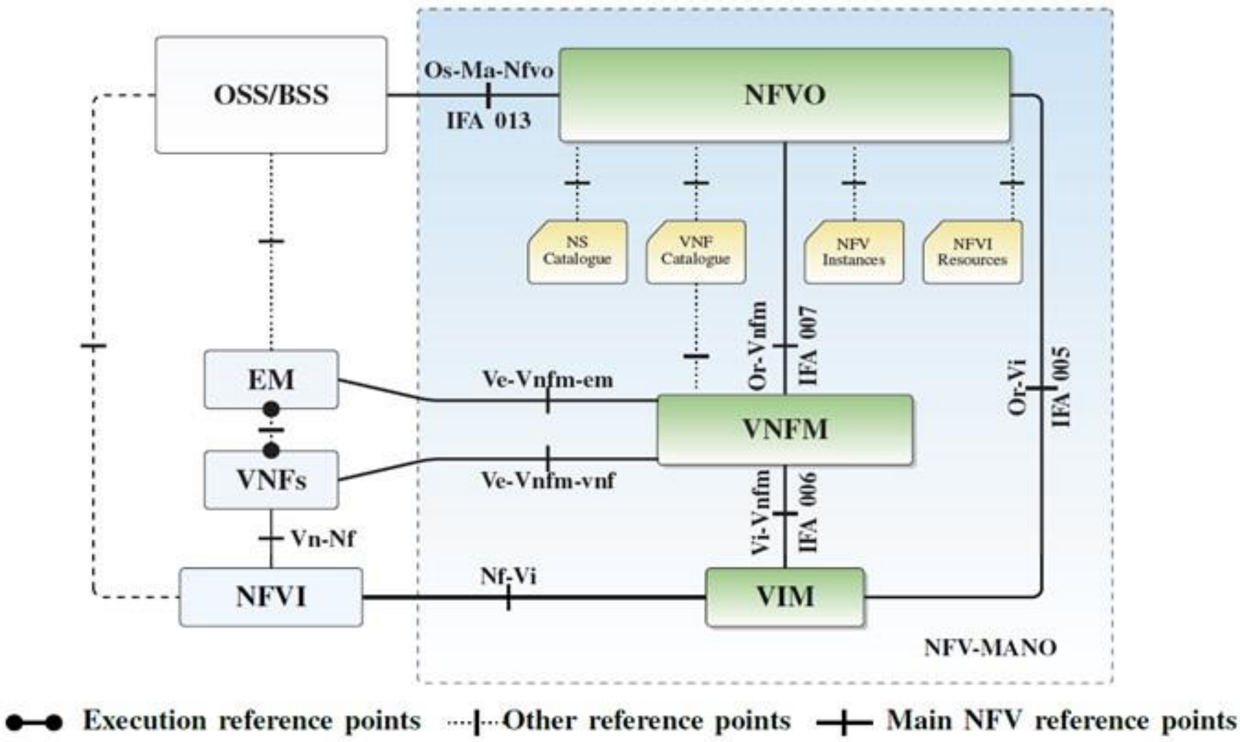

Figure 5: NFV-MANO architectural framework

NFV MANO is responsible for the orchestration and management of the NFVI resources on which it deploys Network Services (NSs). NFV MANO has several responsibilities to automatically instantiate and monitor a NS using available resources. First, it should discover and manage all the available virtual resources. Then, be able to on-board NSs and VNFs. Finally, it has to instantiate those NSs using the discovered resources, and manage their life cycle.

NFV MANO shares these responsibilities among different layers, connected to each other by normalized interfaces. These layers are the Virtual Infrastructure Manager (VIM), the VNF Manager (VNFM) and the NFV Orchestrator (NFVO).

In particular, the role of the NFVO (the Multi-site NSO can be seen as a hierarchical NFVO) is to on-board NSs, and manage their lifecycle. To receive NS creation requests, NFVO must be connected to an external module, which is aware of the provider’s business (e.g., Operating and Business Support System (OSS/BSS) module, or in the case of 5G EVE, the Experimentation Portal). The NS is described via a Network Service Descriptor (NSD). The NSD details the VNFs, Physical Network Functions (PNFs), Virtual Links (VLs) and VNF Forwarding Graph (VNFFG) that compose the NS. The NFVO identifies and orchestrates the instantiation of all these components.

The ETSI NFV Industry Specification Group has defined also the descriptors for all network entities in the NFV environment. A high level, technology agnostic description of the structure and composing fields of each entity is given in NFV-IFA (InterFaces and Architecture) deliverables. More specifically NFV-IFA 011 [2] defines what is included in a Virtual Network Function Descriptor (VNFD) and how to build a VNF Package, which is an archive with the VNFD plus other files (metadata). NFV-IFA 014 [3] defines what is included in a Network Service Descriptor (NSD).

NFV-SOL (SOLutions) deliverables propose an implementation-ready solution for concepts defined in the NFV-IFA documents. The NFV-SOL 001 [4] deliverable defines a mapping of information elements in NFV-IFA 011 and NFV-IFA 014 to data types of TOSCA Simple Profile in YAML [5]. YAML files are provided to be downloaded and used to write custom NSD, VNFD, etc. [6]. To on-board an NSD on the NFVO, the YAML files must be packaged with a proper format. NFV-SOL 007 [7] and NFV-SOL 004 [8] describe the NSD package and the VNF package, respectively. The documents include structure, format, contents and the naming conventions for the different files in each package. The format used is TOSCA CSAR.

Other deliverables describe the operations to manipulate the aforementioned descriptors. NFV-IFA 013 [9] defines the interfaces and exchanged information models (descriptors/entities/instances) between the Operations Support System (OSS)/Business Support System (BSS) and the Network Function Virtualisation Orchestrator (NFVO). The interface is called Os-Ma-nfvo (see Figure 5). While NFV-IFA 013 specification is abstract, a RESTful protocol and data model implementation is given in NFV-SOL 005 [10], which implements all the operations defined in NFV-IFA 013. The API description is also provided in OpenAPI standard [11].

For more clarity, here follows a table with relevant correspondences between NFV-IFA and NFV-SOL deliverables.

Table 1: Mapping between NFV interfaces and architectures (NFV-IFA) and relevant solutions (NFV-SOL)

NFV-IFA	Short Description	NFV-SOL
IFA 011 (VNF Package, VNFD)	MANO; VNF Descriptor and Packaging Specification	SOL 001 SOL 004 (VNF Packaging)
IFA 013 (Os-Ma-nfvo)	MANO; Os-Ma-Nfvo reference point - Interface and Information Model Specification	SOL 005 (Os-Ma-nfvo)
IFA 014 (NSD)	MANO; Network Service Templates Specification	SOL 001 (VNFD, NSD) SOL 007

The close relationship of 5G EVE with the ETSI NFV initiative is obvious for many reasons, mainly because all sites implement some sort of local NFV MANO architecture. Similarly to what they may be supporting in their local North-Bound Interfaces (NBI), the just mentioned NFV-SOL 001 descriptors, the NFV-SOL 004 packages and the NFV-SOL 005 API specification, can be adopted at the North-Bound Interface (NBI) of the Multi-site Network Service Orchestrator at the global Interworking Framework. However, these are not sufficient for the part of the South-Bound Interface (SBI) that interconnects with the local Orchestrators.

For the expansion of composite NSs across multiple administrative domains, NFV-IFA 030 [12] has started the specification of a reference point that ensures the interoperability of NFVOs in such scenarios. This reference point is called Or-Or, and although its specification is still in an early stage, it will be the basis for the communication of the Multi-Site Network Service Orchestrator with the local NFVOs at the 5G EVE sites.

2.1.2 Network Slicing

In the most general case, network slicing aims for building dedicated logical networks that exhibit functional architectures customized for their respective offered services (e.g. eMBB, mMTC, etc.). Slices can provide the full functionality of a complete network, including access and core functions that can be of different vendors. To support slicing, the slice manager creates a group of network resources and instantiates all the network and service functions assigned to the slice.

The key advantage of slicing is that it is designed to satisfy the demand of dedicated telco services with specific Service Level Agreements (SLA). This way, slicing permits to create isolated logical network partitions that ensure the availability of the required resources.

As commented in D3.1, the general reference for modelling and description of end-to-end slices which will be adopted in the I/W Framework is ETSI NFV-EVE 012 [13] (EVE – Evolution and Ecosystem).

2.1.2.1 3GPP SA5 – Network Slicing concept and multi-site scenarios

The scope of 3GPP SA Working Group 5[2] (SA5) – Telecom Management – is the “specification of the requirements, architecture and solutions for provisioning and management of the network (RAN, CN, IMS) and its services”. The Working Group also deals with the definition of associated charging solutions.

SA5 specifies at 3GPP TS 28.533 [14] the Management and Orchestration (O&M) solution for 5G networks using a Service Based Architecture (SBA). According to this architecture, the management services for a mobile network are provided and consumed through APIs. This model foresees that any authorized consumer can access the management services, and those services can also be exposed externally for interworking purposes with other administrative domains. Therefore, the new 5G management architecture is no more based on reference points, as in 4G, statically linking together two NFs. This new flexible approach permits building up a new management model and offers more flexibility to let different administration domains to interwork, as is the requirement in 5G EVE.

To understand at what level the interworking can be performed, it is important to look at the management architecture defined by SA5 in 3GPP TR 28.801 [15]. New Management Functions (MFs) are foreseen to manage the network using a Network Slicing approach. It is important to notice that SA5 does not specify Management Functions but only the Management Services (MnS) according to the SBA approach. It is also useful, therefore, to consider the Management Functions, logic aggregators of the MnSs, to better understand the management roles that compose a 5G Management and Orchestration.

SA5 foresees two levels of management for a Network Slice, defining it as a collection of managed entities called Network Slice Subnets. The Network Slice Subnet, in turn, is a collection of NFs. In SA5 terminology, a deployed Network Slice is called Network Slice Instance (NSI) and a deployed Network Slice Subnet is called Network Slice Subnet Instance (NSSI).

According to this model, a possible deployment of the MnSs can be done grouping them in Management Functions according to the following roles:

- Network Slice Management Function (NSMF): responsible for management and orchestration of Network Slices, it requests the allocation of the network slice subnets to NSSMF.
- Network Slice Subnet Management Function (NSSMF): responsible for management and orchestration of Network Slice Subnets. It takes care of specific management domains dealing with the NFs

[2] https://www.3gpp.org/Specifications-groups/sa-plenary/56-sa5-telecom-management

allocation. In case of a VNF, the NSSMF requires the NFs allocation to a virtualization environment manager (e.g. NFVO or VNFM according to ETSI MANO specifications).

Another higher level management function is foreseen by SA5 to cope with service requests in terms of Communication Service and related SLA. This Management Function is called:

- Communication Service Management Function (CSMF): this function takes care of the management of the communication service and translates the requirements related to the communication service to network slice related requirements, which are sent to the NSMF.

Figure 6 summarizes the objects managed by the Management Functions and their interactions.

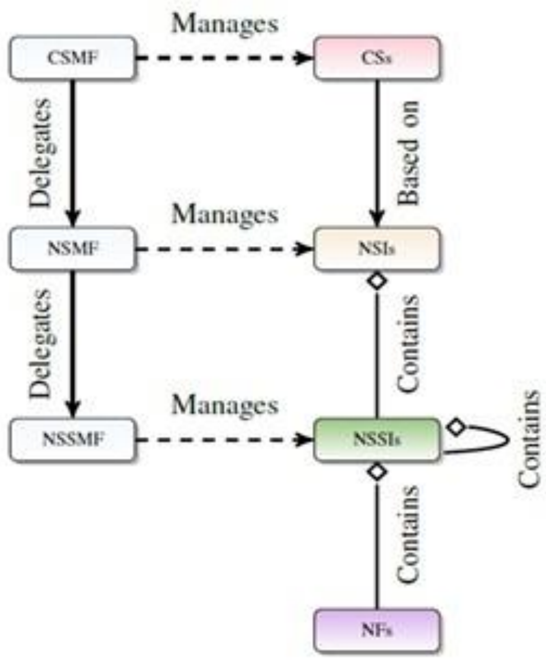

Figure 6: 3GPP Network Slice related MFs

According to this model, different scenarios for interworking can be possible. When a customer requests a Communication Service that leverages on Network Slices that have to be deployed in different administration domains (e.g., different 5G EVE sites), interworking functionalities are required. An interworking feature that is required is the exposure of the slice management capabilities.

SA5 foresees the following options to instantiate a Communication Service that leverages on different NSIs of different operators (e.g., mapped to different 5G EVE sites):

- The Customer, via its CSMF, asks to the different operators' O&Ms the creation of different NSIs all related to the same Communication Service. Each O&M internally requests the creation of the NSI to its NSMF. This option requires the customer's CSMF to manage multiple instances of NSI to make them support a single service instance.
- The Customer asks the creation of the Communication Service to a single operator's O&M. The CSMF of that operator (in O&M #A) can decide to request an NSSI to another operator's management system (O&M #B) together with the exposure of a set of slice management services. For the O&M #B this results into a request for an NSI. It is important to notice that an NSSI is just an O&M internal concept: if O&M #A needs an NSSI from another operator's O&M, the entity that is requested to O&M #B is always an NSI. That NSI is used by O&M #A as an NSSI. In this option, the communication service management functionality is hosted by Operator #A. Operator #A is responsible for the management of the NSI.

- Multi-operator NSI creation by operator NSMF to NSSMF interface. The use case is similar to the previous one but this option foresees a stronger relationship between the operators with direct access to NSSMF. The NSMF of operator A, to build up the requested NSI, decides to use another operator NSSMF to realize the NSSI.

Figure 7 portrays how multiple operators can coordinate to manage instances of network slices.

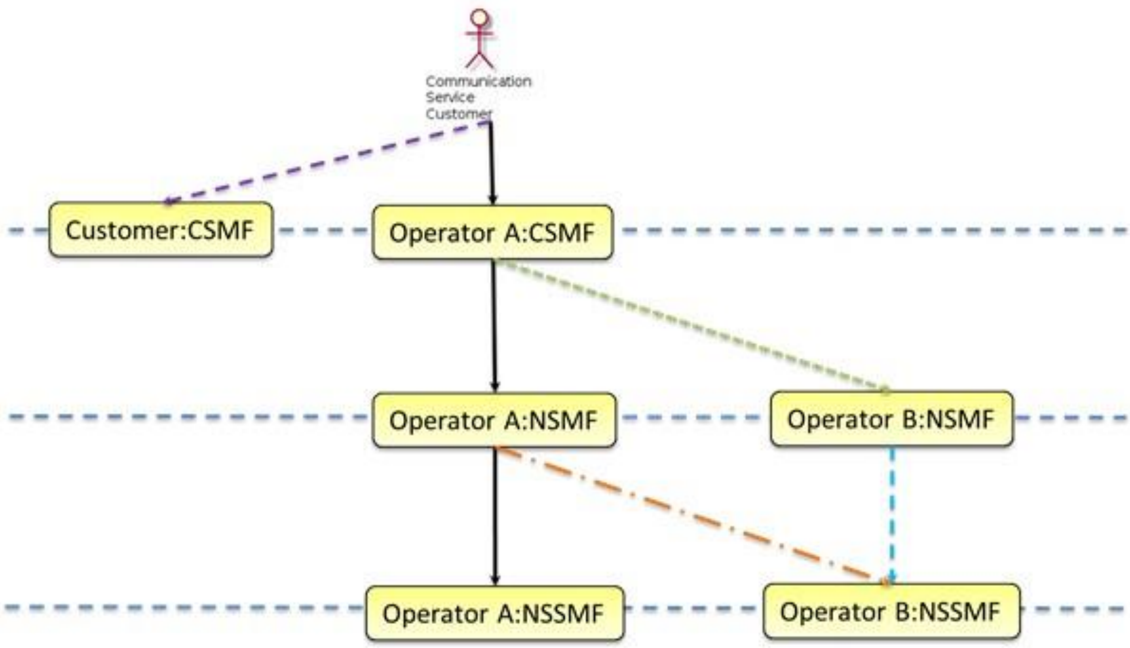

Figure 7: Multiple operator coordination management

The offering of slices implies the exposure of management services to the customer to let him partially control the slice. SA5 defines in 3GPP TS 28.530 [16] a possible deployment scenario with an Exposure Governance Management Function (EGMF), which intermediates the exposure of the management APIs to external consumers. EGMF is a service consumer of other Management Service producers with the aim to expose those management services to other consumers such as the management system of another Operator or to some Vertical industry. The level of service exposure can be different according to the different consumers and can be policy driven.

An important feature for the interworking model in an SBA environment is the Discovery Service. For 5G network functions, the 3GPP SA2 working group defines the Network Repository Function (NRF) for registration, discovery and authorization. For the same purpose 3GPP SA5 is working on a Discovery Service for MnSs. The Discovery service is important in an SBA architecture because any authorized MnS Consumer can be part of the O&M and so it should be able to discover other MnS (it has to be authorized to consume them). The Discovery service is important for the internal interactions inside one O&M but it could be also interesting to study a possible role in an interworking scenario.

2.1.2.2 3GPP SA5 – Slicing Lifecycle Management

According to 3GPP TS 28.530 [16], a NSI lifecycle is composed of four phases as depicted in Figure 8.

i. The **Preparation Phase** includes the design templates, the planning capacity, the on-boarding and the evaluation of the network slice requirements, as well as the preparation of the network environment.
ii. The **Commissioning Phase** consists of NSI provisioning, which includes the NSI creation, wherein all the needed resources for the slice, whether shared or not, are allocated and configured. The creation of an NSI can include the creation and/or the modification of the NSI constituents.

iii. The **Creation Phase** is a composition of several activities namely: activation, supervision, performance reporting, resource capacity planning, modification, and de-activation. The activation makes the NSI ready to support the traffic for Communication Services. The resource capacity planning activity includes any action that calculates resource usage based on NSI provisioning and performance monitoring, and generates modification polices as a result of the calculation. The NSI modification may include capacity or topology changes, and represents the creation or modification of NSI constituents. NSI modification can be triggered by receiving new network slice requirements or as the result of supervision/reporting. The deactivation includes actions that make the NSI inactive and stops the CSs.

iv. The **Decommissioning Phase** consists of the decommissioning of the non-shared constituents and the removing of the NSI specific configuration from the shared constituents. After the decommissioning phase, the NSI is terminated and does not exist anymore.

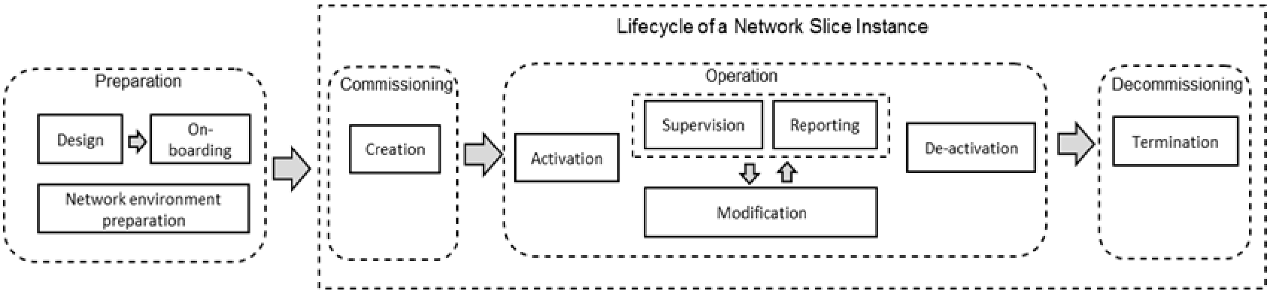

Figure 8: 3GPP Network Slice Lifecycle Management

2.1.2.3 GSMA – Network Slicing templates

Network slice implementations differ among themselves in the set of attributes that define them, and also on the values associated with those attributes. The flexibility that such definition permits might, however, become a constraint with regards to the interaction with customers, as it could become complex to manage. With the objective of characterizing together different types of slices sharing a set of attributes (even if not all of them), GSMA has been proposing the concept of "templates".

A Generic Slice Template (GST) is a common set of attributes that define a type of network slice. These are generic, in the sense that are not related to any kind of deployment. When the list of attributes of a GST is given values, then this creates a NEtwork Slice Type (NEST). The NEST is the input to the NSI Preparation phase described in the previous section, phase where the Network Slice Template (NST) is created. GSMA defines at its Official Document NG.116 [17] a detailed set of GST attributes.

It is worth mentioning, however, that these concepts of GST and NEST are probably more directly applicable to the Experimentation Web Portal and WP4, than to the I/W Framework of WP3. When defining an experiment in 5G EVE, the Vertical is first given the chance to select any of the available "Service Blueprints", which actually match the project's Use Cases (e.g., Industry 4.0, IoT, Smart City…). Within each Use Case, a number of desired tests are meaningful for the experimenter; the additional elements required to execute each test (e.g., a traffic generator, an active or passive probe…) define what is called the "Context Blueprint". Each test will have its own associated Context Blueprint, and the sum of Service and Context Blueprints generate the "Experiment Blueprint". This Experiment Blueprint includes all the components required to execute a test, but the test conditions are not yet defined; in that sense, it is a concept very similar to the GST.

When the experimenter fills in all the parameters that govern the test (e.g., the initial bandwidth of the background traffic, and its change over time), then the "Experiment Descriptor" is obtained. In this case, the definition is very similar to that of the NEST, furthermore considering that the Experiment Descriptor is not even the one required by the Multi-Site Orchestrator to instantiate a NS: it is the Network Service Descriptor (NSD).

2.1.2.4 NGMN – 5G Network and Service Management

The NGMN Alliance (Next Generation Mobile Networks Alliance) is an open forum formed by mobile operators, vendors, manufacturers and research institutes. Its main goal is to evaluate candidate technologies for the next evolution of wireless networks and to ensure that the standards for next generation network infrastructure, service platforms and devices will meet the requirements of operators and users.

In one of their publications, "5G Network and Service Management Including Orchestration" [18], some considerations regarding the impact of network slicing with regards to O&M are given. Based on user stories, both pre-conditions, requirements and/or procedures are outlined for required features like the creation of a network slice, the design of an end-to-end (E2E) Service Management, or the application of slicing to multi-operator environments. Although at a very high level, alignment of the I/W Framework workflows with these procedures will ensure the usability of such workflows beyond 5G EVE.

The document also includes a list of requirements "from a Network and Service Management including Orchestration perspective". Its main value for 5G EVE is that it mentions a set of desired functionalities to ensure the correct operation of the facility. Examples of those include self-healing, scalability, proactive monitoring, capacity planning, etc., and the most immediate ones of them have already been tackled in this Deliverable (see section 6).

2.1.3 SDN Control

SDN is an architectural concept that divides the forwarding plane from the control plane. In the conventional architecture, a network device makes its own decisions about where a packet is forwarded. SDN moves the control plane outside the network device: forwarding decisions are taken by external software, and injected into the device.

In the current state of the 5G EVE facility, only two sites have deployed SDN Controllers. In both sites, these Controllers are under the direct governance of the local Orchestrators, so they are not exposed to the I/W Framework and there is no need to implement a SBI towards them. This is so because they are mainly in charge of controlling the local NFVI, process already supported as part of the VIM relationship with NFVO.

However, this may change in the future. For example, one site deploying an SDN Controller for the WAN may choose between giving the higher control to the local Orchestrator (via the WIM interface) or deploying the Controller standalone (as very few WIM implementations are available up to date). In such cases, it might be required to upgrade the I/W Framework.

2.1.3.1 Available SDN Controller implementations

ONOS

ONOS [19] (The Open Network Operating System) is an open source SDN network operating system released in 2014 as an open-source project by ON.Lab.

ONOS was created as a distributed system in the form of a cluster (Distributed Core), where all the instances have identical functionality. The network operator can add servers incrementally if needed without disruption. Each instance works together as a single platform, applications and network devices not being aware if they are working with a single or multiple instances of ONOS. This brings a highly available environment.

ONOS provides two powerful Northbound abstractions: the Global Network View and the Intent Framework. The Intent Framework allows an application to request a service from the network without having to know the details of how the service will be performed. In other words this allows network operators or application developers to program the network at high level. The Global Network view provides the application with a view of the Network: hosts switches, links and any other state associated with the network.

The ONOS project is supported by a large open-source community, which has improved the platform and has introduced new compelling applications. The ONOS project is run under the governance of ONF (Open Networking Foundation). ONF is a non-profit consortium led by major telecom and service operators.

OpenDayLight

OpenDaylight[3] (ODL) is a modular open platform managing networks according to the SDN paradigm, and focused on network programmability. Its permits customizing and automating network management and control procedures, obtaining flexible and responsive networks. At the same time, ODL is driving network automation to improve operational efficiency.

ODL is also an open source SDN controller, in this case developed by a community under the Linux Foundation umbrella. The community is collaborative and is supported from vendors and organizations addressed to the SDN and NFV industry. The code is integrated in more than 35 vendor solutions and addresses a variety of use cases in existing network environments. It is also part of open source frameworks such as ONAP, OpenStack, and OPNFV.

The main technical concept behind the ODL platform is the Model-Driven Service Abstraction Layer (MD-SAL). The SAL architecture is a publish/subscribe model that exchanges information from its data-stores. Each data in the SAL is a data processing and adaptation between YANG models representing network devices and applications.

Finally, ODL exposes a set of APIs at the North-Bound level that provide the network programmability through RESTCONF or NETCONF APIs. The modularity and flexibility of ODL allows to be integrated in OpenStack by a dedicated plug-in which provides programmability to Neutron networks. Northbound protocols and control plane services are connected by the MD-SAL.

2.1.3.2 Standardization activities at ONF

T-API [20] (Transport API) is a standard API defined by the Open Networking Foundation (ONF) that allows a T-API client, such as a carrier's orchestration platform or a customer's application, to retrieve information from and to a transport domain controlled by a T-API server, such as a Transport SDN Controller. This could be applicable in 5G EVE in case any site deployed an SDN-controlled optical WAN.

T-API supports the following services:

- Topology: it supports the retrieval of topology information from the Controller in the form of node, link and edge-point details.
- Connectivity: it allows the client to retrieve information and request point-to-point, point-to-multipoint and multipoint-to-multipoint connectivity services across the transport network.
- Path Computation: it allows the client to make a request for paths Computation and Optimization.
- Virtual Network: it allows the client to create, update and delete Virtual Network topologies.
- Notification: it allows the client to subscribe for events from the server.
- OAM: it allows the client to determine where monitoring points may be present as well as to start, terminate, enable and disable measurement services between specified points in a connection.

2.1.3.3 NETCONF and YANG models

As stated above, the original concept of SDN was about removing the forwarding decisions outside of the deployed nodes, and executing them in external software. Afterwards, the result of those decisions was pushed towards the nodes, using typically the OpenFlow protocol.

While OpenFlow is still used in simple NFV environments (e.g., those based on VLAN forwarding), the increased complexity of the datacentres is producing a migration towards Layer 3-based solutions in the fabric stages. These have scaled for years in the operators networks (e.g., BGP), and are now being adopted even by the most important hyper-scalers. In this sense, the definition itself of SDN is also changing.

In many aspects, SDN in the transport networks, and by extension in the datacentre, is about automation, mainly regarding the deployment of connectivity services, being NETCONF/YANG one of the potential solutions. The

[3] https://www.opendaylight.org/

Network Configuration Protocol (NETCONF) is being standardized by the IETF (RFC 6241 [21]), and permits the installation of configuration data – encoded as XML – in network nodes.

YANG, instead, is a data modelling language (also defined by IETF – RFC 6020 [22]) complementing protocols like NETCONF. "Configuration templates" required for specific services can be modelled using YANG, which permits interoperability by adoption of the same models. IETF has specified, and still is specifying, models for numerous features. However, it is not the only initiative standardizing YANG models: others like OpenConfig[4] have put their main focus on compiling a set of data models driven by use cases from network operators.

5G EVE sites moving from VLAN-based forwarding to new solutions like E-VPN over VxLAN, for example, will require the support of this type of features if automation is desired. The degree of support by the industry and by SDN Controllers like those mentioned above (ONOS & ODL), fortunately, is growing very fast, although there is still space for improvement, and for contributions in the standardized YANG models for the different services.

2.2 European funded projects

The following sub-sections include brief descriptions of other European funded activities tackling some of the 5G EVE Interworking Framework's concepts, like interfaces, components or workflows. Beyond these descriptions, the section's objective is to clearly identify points of coincidence on which we can leverage, and key differences that clearly state where projects diverge. This analysis has permitted the consortium to focus on those items which required innovation beyond current State of the Art.

The first three projects (5G-TRANSFORMER, 5GinFIRE AND 5G-MEDIA) are those which are more related to WP3 activities and its participating partners, so the descriptions are more detailed. Also, as an initial summary, Table 2 permits a quick understanding of what these projects bring to the declared objective: identify which results are or not reusable for the I/W Framework.

Table 2: Management and Orchestration mechanisms comparisons – 5GinFIRE and 5G EVE

Component	Description	Are results reusable?
5G EVE		
Multi-site NSO	• Centralized I/W Framework must orchestrate local MANO implementations and potentially other Controllers, like SDN, which may differ from site to site	
Runtime Configurator	• Required for Day2 configurations, configurations of non MANO/SDN controlled nodes, and to automatically vary experiment conditions	
Data Collection Manager	• Used both for Platform and Experiment monitoring • Capability to provide rests results for report generation	
Multi-site Catalogue	• Integrated with Portal	
Multi-site Inventory	• Required for automated control of active experiments	
NBI	• Web Portal interfacing multiple components, including Multi-site NSO	
SBI	• I/W Framework mainly interfacing NFVO components at sites	
5G-TRANSFORMER		

[4] http://www.openconfig.net/

Multi-site NSO	• Multi-domain service orchestration is performed in a distributed manner through horizontal interfaces among service orchestrators	**NO**
Runtime Configurator	• Not present	**NO**
Data Collection Manager	• Integrated monitoring infrastructure for NFV and service related metrics collections (based on ETSI NFV-IFA operations and information models)	**PARTIALLY**
Multi-site Catalogue	• Not present, each single domain service orchestrator stores partial multi-domain catalogue information (based on ETSI NFV-IFA information models)	**NO**
Multi-site Inventory	• Not present, each single domain service orchestrator stores partial multi-domain inventory information (based on ETSI NFV-IFA operations and information models)	**NO**
NBI	• Based on ETSI NFV-IFA operations and information models	**PARTIALLY**
SBI	• Mostly providing access to VIM interfaces (following ETSI NFV approach) and transport network provisioning interfaces	**PARTIALLY**
5GinFIRE		
Multi-site NSO	• Centralized MANO (OSM) and distributed VIMs • Only local SDN support at some sites; no centralized control	**NO**
Runtime Configurator	• Not present; pure NFV testbed; experimenters have access to their VNFs for configuration	**NO**
Data Collection Manager	• Only Platform monitoring	**PARTIALLY**
Multi-site Catalogue	• Integrated with Portal	**YES**
Multi-site Inventory	• Not present; manual control of active experiments	**NO**
NBI	• Web Portal only implementing NBI of OSM	**PARTIALLY**
SBI	• Only OSM-VIM interface towards sites	**NO**
5G-MEDIA		
Multi-site NSO	• Not present, the focus is on single domain orchestration based on OSM	**NO**
Runtime Configurator	• Integrated Day2 configuration platform based on Kafka message bus and executor services for configuration of heterogeneous functions (pure NFV and media related)	**PARTIALLY**
Data Collection Manager	• Integrated monitoring infrastructure based on Kafka message bus and data collection points following common models	**PARTIALLY**
Multi-site Catalogue	• Centralized catalogue able to sit on top of several technology and domain specific catalogues	**YES**
Multi-site Inventory	• Not present	**NO**
NBI	• Based on OSM NBI for single domain scope	**NO**

SBI	• Only OSM-VIM interface towards heterogeneous NFV infrastructures	**NO**

Finally, in section 2.2.4, an additional very specific list of projects is included, focusing on other scope-limited features but which are also of interest for 5G EVE.

2.2.1 5G-TRANSFORMER

5G-TRANSFORMER[5] is a 5G-PPP Phase 2 project that aims at developing an SDN/NFV-based platform for the delivery of vertical-tailored network slices as a mean to facilitate verticals industries in provisioning their services over mobile transport networks. In particular, 5G-TRANSFORMER enables verticals to easily meet their service requirements through customized 5G end-to-end slices. This is done through aggregation and federation of transport networking and computing fabric resources from the edge up to the core and cloud, for creation and management of slices on a federated and virtualized infrastructure.

This requires a transformation of current mobile transport networks into an SDN/NFV-based Mobile Transport and Computing Platform (MTP), bringing network slicing as a core paradigm for provisioning MTP slices tailored to the specific needs of vertical industries. To do this, the 5G-TRANSFORMER platform is based on extensions of the ETSI NFV MANO architecture: as depicted in Figure 9 below, it consists of three main functional components:

i. A Vertical Slicer (VS) for service and slice management [23].
ii. A Service Orchestrator (SO) for service and resource orchestration in multi-domain, federated scenarios [24].
iii. The MTP, representing the underlying unified transport stratum, responsible for providing the networking and computing resources required by the NFV NS orchestrated by the SO [25]

The main touching points of 5G-TRANSFORMER and 5G EVE (targeting its interworking framework) are the SO and the VS principles and functionalities. For what concerns the SO, 5G-TRANSFORMER addresses end-to-end service and resource orchestration across different administrative domains through a federation paradigm. In particular, end-to-end services are split into multiple segments deployed in different administrative domains, based on service requirements and resource availability. Federation is managed at the interface between SOs belonging to different domains and handling abstraction of services and resources. Therefore, it is clear that 5G-TRANSFORMER follows an approach with a fully distributed end-to-end service orchestration, without any multi-site Network Service Orchestration (NSO) as in the 5G EVE interworking framework case.

[5] http://5g-transformer.eu/

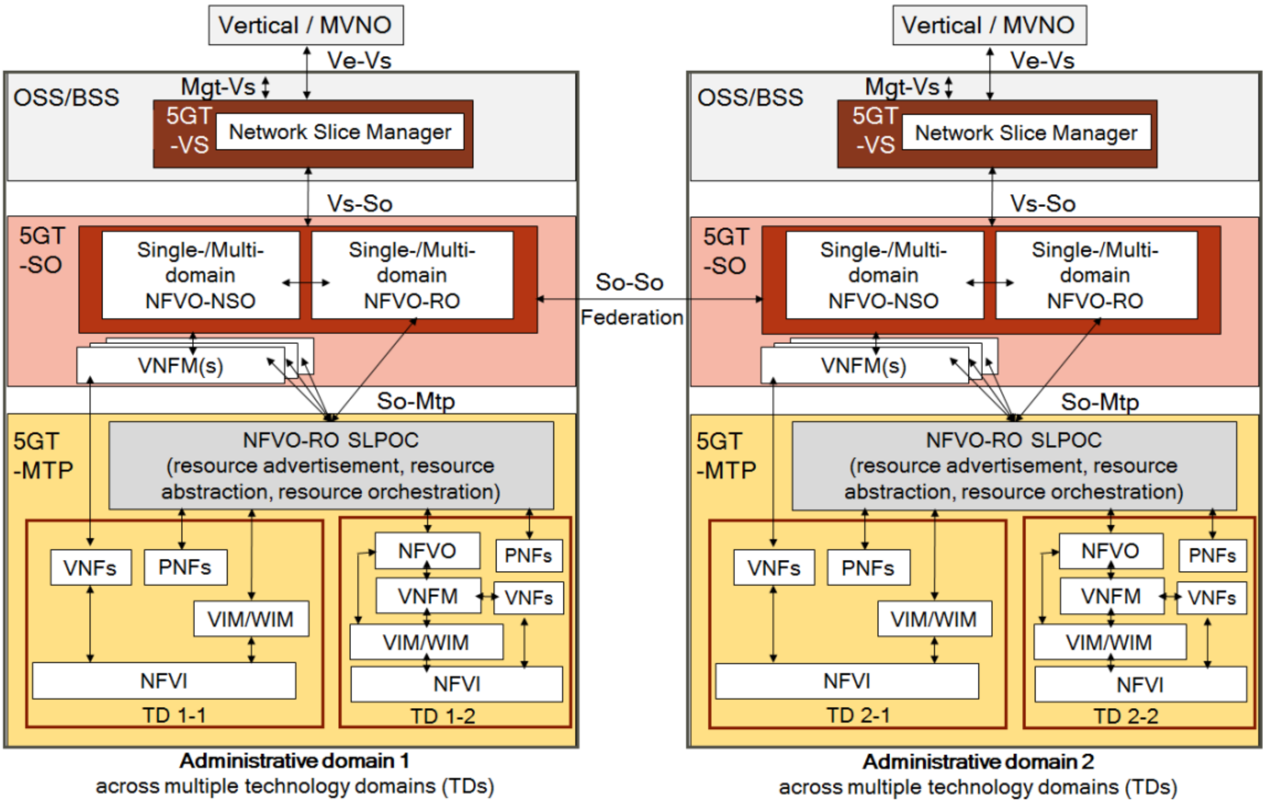

Figure 9 5G-TRANSFORMER architecture

On the other hand, the VS aims at facilitating the specification, instantiation, monitoring and management of vertical services through network slices, and it creates and maps the services onto network slices according to the verticals' requirements, managing their lifecycle. In this case, the VS can be considered as valuable reference baseline for the 5G EVE interworking framework (and possibly the 5G EVE portal) in terms of network slicing features and lifecycle workflows. Following the VS approach, the 5G EVE interworking framework (or the 5G EVE portal) could translate the vertical experiments, service and slicing requests into NFV NSs to be deployed in the proper 5G EVE sites by interacting with the dedicated per-site orchestrators.

2.2.2 5GinFIRE

The concept of 5GinFIRE[6] is very similar in many elements to that of 5G EVE. The main objective of 5GinFIRE is to provide a multi-site testbed for NFV experimentation, experiments in this case determined by the result of Open Calls. The initial testbed was composed of four sites, 5TONIC, University of Bristol, IT-Aveiro and the University of Uberlandia, at Brazil. Through the successive Open Calls, five new sites have been (or are being) integrated. Figure 10 below represents the status of the 5GinFIRE testbed in January 2019, as described in their public Deliverable D4.2 "Intermediate Report on the MANO Platform" [26].

[6] https://5ginfire.eu/

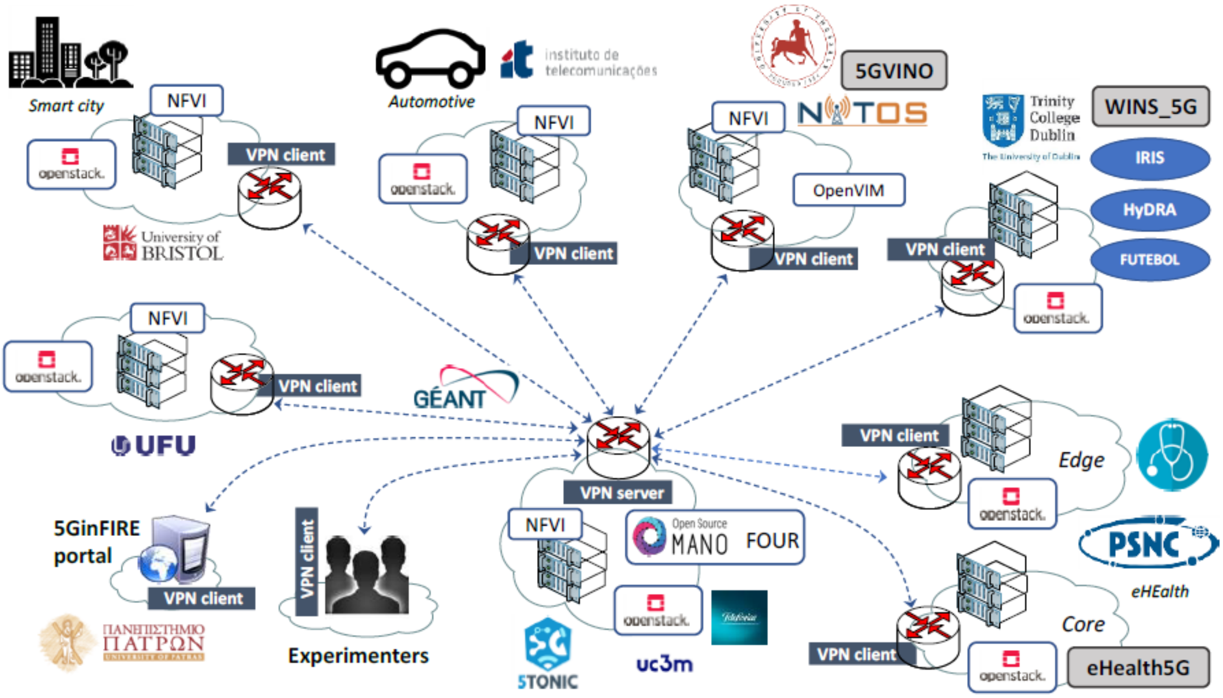

Figure 10: Overview of the 5GinFIRE testbed in January 2019

When compared with the Interworking Framework (I/W FW) proposed in 5G EVE, a few similarities and differences are worth mentioning, as summarized in Table 2.

First of all, there are important aspects where 5G EVE clearly differs from 5GinFIRE, the most important one being the Orchestration approach. In 5GinFIRE, there is a single Orchestrator, which interacts with the local VIMs deployed at each of the project sites. Instead, in 5G EVE each site already counts with its own NFVO component orchestrating the local VIMs. This provides sites with more capabilities, but is an added complexity for the I/W Framework since the interface between Orchestrators (proposed in 5G EVE for the Multi-site NSO SBI) is just starting to be standardized [12] at ETSI, and is not supported by available implementations. Therefore, extensions will be required.

Other important differences are the lack of global SDN support at 5GinFIRE and that testing/validation is not automated: Open Call experimenters run their experiments by connecting (via SSH) to their VNFs in the management plane, to which they access using a centralized VPN service. This has an impact on the direct reusability of 5GinFIRE concepts for the Runtime Configurator module and the KPI analysis and validation related part of the Data Collection Manager.

On the other hand, the main component that 5G EVE will be able to globally adapt, and apply lessons learnt, will be the Platform monitoring scheme[7]. All Control Plane components in 5GinFIRE are continuously monitored, with alarms (Bugzilla tickets) being generated each time any of them is not accessible. There are also VNF-based tools to measure the inter-site performance, and OSM Monitoring capabilities are being currently deployed. This makes that, even in a best-effort support scheme, there are many components reaching +95% availability figures. All monitoring tools are open source.

Finally, regarding local sites, 5TONIC is present in both projects, so all the local features deployed within 5G EVE will be built on top of the infrastructure which is already available for 5GinFIRE (e.g., OSM as NFVO).

[7] The Web Portal and associated components are also applicable to 5G EVE, but these are covered by WP4.

2.2.3 5G-MEDIA

5G-MEDIA[8] is 5G-PPP Phase 2 project that aims at innovating media-related applications by investigating how these applications and the underlying 5G network should be coupled and interwork to the benefit of both. The 5G-MEDIA framework leverages on SDN/NFV principles and technologies to offer an integrated Service Virtualization Platform (SVP) capable of handling the lifecycle management and monitoring of media applications in 5G networks. The project's main goal is to deliver an integrated programmable platform for the design, development and operations of media applications. This is achieved by providing mechanisms to flexibly adapt services to network changing conditions, and react upon them (e.g. to transparently accommodate auto-scaling of services, VNF re-placement, media applications re-configuration, etc.).

The 5G-MEDIA platform follows a layered architecture that goes beyond the ETSI NFV MANO architecture, and is composed by the following main building blocks (as depicted in Figure 11):

i. An SDK tool for developing, validating, profiling and testing media services and applications, following a DevOps approach.
ii. A generalized 5G Apps and Services Catalogue for validating and storing heterogeneous 5G applications and service descriptors, and translating them in NFV technology specific descriptors.
iii. An SVP including an ETSI NFV MANO framework (ETSI OSM) integrated with a Monitor Adapt Plan and Execute (MAPE) component for automated lifecycle management of media services applying cognitive optimization logics based on collected and aggregated monitoring data.
iv. Heterogeneous Network Function Virtualization Infrastructures (NFVIs), administrated by different types of VIM (e.g., OpenStack, OpenNebula and OpenWhisk), providing computing resources by different operators and supporting different cloud technologies, including server-less computing.

Despite its focus on 5G media vertical applications and service orchestration, the 5G-MEDIA platform has still some common points with the 5G EVE interworking framework due to its NFV-based approach. While it mostly targets single-domain 5G media services deployment and operation as NFV network services orchestrated by a centralized ETSI OSM orchestrator, 5G-MEDIA is developing some 5G management tools that are relevant to 5G EVE. Among these, the 5G Apps and Service Catalogue is indeed a generalized catalogue that allows service customers to bring their own VNFs and services into the 5G management and orchestration framework, with unified ETSI NFV standard format for descriptors and MANO specific translation towards ETSI OSM and other NFV-MANO stacks. In particular, the 5G-MEDIA catalogue can act as a centralized catalogue on top of several domain- and technology-specific catalogues, to support on-boarding and exchange of slice, service and function descriptors from multiple domains. Therefore, this 5G Apps and Service Catalogue can be considered as a valuable reference baseline for the 5G EVE Interworking Framework Multi-Site Catalogue.

[8] http://www.5gmedia.eu/

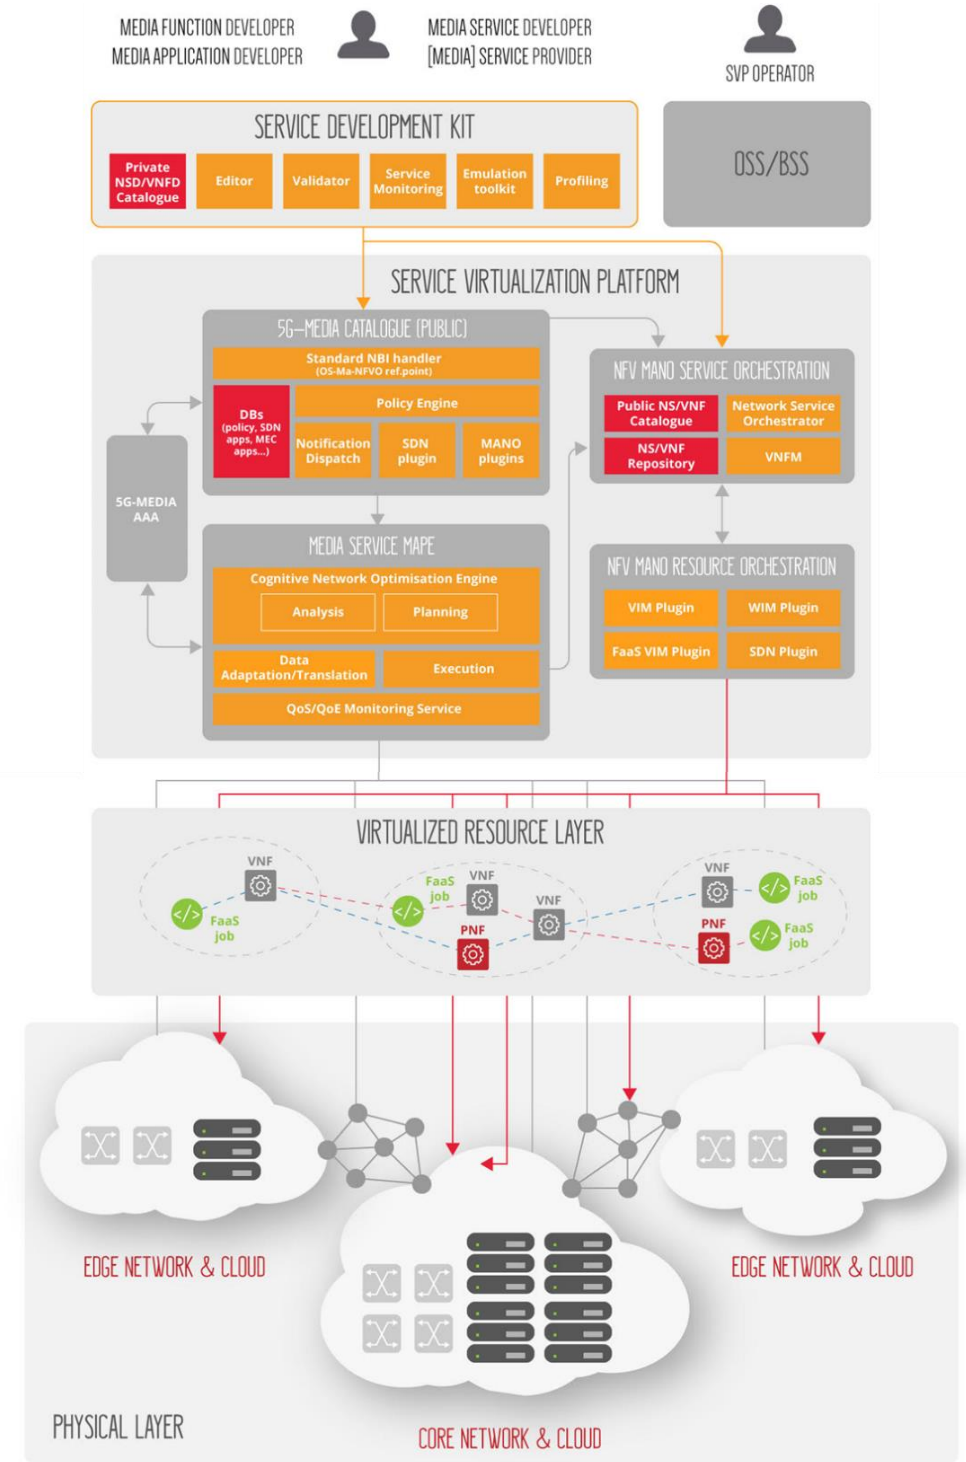

Figure 11: 5G-MEDIA architecture

2.2.4 Other Projects

This section reviews the state-of-the-art of other 5G architectures oriented for specific issues and challenges.

2.2.4.1 5G-Exchange

5GEx[9] is a 5G-PPP project aiming at providing an exchange mechanism facilitating and accelerating the deployment of network services by operators [27]. The project especially focuses on services spanning across multiple domains under the same or multiple administrations and heterogeneous technologies. Therefore, the rapid creation of reliable and controlled multi-domain services is the main business incentive of the 5GEx, and it forms a good basis for the connectivity services to be offered by 5G EVE. The project has established a list of use cases categorized into three families: connectivity, VNF as a Service (VNFaaS) and Slice as a Service (SlaaS). Connectivity may either target core services or directly end customers. VNFaaS provides a customer with a VNF. SlaaS is a combination of connectivity and VNFaaS.

The Multi-domain Orchestrator (MdO) is the central component of the 5GEX architecture, although in this case it is composed of the NFVO and the VNFM of the ETSI NFV MANO framework. The architecture also describes a catalogue containing all the necessary information regarding the on-boarded NSs and VNFs. The NFVO is split into Network service Orchestrator (NSO) and Resource Orchestrator (RO) entities. The NSO is responsible of the NSs; it uses the catalogue to translate the received NS requests into subsets of resources, VNFs and NSs. The RO manages the domain orchestrators, which in turn are responsible for the resources of their domains. When the RO receives a resource request from the NSO, it maps that request to a set of domains, and reserves a quota of resources.

Although the orchestration architecture is not exactly equivalent to that of 5G EVE, the workflows between the Multi-Site NSO and Multi-Site Catalogue can be of use for the project.

2.2.4.2 5G!Pagoda

5G!Pagoda[10] is a Japanese-European project that aims to develop a 5G network based on the network slicing concept [28]. The proposed architecture especially focuses on multi-domain orchestration and slice dynamicity. The project has made a deep description of the different building blocks of a slice, which could be an input for the slice definition at 5G EVE.

Regarding use cases, it especially focus on the mobile network aspect, with specific consideration for the RAN slicing, concept to be integrated also at 5G EVE. Besides the ETSI NFV MANO framework, 5G!Pagoda includes additional building blocks, namely: the Business Service Slice Orchestrator (BSSO), the Multi Domain Slice Orchestrator (MDSO), the Domain Specific Slice Orchestrator (DSSO) and the Policy Based Management (PBM).

Of particular interest is the MDSO, which performs the slice placement. It receives slice creation requests from the BSSO, and uses information regarding resources and cost obtained from the DSSOs to make the placement decision. The DSSO is responsible of single domain slices that it creates when the MDSO has made the placement decision. This module runs on top of the domain NFVOs.

Finally, the PBM is an element that runs across all layers. It receives policies from the network administrator, and applies them to the system. To do so, the PBM monitors various elements, collecting and analysing metrics, and reacting according to its defined policy (similarly to the Data Collection Manager).

2.2.4.3 5G NORMA

5G Norma[11] is a 5G-PPP project that aims to develop an adaptive 5G network architecture enabling network customization and ensuring the tight performance, security, and energy requirements to be met. For this purpose, a novel concept of network control has been introduced, extending the software-defined routing approach to all kinds of mobile NFs (both from data and control layers), with a focus on wireless control functions. Two of its four architectural layers may be of interest for 5G EVE.

[9] http://www.5gex.eu/

[10] https://5g-pagoda.aalto.fi/

[11] http://www.it.uc3m.es/wnl/5gnorma/

The MANO layer is an extension of the ETSI NFV MANO framework towards multi-tenant and multi-service networks. It provides the Software-defined Mobile network (SDM) Orchestration (SDMO) concept. Besides the MANO entities, this layer includes two components; the Inter-Slice Resource Broker (ISRB) and Service Management. The ISRB determines and enforces policies for cross-slice resource allocation, particularly in the case of shared network functions, while the Service Management mediates the service layer with the ISRB, transforming consumer-facing service descriptions into resource-facing service descriptions and reporting about selected KPIs.

The Control Layer is represented by two controllers: the SDM Coordinator (SDM-X) and the SDM Controller (SDMC). The former is used for the control of common (shared) NFs, while the latter is for dedicated NFs. The two controllers abstract from the technology and implementation-related details of controlled network functions. They translate the application control decisions into commands towards VNFs and PNFs.

2.2.4.4 SliceNet

SliceNet[12] is a 5G-PPP project that aims to develop an orchestration and management stack for 5G networks. The solution envisions to use Artificial Intelligence (AI) to enhance the efficiency of the management plane, concept which is not specifically included in 5G EVE. However, there are other components of its architecture which have a clear relationship.

In particular, the SliceNet information sub-plane is a set of inventories and catalogues that store the data required by other sub-planes. Also, the monitoring sub-plane disseminates a set of sensors to measure relevant indicators (determined by the cognition sub-plane) at service, slice, and resource levels. From both of these sub-planes 5G EVE can obtain valuable experience.

The cognitive management sub-plane in SliceNet concentrates all the AI that is put into its MANO system. It collects and analyses data gathered by the monitoring sub-plane to optimize the network by selecting the policies to apply. The AI to implement for this sub-plane relies on the CogNet [29] project.

The orchestration sub-plane is made of two main components: the Slice Service Orchestrator (SS-O) and the Resource and multi Network segment Orchestrator (NMR-O). The SS-O takes care of end-to-end services and multi-domain slices, while the NMR-O orchestrates slices and sub-slices over a single administrative domain. The SS-O enriches the Service Templates (STs) and turns them into Service Descriptors (SDs). Then, the SDs are divided into several resource oriented Network Slice Templates (NSTs), which are sent to the corresponding NMR-Os. The NMR-O architecture is closely related to ETSI NFV MANO framework that includes a Mobile Edge Application Orchestrator (MEAO) for Distributed Cloud/MEC applications.

2.2.4.5 5Gtango

5Gtango[13] is a 5G-PPP project enabling the flexible programmability of 5G networks. The goal is to develop network slicing features as part of a Service Platform (SP), including the development and deployment of complex 5G services empowered by SDN and NFV technologies. The 5Gtango architecture is based on two tenets: (i) the support of flexible verification & validation (V&V) processes addressing different verticals and NFV platforms, as well as different roles in the design and operational workflow of NFV services; (ii) the use of micro-services concept to separate the concerns, giving single responsibility per entity. SDN and validation topics at 5Gtango can definitely represent a good starting point for 5G EVE.

More concretely, the V&V platform provides a verification and validation service that dissects and profiles the submitted VNFs and NSs to ensure they pass a range of tests. The V&V is composed of three central components (Test Invoker, Test Engine and Test Platform Manager) connected to a public Graphical User Interface (GUI) Application Programming Interface (API) and V&V integration, interacting with a NS catalogue, test repository, and referencing a test result repository.

[12] https://slicenet.eu/

[13] https://5gtango.eu

The SP handles the management and orchestration of NSs and provides the needed complementary and supporting features such as ensuring resource isolation (i.e., network slicing management), policy and SLA management, and user access management. Focussing on network slicing, the SP is able to create NSTs, instantiate, scale, heal, and terminate multiple Network Slice Instances (NSIs). The Slice Manager, a component of the SP, is responsible for the management of network slices. It maps between the NSIs and the NSs via a Slice2NS Mapper, and manage the lifecycle of the NSI by a Slice Lifecycle Manager. The NSTs and the running NSIs are stored in a Catalogue and the NSI repository.

2.2.4.6 5G MoNArch

5G MoNArch[14] is a 5G-PPP project addressing the resilience and elasticity of network slices, from which 5G EVE can also leverage. The 5G MoNArch architecture was designed in modular manner: it depicts several layers, each of which defines the architectural elements that the system should deliver, including the functional elements, the exposed interfaces, and the interactions between them.

For example, the Service Layer comprises both the BSS, the business-level policy, and the decision functions in one hand. On the other hand, this layer holds the applications and services operated by tenants and other external entities. All these functions interact with the Management & Orchestration layer via the Service Management function, work which can be applicable to the North-Bound Interface of the 5G EVE Interworking Framework.

Another example is the Network Layer, which comprises the VNFs/PNFs needed to control and carry the user data traffic, including the Access and Mobility management Function (AMF), Session Management Function (SMF), and Authentication, Authorization, and Accounting (AAA) for the control plane, and S/P-GW, and UDP user plane function for the user plane. Some of these could also be the basis for supporting services of the 5G EVE facility, additional to those defined in section 6.

2.2.4.7 5G ESSENCE

5G ESSENCE[15] is a 5G-PPP project addressing the paradigms of Edge Cloud computing and Small Cell-as-a-Service to provide a highly flexible and scalable 5G small cell platform, able to support new business models and revenue streams. The project is built on the distributed and network-integrated Cloud inherited from various 5G-PPP projects, particularly the SESAME[16] project. It aims to decouple the control plane from the user plane of the RAN to get benefits from the Cloud-RAN and avoid the front-haul latency.

The architectural design for the 5G ESSENCE is also based on the SESAME architecture. It combines the 3GPP framework for network management in RAN sharing scenarios and the ETSI NFV MANO. This architecture is designed in a way that multiple network operators are able to provide services to their clients via a set of Cloud-Enabled Small Cell (CESC) deployed, owned, and managed by the CESC provider. In this way, operators can extend their 5G RAN capacity.

The main 5G ESSENCE result which could be reused in 5G EVE is its capacity to provision the Edge Datacentre with Distributed Cloud/MEC capabilities to enhance the user experience and the agility in the service delivery. This centralized control of Distributed Cloud/MEC by the Interworking Framework is yet another topic that could be valuable in the future for 5G EVE.

2.2.4.8 MATILDA

[14] https://5g-monarch.eu/

[15] http://www.5g-essence-h2020.eu/

[16] http://www.sesame-h2020-5g-ppp.eu/

MATILDA[17] is a 5G-PPP project having a vision of holistic 5G end-to-end services operational framework, which tackles the design, development, and orchestration lifecycle of 5G network services and 5G-ready applications over programmable infrastructure, which is in fact quite similar to the objectives of 5G EVE. The project aims to devise a software for 5G-ready applications as well as VNFs/PNFs and NSs by using a unified programmability model, a proper abstractions definition, and a creation of an open development environment. The architectural design of the MATILDA project is done over four layers from the top-down.

The Application Layer corresponds to the business service and business function layer. It considers the design and development of 5G-ready applications, along with the associated networking requirements, per vertical industry. These requirements associated with the respective 5G-ready applications' graph define the business functions and service qualities of an individual application.

The Orchestration Layer includes the application components, VNFs, and the set of networking resources needed to chain the network functions together to provision a complete NS package. The layer supports dynamic, on-the-fly, and optimal deployment and adaptation of the 5G-ready applications to its service requirements, as well as strategic placement, runtime policies enforcement, data mining, analysis, and context awareness support for holistic 5G end-to-end network services.

The Network Functions and Resource Management Layer positions network functions and resources closer to the Orchestration Layer, hence having almost a native impact on the virtual and the physical resources used by the orchestrator to provision the 5G end-to-end network slices. This layer includes the implementation of the resource management functionalities over the available programmable infrastructure and to the lifecycle management of the activated VNFs. It uses both physical and abstracted virtualized resources, along with VNFs and PNFs, to deliver application-aware network slices.

Last, the Infrastructure Layer consists of the data communication network spanning across a set of Cloud computing and storage resources such as C-RAN, Edge/Fog computing resources, Edge network, and Transport/Core network.

[17] http://www.matilda-5g.eu/

3 5G EVE Interworking Services

This section describes the Interworking Services that the 5G EVE platform must support. "Interworking Services" is a concept that is not strictly related with connectivity services, although these are of course included; "Interworking Services" are also the access to KPIs or the automated deployment of a VNFs for experiments that may, or may not, expand over multiple sites.

As already stated, the I/W Framework is an abstraction layer in 5G EVE which is present independently of whether experiments are single-site or multi-site. In both cases, this layer is hiding the implementation details of each site from the global Web Portal. The main difference is that multi-site experimentation will require some additional functionalities which are not required for single-site.

Two main blocks are covered in this section. First, the high-level capabilities offered by the platform to experimenters, which is a global definition divided in capabilities required for single-site, and extensions needed to implement multi-site. Second, the connectivity services that will support the interconnection of the VNFs of an experiment will be also be described. Although the final goal of 5G EVE is that these connectivity services are "multi-site slices", alternative simpler solutions will have to be implemented to support experimentation until the former are totally supported by the platform.

3.1 Interworking Framework Services

The main services offered by the Interworking Framework for experimentation can be classified in three categories: connectivity, operation and monitoring. There are different degrees of supported features in each of them, depending on the maturity of the developed solution but also on the type of experiment.

3.1.1 Single-site scenarios

The following table summarizes the features for single-site scenarios.

Table 3: Interworking Framework key features for single-site scenarios

Key features	Category	Brief Description
Orchestration Plane Interworking	Connectivity	Low bandwidth performance but secure connectivity among sites for orchestration traffic.
Single-site Experiment Monitoring Support	Monitoring	Capability of translating the monitoring requirements defined by experimenters (based on selected KPIs) to the requested site[18]. Sites will typically have different local monitoring tools and mechanisms.
Single-site Applications Deployment Support	Operation	Capability to deploy the required VNFs, hosted in the 5G-EVE Catalogue, at the requested site. Sites will typically have different local orchestrators.
Single-site Network Automation Support	Operation	Capability to deploy the required connectivity services (first phase) and slices (second phase) to the requested site. Sites will typically have available different local controllers and network infrastructure.

[18] The term "requested site" is used in this section to define "the site where the experiment will be deployed".

3.1.2 Multi-site scenarios

Equivalent to the previous section, the following table summarizes the features for multi-site scenarios. It is worth mentioning that multi-site scenarios require of the capabilities of single-site scenarios to work; Table 4 includes only the incremental features.

Table 4: Interworking Framework additional features for multi-site scenarios

Key features	Category	Brief Description
Control Plane Interworking	Connectivity	Low bandwidth performance but secure connectivity among sites for control traffic.
Data Plane Interworking	Connectivity	Secure connectivity among sites for user traffic. Low bandwidth performance experiments will employ best effort connectivity. High bandwidth performance experiments will employ a parallel high bandwidth low latency network, which will be available at least between two sites.
Multi-site Experiment Monitoring Support	Monitoring	Capability of translating the monitoring requirements defined by experimenters (based on selected KPIs) to the sites taking part in the same experiment. Sites will typically have different local monitoring tools and mechanisms.
Multi-site E2E Orchestration Support	Operation	Capability to deploy the required slices, and VNFs hosted in the 5G-EVE Catalogue on top of them, to the sites taking part in the same experiment. Sites will typically have different local orchestrators, controllers and network infrastructure.

3.2 Data Plane Connectivity Services

Support for slices in 5G EVE is not in the roadmap for the first version of the Interworking Framework, mainly due to the huge implementation effort that such feature demands (even in single-site experiments). It is true that some partial slicing capabilities may be already available (e.g., slices in the RAN, or slices in the datacenter infrastructure, at one or several sites), but limiting the "slice" concept to those specific scopes is not the goal of 5G EVE. Therefore, the support of slicing is expected for Deliverable D3.4 (Month 24) – *Second implementation of the interworking reference model* –, which provides the fully-capable fully-tested software version of the I/W layer.

Meanwhile, the data plane among VNFs in an experiment will be built using more traditional datacenter mechanisms like VLANs or L2/L3 VPNs. These are briefly described below.

3.2.1 Traditional services for single-site scenarios

At the time of writing of this Deliverable, as stated in section 1.1.1, there have been no updates to D3.1 concerning the interworking requirements. This way, the project will concentrate on the implementation of single-site data plane connectivity solutions, like the ones mentioned above.

The expectation, in any case, is that when multi-site Use Cases are finally demanded and ready to be executed, 5G EVE is ready to provide multi-site slices. That is the reason why we are not considering multi-site VPNs, for example, in this document. If this changes, these assumptions will be updated in future WP3 Deliverables, mainly D3.3 (M16), the first version of the I/W layer.

Under these considerations, the simplest scenario that could be deployed in any of the sites would be the usage of VLANs to provide Layer 2 connectivity among VMs (Figure 12, which includes the commonly accepted Leaf & Spine fabric).

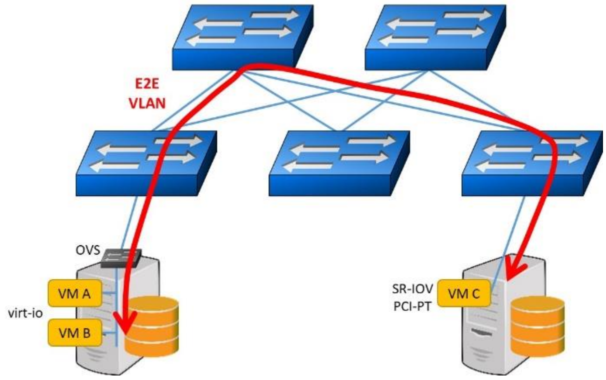

Figure 12: **Single-site VLAN connectivity**

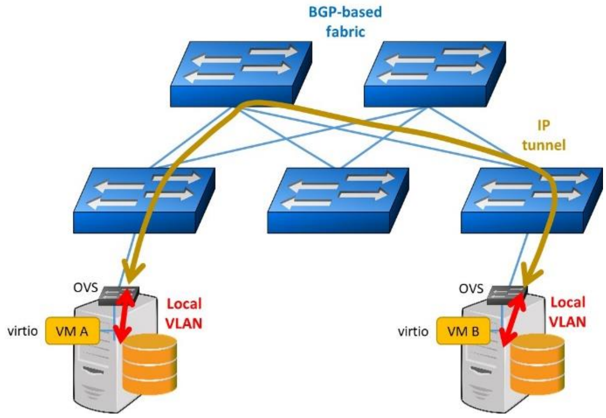

Figure 13: **Single-site VxLAN connectivity**

In the most general case where VMs can be based either on virt-io (and connected to the fabric via an Open Virtual Switch – OVS), like VMs A and B in Figure 12, or on SR-IOV/PCI-PT, like VM C, the VLAN will have to expand across all the infrastructure.

This type of solution was very typical in the first datacentres since they were not very complex, and there were few technical alternatives. The reality today is that datacentres based on VLANs are not very common in production networks, due to some limitations and operational issues. However, and this is why this approach needs to be covered in 5G EVE, it is still used at small locations, normally dedicated for testing, like the 5G EVE sites.

Lately, fabric stages have migrated to support L3 protocols in the underlay. The improved scalability, together with the reduction of prices for basic L3 features, have made the BGP-based Leaf & Spine architecture very popular. Furthermore, such a L3 fabric permits the provision of overlay services among OVS of different servers without the need to make any configuration at the fabric itself (see Figure 13 above, where both OVS have L3 reachability, and can therefore set up an IP tunnel), which is an important operational advantage.

Depending on the type of end-to-end service, different tunnelling technologies exist: normally GRE tunnels will be used to encapsulate L3 over L3, while VxLAN will be used for encapsulation of L2 over L3. A very good advantage, however, of VxLAN is that it supports E-VPN, which equally permits the provision of L2 or L3 VPN services (e.g., for VMs needing SR-IOV or PCI-PT interfaces, normally required by applications which are intensive in the data plane). This common support homogenizes the provisioning solution.

In 5G EVE, and according to the capabilities of the different facilities, all these approaches will be supported. However, it is expected to provide only VLAN support in the first phase, as this is supported by all sites.

3.2.1.1 Service provisioning automation

Regarding automation in the provisioning of any of the services above, it must be noted that current VIMs like Openstack can actually automate the provision of connectivity services for virt-io based VMs, by configuration of their associated OVS. This is supported both using internal components like Neutron, and using plug-ins to interact with SDN controllers like ONOS or ODL.

However, VIMs cannot directly interact with the fabric using Neutron. Therefore, for VMs using SR-IOV or PCI-PT, the only solution for automation is the usage of an SDN Controller. The NBI of such SDN Controllers may be typically facing the VIM itself (e.g., with plug-ins as in the previous case), or the NFVO (like in the "SDN Assist" mode of ETSI OSM).

These, in fact, are the approaches currently followed by the 5G EVE sites, which explains why the local SDN Controllers are not be exposed to the I/W Framework (they are only exposed to local modules of the MANO stack). This way, the automated provision of connectivity services among VMs at single-site scenarios will entirely depend on each site's capabilities.

3.2.2 Multi-site slices

The concept of slicing, and mostly the “extension” of the concept, varies depending on the context where it is being used: datacenter slice, RAN slice, WAN slice, etc. In 5G EVE, a specific working group has been defined to elaborate on what exactly Multi-site Network Slicing means for the Project. The activity of this group, which is critical for the evolution of the Interworking Framework, has not yet been finished. Therefore, only a preliminary outline is included here, which will be updated in the future.

In 5G EVE, multi-site experimentation is based on the concept of multi-site Vertical slices, which span across multiple 5G EVE sites. The Vertical slices definition relays in the Interworking Layer capabilities for defining an end-to-end slice that satisfies the requirement of the Vertical. The Interworking Layer provides an abstraction layer of the 5G EVE sites and it builds the Vertical slices using the capabilities that each 5G EVE site provides, as well as the interworking capabilities provided by the Interworking Framework itself. With this approach, 5G EVE project decouples the definition of the end-to-end Vertical slice from the specific solution used in each site, allowing different approaches and vendor solutions.

Although the main idea behind the multi-site slice is to allow the interconnection of any VNF from any 5G EVE site, we have established some restrictions due to operation and maintainability reasons. Each 5G-Site network is operated as an independent network associated to a specific operator, following the real deployments that we will find in a near future. For this reason, we will support the following composition of multi-site deployments:

- Access Network: the complete Access Network will belong to a specific 5G-Site. The end-to-end Vertical slice could include Access Networks from different sites but cannot define an Access Network expanded through several sites.
- Core Network: the core network can span through two sites using the standard roaming deployments, i.e., Local Break-Out or Home Routed, according to 3GPP TS 23.501 [30].
- VNF in data network: VNFs belonging to the Data Network can be deployed in any of the 5G EVE sites without any restriction.

4 Interworking Reference Model

The joint operation of 5G multi-site facilities requires a homogeneous definition of procedures and services, plus some level of adaptation to the different environments composing the overall experimentation environment. This way, the complexity derived from the specific implementations in each site is hidden to the Verticals executing experiments on top of our infrastructure.

The main service offered by 5G EVE to Vertical customers is a Network Slice providing a virtual, logical and isolated network with tailored requirements over an accessible network infrastructure. All customers' Slices will coexist in a transparent and isolated mode, even with very distinct requirements.

Slice provisioning, scaling, assurance, or decommissioning are part of the actions that must be supported by the overall experimental facility. In this section we present the multi-site interworking framework being defined in the 5G EVE project. The design starts from single-site scenarios and builds additional capabilities on top to account for multi-site particularities, when they apply.

4.1 Design criteria

In order to minimize excessive integrations and particularities, and guarantee a future-proof specification, the Interworking (I/W) Framework design will be based on existing standards, even if they are in a draft status. That approach permits to have a clear reference for the functional needs of the different internal components, and to leverage on industry-lead advances, guaranteeing future adoption by the market. It also allows for an independent evolution of the local implementations, which could evolve along time under a stable and well-defined upper layer (i.e., the I/W Framework's South-Bound Interface).

Another important part of the design strategy is the adoption, when feasible, of standard interfaces for the interaction between components. Whenever those interfaces do not cover all the needs of 5G-EVE, then some extension or additional elaboration is required to adapt them to the purpose of the project. In those cases, such extensions could be potential subject of contribution from the project to the state-of-the-art, e.g. by contributing to the corresponding SDOs.

4.2 Interworking Framework architecture

This section describes the architecture and components of the I/W Framework. Essentially, these have not changed from those outlined in Deliverable D3.1 [1] and depicted in Figure 1 of the present document. This section, therefore, extends those descriptions and provides initial insight about the relationships among the different components (which will be further covered in section 4.3).

A key objective driving some of the architectural decisions is trying to simplify the integration efforts towards the sites. In that sense, there are components, like the Multi-site NSO for example, which will undoubtedly interact with local components. However, this can be argued for the Multi-site Catalogue and Inventory. The following discussion is based on this subset of components of the I/W Framework, depicted in Figure 14.

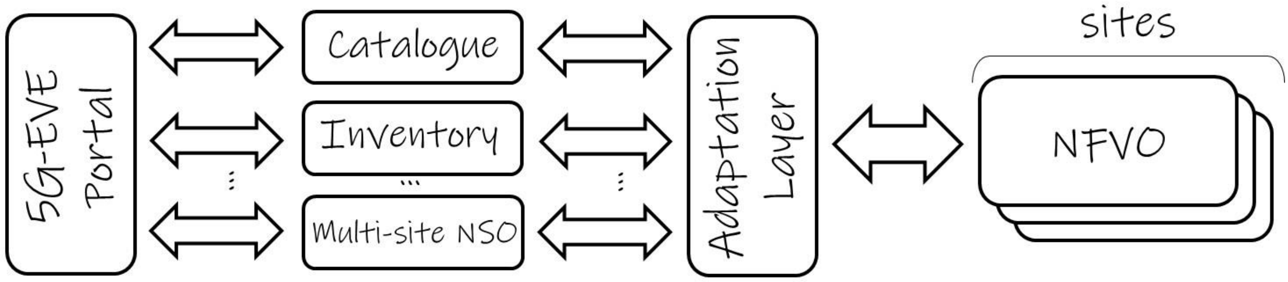

Figure 14: Original approach for interworking with 5G EVE sites

The Adaptation Layer has the main goal to abstract the heterogeneous capabilities and APIs exposed by the 5G EVE site facilities. This layer offers to the components of the I/W Framework a set of common internal APIs and models for accessing per-site management, control, orchestration and monitoring services in a transparent way. Additionally, it translates those common APIs and models into site-specific APIs and models, through drivers providing the required translation. While the APIs could represent syntactic differences among sites, the functional behaviour will be the same; in other words, the semantic capabilities of those APIs will be the same independently of the implementation.

Due to the normally supported capabilities of deployed NFVO components at sites, two different models can be considered with regards to the Multi-Site Catalogue and Inventory relationship with the Adaptation Layer:

- Option A: all components independently interacting with the local NFVOs via the Adaptation Layer (Figure 15).

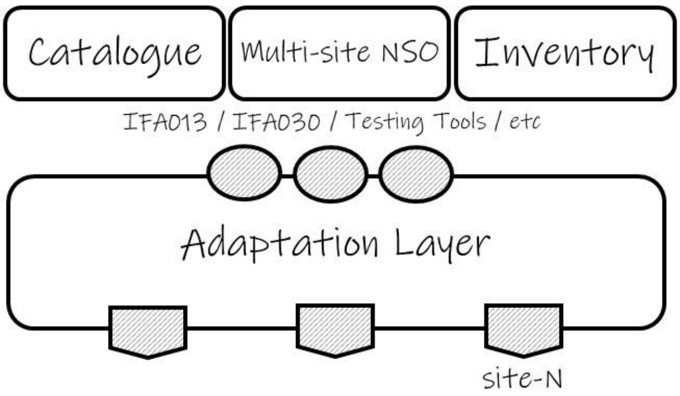

Figure 15: Direct interaction of Multi-site Catalogue and Inventory with sites

In this case, the Adaptation Layer has to incorporate logic for the interaction of all the components with their local counterparts. The drivers per site should be enriched to support all the transactions from multiple components, then making the Adaptation Layer more complex, although more flexible.

- Option B: Multi-site NSO acts as a proxy for the interactions with the local NFVOs (Figure 16).

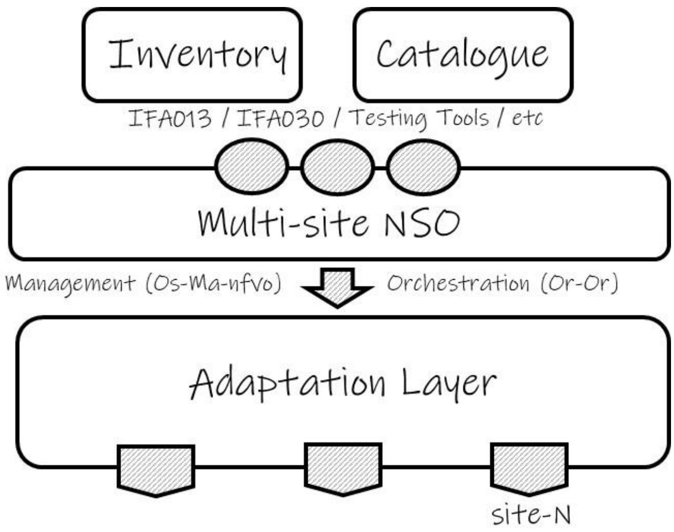

Figure 16: Interaction via proxy of Multi-site Catalogue and Inventory with sites

In this other case, the Multi-site NSO is the component absorbing all the complexity, offering a single point of contact for whatever interaction is needed. This approach, therefore, makes the Adaptation Layer much simpler, making it a lightweight component centred on effective translation between orchestration environments. This can provide stability to the interaction among sites and the centralized I/W Framework, as well as simpler evolution pace for components (less impact).

None of the options is clearly superior to the other in a scenario like that of 5G EVE, with multiple sites, each one implementing its own orchestration mechanisms and deploying its own local catalogues. In fact, the decision could be even different depending on the site. For example, if the local NFVO supported catalogue-related operations, option B could be the most appropriate. However, a commercial NFVO with no support of certain feature could be very difficult to upgrade, that leading to a scheme like that of option A being better. It could also be possible that each option demonstrated to be better for one of the components.

To account for this, both schemes have been considered in this Deliverable:

- All the workflows, and the SBI, are specified based on a scheme like that of option B.
- The solution in case one component demanded option A would be to replicate the North-Bound Interface in the Adaptation Layer, exclusively for the requiring component.

4.2.1 Multi-site Catalogue

The 5G EVE Multi-site Catalogue is a core functional component within the I/W Framework. It is responsible for storing and maintaining all the information related to the available Networks Services (NSs) and VNFs that can be used for the provisioning of single-site and multi-site experiments. In practice, the Multi-site Catalogue decouples the Network Service Descriptors exposed to the 5G EVE Portal (modelling single-site and multi-site services and slices) from the Network Service Descriptors and VNFs collected from each of the 5G EVE site facilities (modelling the actual capabilities of each site). In summary, the Multi-site Catalogue maintains the following information:

- *VNF Descriptors*, collected from site facilities and exposed to the 5G EVE Portal. These could include vertical-related VNFs (on-boarded directly on the 5G EVE site facilities), as well as network-related VNFs (e.g. related to 5G Core or EPC) and virtualized testing/experimentation tools.
- *PNF Descriptors*[19], collected from site facilities, modelling non-virtualized testing/experimentation tools, as well as regular network functions. These are then exposed to the 5G EVE Portal for their usage in single- and multi-site Vertical experiments.
- *Network Service Descriptors*, collected from site facilities and exposed to the 5G EVE Portal. These descriptors model existing available services in the sites, as a composition of VNFs that are offered to the Interworking Layer for their usage in the context of single- and multi-site Vertical experiments.
- *Network Service Descriptors* that are on-boarded by the 5G EVE Portal, as a result of the definition and composition of single- and multi-site Vertical experiments as a collection and interconnection of existing descriptors from the three categories above.

The information stored in the Multi-site Catalogue has two main purposes. On the one hand, it enables the 5G EVE Portal to retrieve the available capabilities in the end-to-end 5G EVE infrastructure (i.e. within all the site facilities) in terms of Network Service Descriptors and VNFs to be composed for the deployment of Vertical experiments. On the other hand, the Multi-site Catalogue maintains the relevant information that the Multi-site NSO uses to manage the lifecycle of single- and multi-site Network Services, including which local orchestrators have to be involved for end-to-end provisioning of the Vertical experiments. In particular, the Network Service Descriptors on-boarded in the Multi-site Catalogue by the 5G EVE Portal include the references to the individual local components of the end-to-end services.

---

[19] Testing tools will be modelled in the Multi-site Catalogue as PNFs or VNFs, depending on their nature, so a special category is not required for them.

For these purposes, the Multi-site Catalogue implements the various Network Service and VNF on-boarding specific workflows described in section 4.3 for:

i. **Catalogues synchronization**, to let the Multi-site Catalogue to automatically collect available Network Service and VNF descriptors from each site facility.
ii. **VNF on-boarding**, to allow dynamic and automated retrieval of new VNF descriptors that are on-boarded directly in the 5G EVE sites at runtime.
iii. **VNF removal**, to allow dynamic and automated removal of existing VNF descriptors in the Multi-site Catalogue when the related VNFs are no longer available in the 5G EVE sites.
iv. **Network Service Descriptor on-boarding**, to allow the 5G EVE Portal to on-board in the I/W Framework new descriptors for single- and multi-site NSs modelling Vertical experiments.
v. **Network Service Descriptor removal**, to allow the 5G EVE Portal to remove existing descriptors whenever the related single- or multi-site Vertical experiment is no longer available in the Portal.

These workflows are enabled through a well-defined set of external interfaces, both at the northbound and southbound reference points (the latter initially via the Multi-site NSO, as per the discussion in section 4.2). In both cases, the Multi-site Catalogue leverages on the ETSI NFV information models and APIs for the catalogue and on-boarding related operations, in support of the required workflows.

In particular, at its northbound reference point the Multi-site Catalogue allows the Portal to retrieve the available NS and VNF Descriptors for their composition in support of single- and multi-site vertical experiments. Additional information related to UE and subscriptions available in the 5G EVE sites is also stored by the Multi-Site Catalogue and exposed to the Portal for selecting proper end-devices and related configurations for each experiment. It also offers the possibility to on-board new NS Descriptors. These operations and information models contribute to the Interworking API, and follow the ETSI NFV SOL005 [10] and SOL001 [4] specifications for on-boarding related APIs and packaging structure, as detailed in section 4.4.1.

At the southbound reference point, the Multi-site Catalogue follows the abstraction and adaptation principles discussed in section 4.2, that is, interacts with local components via the Multi-site NSO that acts as a proxy. Therefore, it supports the on-boarding related operations and information models defined in section 4.5.2.1 for a common and unified interaction with the different catalogue services exposed by the sites, irrespectively of the specific local orchestrator technologies.

4.2.2 Multi-site Inventory

The 5G EVE Multi-site Inventory is also a core functional component within the I/W Framework, conceptually similar to the Multi-site Catalogue: if the Catalogue is responsible for storing information about available Networks Services (NSs) and VNFs that can be used for the provisioning, the Inventory is responsible for storing information about the NSs and VNFs that have been effectively instantiated in any of the sites.

The Multi-site Inventory is fully managed by the Multi-site Network Service Orchestrator, who is in charge of notifying of all the changes that have to do with service provisioning. In fact, the Multi-site NSO will act as a proxy of any inventory-related operation towards the sites (although during the first stages of the project some manual operations are also envisioned, mostly regarding the PNF management).

It is worth noting that, on the northbound reference point, not necessarily the Experimentation Portal will be the entity interacting with the Multi-site Inventory. In principle, global information about instantiated slices and services will not be accessible to experimenters, so it is more normal that an Administrative Portal is the one interfacing the Inventory. This, of course, does not prevent the Experimentation Portal to provide specific information about their own services to experimenters.

The Multi-site Inventory maintains the following information:

- *VNF Instances* deployed and configured in the 5G EVE facility and running in one or more sites. These could include both Vertical-related VNFs and network-related VNFs (e.g., related to 5G Core or EPC), together with virtualized testing/experimentation tools.
- *PNF Instances* deployed and configured in the 5G EVE facility and used as part of Vertical experiments. These could include regular network functions as well as non-virtualized testing/experimentation tools.

- *Network Service Instances* deployed in support of Vertical experiments. These could include existing available services in the sites, as a composition of VNFs, and composite Network Services defined as a collection of local Network Services.

The information stored in the Multi-site Inventory has two main purposes. On the one hand, it enables an administration portal to present the current status of the 5G EVE facility to the operational team, which helps in the understanding and troubleshooting of the infrastructure. On the other hand, the Multi-site Inventory maintains the relevant information that the Multi-site NSO uses to make sure that any given site has the available resources for new deployments.

The amount of information (list of attributes) that the Multi-site Inventory will maintain with regards to the different Instances will be called "Level of Inventory" (LoI). The LoI may change in the future, if new demands appear, but the minimum set of attributes required within the context of 5G EVE is defined in Table 5.

Table 5: Multi-site Inventory minimum Level of Inventory

Multi-Site Inventory minimum Level of Inventory
NS Instance LoI
• NS Instance identifier • NS type • Operational status • List of VNF Instances part of the NS Instance (identifiers) • List of PNF Instances part of the NS Instance (identifiers) • NS Forwarding Graph • List of nested NS Instances
VNF Instance LoI
• VNF Instance identifier • VNF type • Operational status
PNF Instance LoI
• PNF Instance identifier • PNF type • Operational status

4.2.3 Multi-site Network Service Orchestrator

The Multi-site Network Service Orchestrator (NSO) is another core element of the I/W Framework and provides the single interface for orchestrating the multi-site network services. This interface is common for the single site deployments as well as the multi-site deployments.

The Multi-site Network Service Orchestrator allows to reference the information available in Multi-site Catalogue and Multi-site Inventory in the Network Service orchestration operations, in order to deploy or to use the NSs and VNFs stored in those components.

The services offered by the Multi-Site Network Service Orchestrator are:

- Network Service Identifier management: allows creating and deleting a Network Service in the I/W Framework. The creation of a Network Service references to a NSD previously on-boarded in the Multi-site Catalogue, and can incorporate existing VNF and PNF instances.
- Network Service management: using this service the 5G EVE portal can deploy a Network Service in the 5G-Site facilities. The deployment includes the instantiation of all the elements that are included in

the Network Service descriptor, as well as the Day0/1 configuration. The NS management also includes the operations of querying (information), scaling, updating and terminating a NS.

The northbound interface that supports these services is defined in section 4.4.3.1. The interface is mainly based on ETSI NFV-SOL 005, although it will be extended for including 5G EVE specific information fields.

4.2.4 Runtime Configurator

The Runtime Configurator allows to apply tailored runtime configurations to the provisioned end-to-end NSs and NFs in support of the Vertical use case experiments, acting as Day2 configurator. The following operations stand out among the configurations that can be handled by this component:

- Dynamic configuration of network or vertical functions, modifying the parameters of the required set of functionalities at runtime.
- Dynamic network or application state modification, in order to adapt to changing external conditions.
- Notification of changes or updates of functionalities in network or vertical functions.
- Enrichment of deployed experiment components by providing additional application level configuration to complete the Day2 configuration.

According to its scope of application, the Runtime Configurator is strongly related to the experiment definition and management, implemented in the 5G EVE Portal. In that sense, after deployment of the NS in the corresponding site facility/facilities for a Vertical experiment, all the required Day2 configurations defined by the Vertical in the Portal must be forwarded to the Runtime Configurator in order to translate them into specific commands and actions to be executed in the corresponding components.

These interactions are reflected in the diagram depicted in Figure 17 below, which shows how the Portal interacts with the Interworking Layer in order to deploy an experiment:

- First of all, the Experiment Life Cycle Manager component in the Portal provides the *service definition*, with all the Day0/1 configurations, to the Multi-site NSO in order to deploy the different NSs of the experiment within the specified site facilities. As a result, VNFs are instantiated in these site facilities and provided a preliminary configuration.

- After deploying the NSs that are going to be tested, Verticals can also complete the *experiment definition* through the Experiment Execution Manager, which will provide the Day2 configurations to the Runtime Configurator. This will consequently update the required deployed components.

It is worth noting, as it will be seen later on, that there are no standard specifications for the northbound or southbound reference points of this module. In fact, the range of components in the sites to which it might have to interact with, and the type of operations it might have to execute, is so broad that it is impossible to cover all the potential cases. As a consequence, the consortium will implement features of the Runtime Configurator as they are demanded by the participating Verticals, and will build on top of those to satisfy new demanded functionalities.

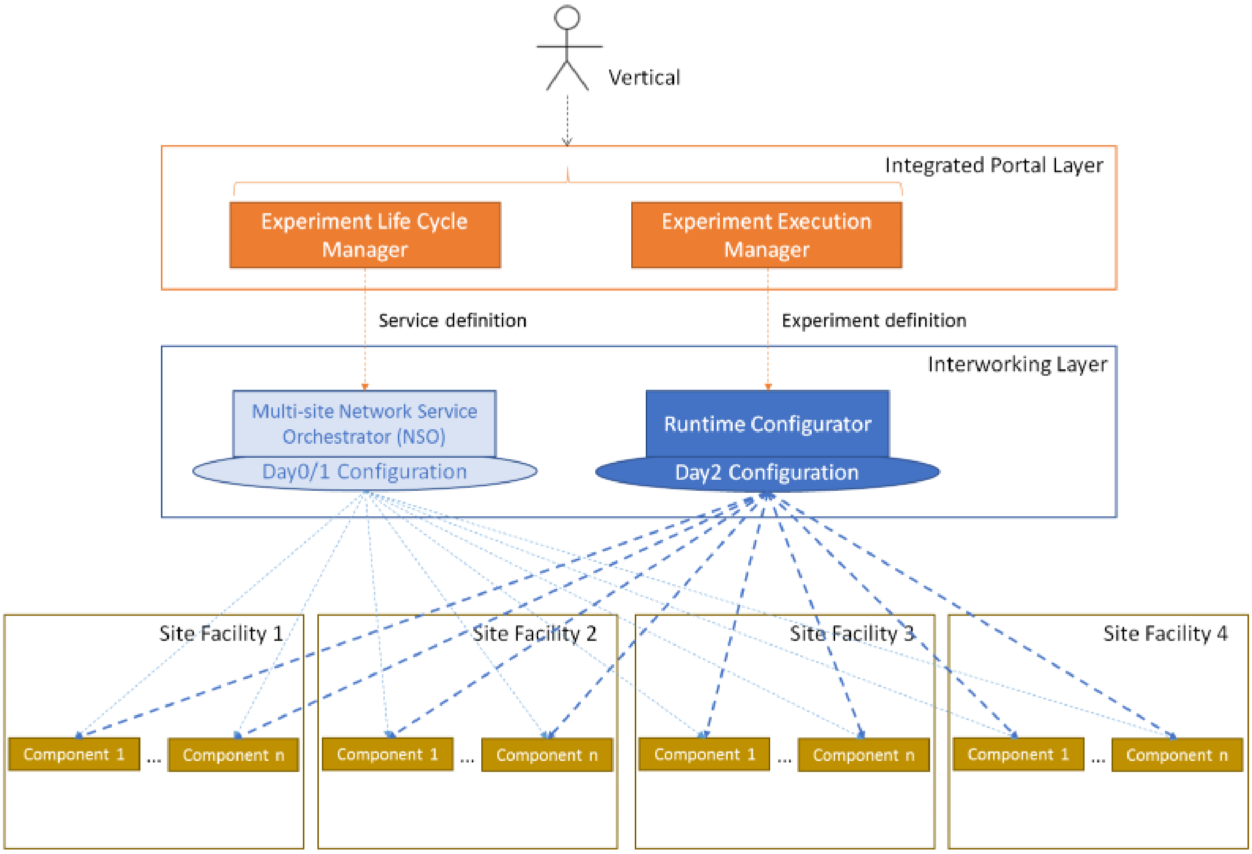

Figure 17: Scope of the Runtime Configurator in the I/W Framework

4.2.5 Data Collection Manager

The Data Collection Manager is responsible for the collection and persistence of all the network and vertical performance metrics that are required to be gathered during the execution of experiments, with two purposes:

- The monitoring of the experiment
- The validation of the targeted KPIs

It is a central component in the 5G EVE architecture, as it has to obtain the necessary metrics from site facilities and deliver them to the components which demand them: the Portal developed in WP4, and the KPI Analysis and Validation framework developed in WP5.

4.2.5.1 Information types

As a first step, it is important to clarify what kind of information could be handled by the Data Collection Manager, as there are significant differences between what is obtained from site facilities (logs) and what is delivered to other 5G EVE components (data):

- **Logs** will be obtained during the execution of the use cases, coming from different sources, for example: activity log files, configuration data, active/passive probes or monitoring devices. Generally speaking, these sources can be grouped in two categories: sources from 5G infrastructure and sources from Verticals applications. In both cases, it will be required to define the specific format in which they are stored in each site facility, as they need to be exposed to the Data Collector Manager in a homogeneous way.

- **Data** requires a transformation function taking logs as input, and providing the processed data ready to be used by upper layers as output. For example, the visualization component in the Portal will demand processed data, representing the targeted KPIs, instead of raw data.

The format of the KPI-related data provided by the Data Collection Manager to other components will be based on the format of the "targeted KPIs" defined within the scope of 5G EVE, which can be found in D1.1 – *Requirements definition and analysis from participant vertical industries* – [30].

4.2.5.2 Proposed architecture

The chosen architecture for the Data Collection Manager is based on a publish/subscribe messaging/message queuing paradigm implemented at each site facility. In that sense, each queue will be subscribed to the KPIs related data in each Vertical experiment. When data is published in the queue, it is directly delivered to the Data Collection Manager, which is able to store it or deliver it to upper layers through yet another queue, following again a publish/subscribe paradigm.

This kind of architecture is closer to a Big Data pipeline, rather than to a classic relational database, as streams and streams of data are pushed from site facilities without a specific format which fits in a relational model. For this purpose, a distributed system with scalability and parallel data processing capabilities is needed, which is reflected in the following diagram:

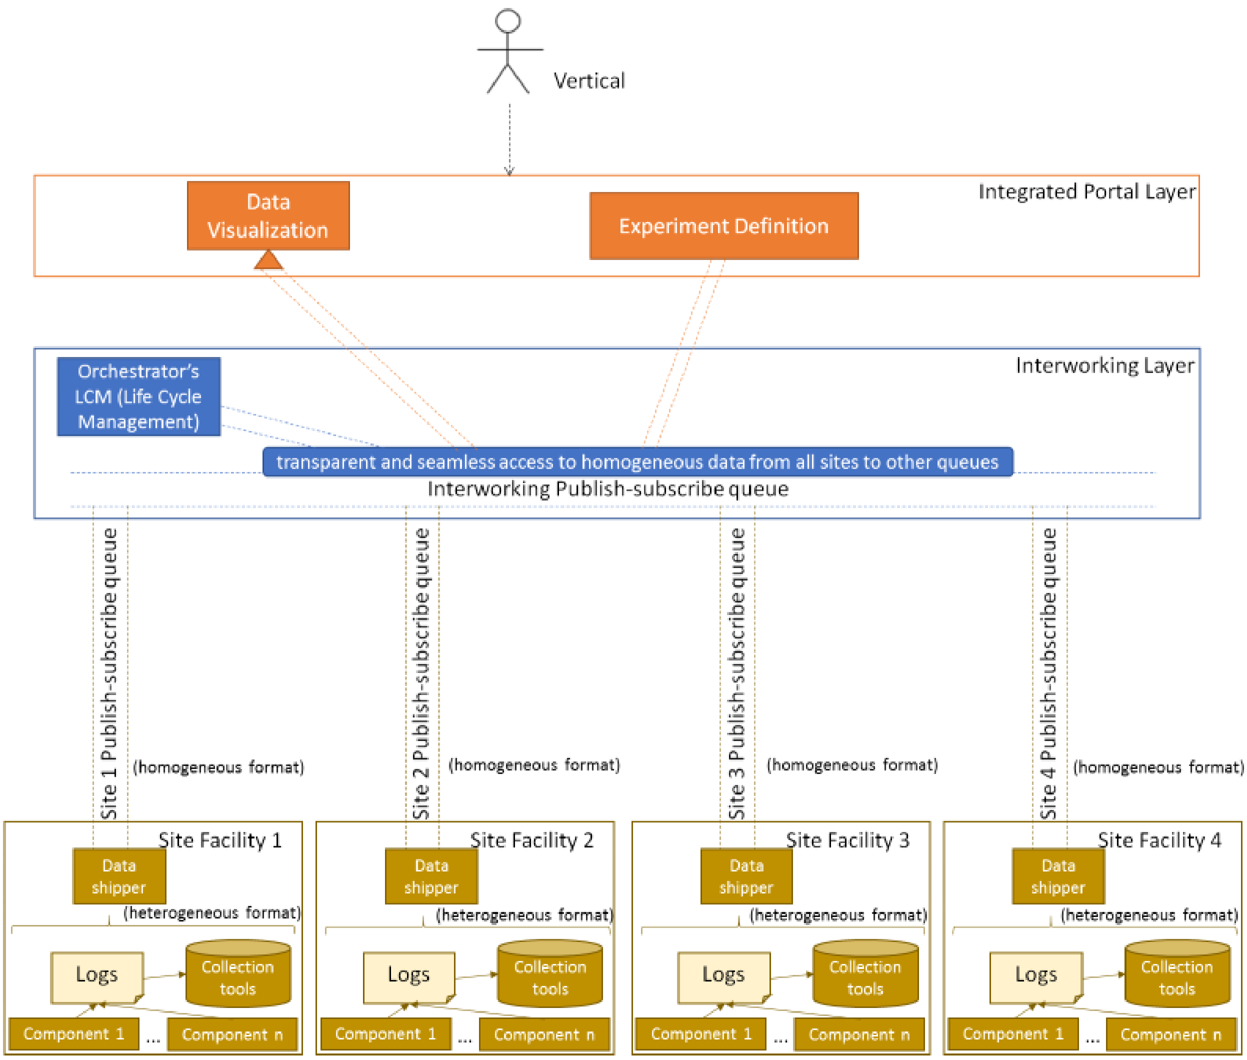

Figure 18: Runtime Configurator publish/subscribe mechanism

The basic workflow would be the following:

i. Each site will deploy an instance of the publish/subscribe queue in their facilities, connected to the components that generate or provide logs through the "data shippers". The main function of data shippers is to transform the heterogeneous, raw metrics obtained from components and collection tools, into data with a common, homogeneous format. As in all publish/subscribe schemes, there must be a subscription to a topic in order to obtain the data.

ii. When use cases are executed in site facilities, the different components used during the experiment (both physical and virtual) generate logs (raw information). As there is a connection with the publish/subscribe queue (again, by subscribing to a topic), logs can be delivered to their corresponding queue.

iii. When raw metrics are sent to the corresponding publish/subscribe queue, they are transformed by data shippers into homogeneous data and finally published in the queue. This provides a persistent storage (using disk for saving data) and results in a distributed system managed by sites: since each site manages its own queue, they can decide what data can be shared. Of course, it is not necessary to replicate all data to all queues. This can be achieved by making use of the different topics and defining access control policies, so the queues will only save information from the topics to which they are subscribed (and allowed to access to).

iv. All the data provided by each site facility is gathered in a common queue, called the Interworking publish/subscribe queue, which also provides persistence for all the data there published.

v. If other components (e.g. monitoring components in the Portal, experiment tools or the orchestrator itself) want to receive the data saved in the Interworking queue, they only have to subscribe to the queue to obtain the desired data. This is represented in Figure 18 with the block called "transparent and seamless access to homogeneous data from all sites to other queues".

It has to be noted that this proposed architecture derives in a Data Collection Manager module in the I/W Framework working only with processed data, not with logs (see section 4.2.5.1). This will be the initial approach in 5G EVE, but in case it was required for any reason, the module could be upgraded to work with logs as well and perform itself the functionality of the data shippers.

Other potential capabilities that could be included in the module, as depicted in Figure 19 below, would be based on additional internal components, such as:

- Another queue dedicated to fast data processing (using in-memory data storage, for example), in case real time processing was needed. The publish/subscribe queue would then be used for normal processing, without time constraints.
- A "cold storage" to save data from the publish/subscribe queue and maintain it for a long time, for example as backup. A potential implementation for this could be a NoSQL database, as there is not a clear relational model for this case.

It is worth noting that access to the queues does not change with these new components. The Interworking queue would only have to notify which was the desired queue (fast or normal) in order to obtain the data.

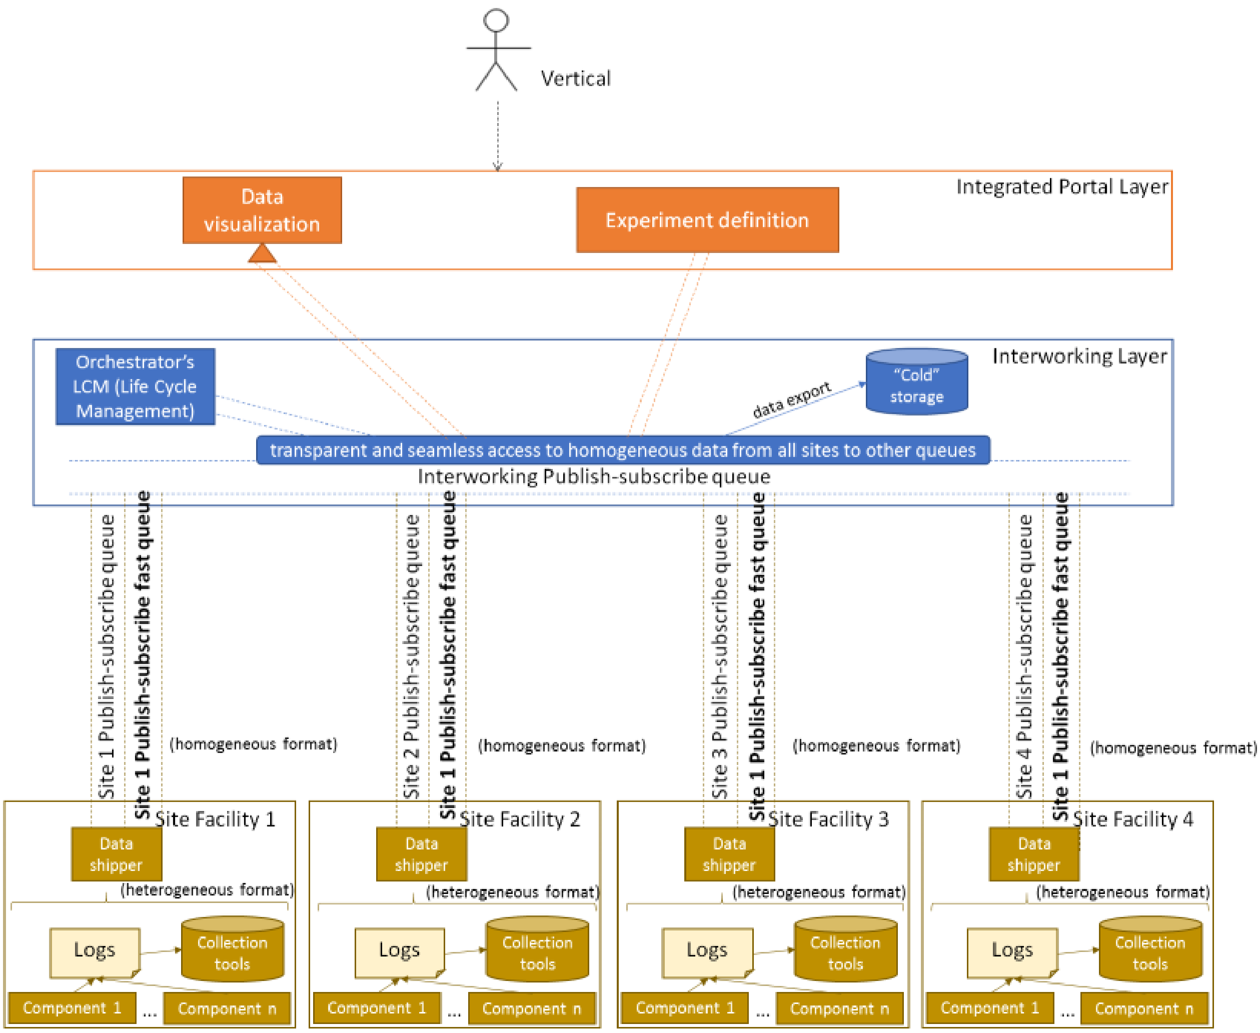

Figure 19: Runtime Configurator architecture

4.2.6 Summary of architectural design agreements

Once the main components of the I/W Framework have been described, it is important to list a number of agreements that are applicable to its architectural design. These have been elaborated during the technical discussions in Task 3.2[20] – *Interworking reference model definition, including multi-site slicing & orchestration* –, and have been agreed when needed with WP2 – *Implementation, pilot execution and validation* –.

These are as follows:

- The I/W Framework will be based on standard components and interfaces when those exist (as described and justified in section 4.1).
- In general, the Multi-site NSO will act as a proxy for the interaction of the Multi-site Catalogue and Inventory with the sites (as described and justified in section 4.2).
- The Multi-site NSO will interact with "orchestrators", not with "sites". This implies that, if any site requires to instantiate a NS in local infrastructures managed by different orchestrators, the Multi-site NSO will be responsible of managing the required local intra-site connectivity. The main advantage of

[20] Task responsible of producing this Deliverable

this approach is that we avoid the need for a per-site aggregation layer to wrap the orchestration granularity within the site.

- At least for the first version of the I/W Framework, it will not interact with local site components in control of the RAN and MEC, nor with local SDN controllers. These will stay under the control of the local NFVOs. The only considered exception to this is the potential interaction of the Multi-site NSO with a local SDN WAN controller, if such scenario is implemented at any of the sites.
- At least for the first version of the I/W Framework, the NS and VNF on-boarding from the Portal to the sites catalogues (via Multi-site Catalogue and Multi-site NSO) will not be fully automated; it will require at least the approval of the site managers.

Any of these agreements may be revisited in the future in case interworking requirements change. Such changes, if any, will be documented in subsequent WP3 Deliverables.

4.3 Interworking Framework workflows

The workflows of the 5G EVE Platform will assist on the identification of the interactions among components of the I/W Framework and components in the sites. From the I/W Framework perspective, the identification of these interactions is essential to understand the internal architecture and logics of the components, and their functional behaviour.

The purpose of 5G EVE is to provide a multi-site experimental facility for the deployment of 5G services as requested by Vertical customers. As such, the system has to provide capabilities for instantiating those services, operating them along time obtaining meaningful data through monitoring and testing capabilities, and finally decommissioning them. A Network Service (NS) Instance in 5G EVE, as in single NFV administrative domains, will be assumed to have two elementary state values in the NFVO:

- INSTANTIATED
- NOT_INSTANTIATED

Of course, a centralized orchestrator will take care of the NS lifecycle, becoming a key element in charge of managing and controlling all the services. Such is the role of the Multi-Site Network Service Orchestrator, which becomes a central piece on the 5G EVE architecture.

One of the design criterion adopted in the project has been to follow standardized specifications whenever they were available. This is especially relevant in the case of the I/W Framework since different sites, each of them representing independent administrative domains, need to interact in a well-defined manner. Adopting standardized-oriented solutions greatly facilitates such approach.

To this respect, it is required to focus on the purpose of the I/W Framework to elicit the main specifications of interest. On the one hand, it is the 5G EVE system for the orchestration of services as result of Vertical customer's requests. On the other, it must deal with different sites, each of them representing a different administrative domain. By scanning the progress on ETSI NFV standardization, the specifications appearing as relevant to be followed in the I/W Framework workflow specifications are:

- ETSI NFV-IFA 028 [32], "Report on architecture options to support multiple administrative domains", which analyses the requirements and needs to be supported by an NFV architecture enabling orchestration of services and resources across multiple independent administrative domains.
- ETSI NFV-IFA 013 [9], "Os-Ma-Nfvo reference point - Interface and Information Model Specification", that specifies the interaction of the OSS/BSS and the orchestration components, for handling the management, requests and lifecycle of the service requests.
- ETSI NFV-IFA 030 [12], "Multiple Administrative Domain, Aspect Interfaces Specification", which elaborates on the interface needs among orchestrators laying in distinct domains.
- Draft ETSI NFV-SOL 011, "RESTful protocols specification for the Or-Or Reference Point", which details the implementation of the inter-orchestrator interface. Note that this specification is in draft status at the time of writing this deliverable, then subject to evolve.

These specifications consider a number of orchestration actions, but are not always coincident since the number and purpose of actions supported can be different. The following is a comparison between NFV-IFA 013 and NFV-IFA 030.

Table 6: NFV-IFA specifications comparison

NFV-IFA 013	NFV-IFA 030
NSD Management	NSD Management
NS Lifecycle Management	NS Lifecycle Management
NS Performance Management	NS Performance Management
NS Fault Management	NS Fault Management
Policy Management	Policy Management
VNF Package Management	-
NFVI Capacity Information	-
VNF Snapshot Package Management	-
-	NS Lifecycle Operation Granting
-	NS Instance Usage Notification

At the sight of these differences, which may be important depending on the use case, it is necessary to clearly identify which specification to follow, how the specifications complement each other, and how these specifications map onto the Multi-Site Network Service Orchestrator approach. Based on those discussions, this section analyses the workflows from the orchestration perspective.

For that purpose, we have focused on the NS lifecycle management process. Other procedures, such as performance or fault management can have dependencies on the developments from other WPs. Additionally, it is assumed that procedures internal to each of the sites follow the same specifications and operations as the ones refereed above.

Finally, although not included in the workflows for simplicity, it must be noted that all the instantiation-related operations include the notification of the Multi-site Orchestrator to the Multi-site Inventory, so that the latter has a clear view of the active resources and services.

4.3.1 NS Lifecycle Management interface

The management of a NS involves the following actions, extracted in this case from NFV-IFA 013, many of which are depicted in Figure 40:

- Create NS Identifier (clause 7.3.2 of IFA 013)
- Delete NS Identifier (clause 7.3.8 of IFA 013)
- Instantiate NS (clause 7.3.3 of IFA 013)
- Terminate NS (clause 7.3.7 of IFA 013)
- Scale NS (clause 7.3.4 of IFA 013)
- Heal NS (clause 7.3.9 of IFA 013)
- Query NS (clause 7.3.6 of IFA 013)
- Subscription/Notification (clause 7.3.11 to clause 7.3.14 of IFA 013)

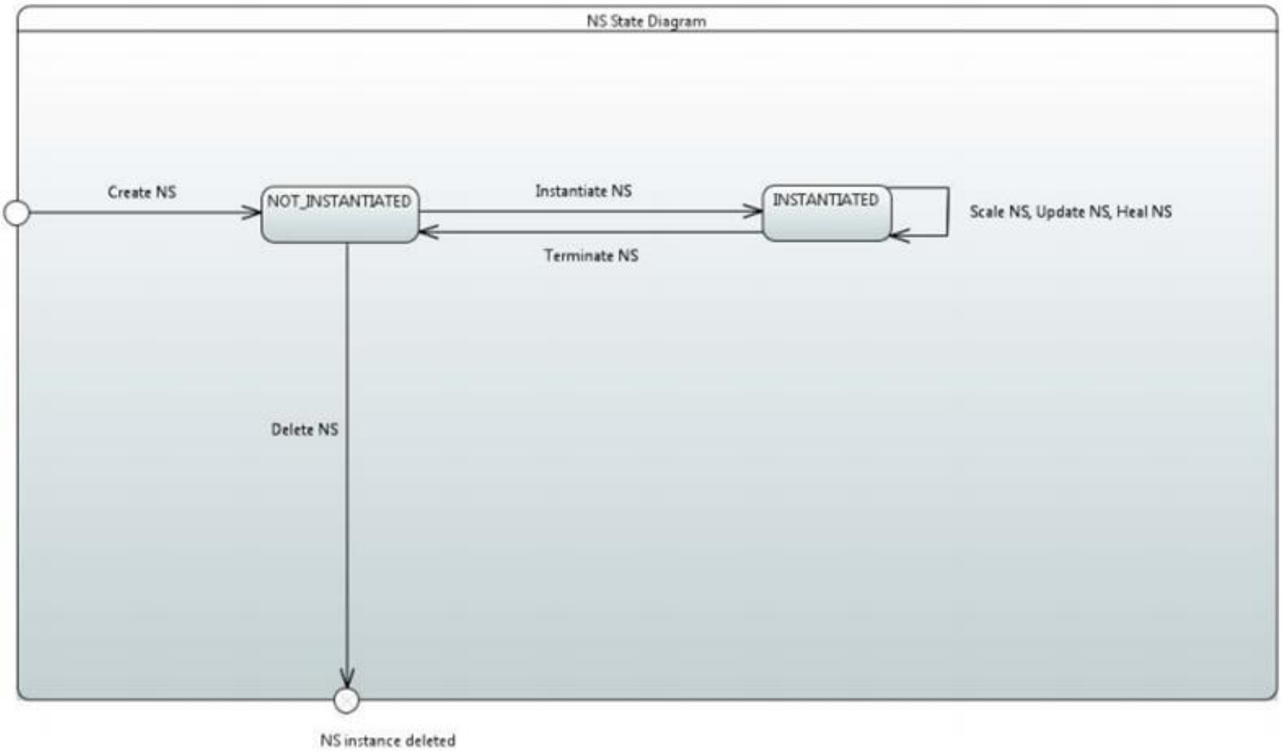

Figure 20: NS State diagram

4.3.1.1 Create NS Identifier

The operation (Figure 21) creates an Identifier for a NS Instance, and an associated instance of an NsInfo information element, in the NOT_INSTANTIATED state. At this stage no additional lifecycle operation is performed (e.g., there is no instantiation of the NS).

After the creation, there is an immediate return of an NS Instance Identifier to be used in any following lifecycle operations, such as the instantiation of the NS itself.

When the operation is successful, the NS Instance Identifier and the associated instance of an NsInfo information element are created. In case of failure, an error notification is returned.

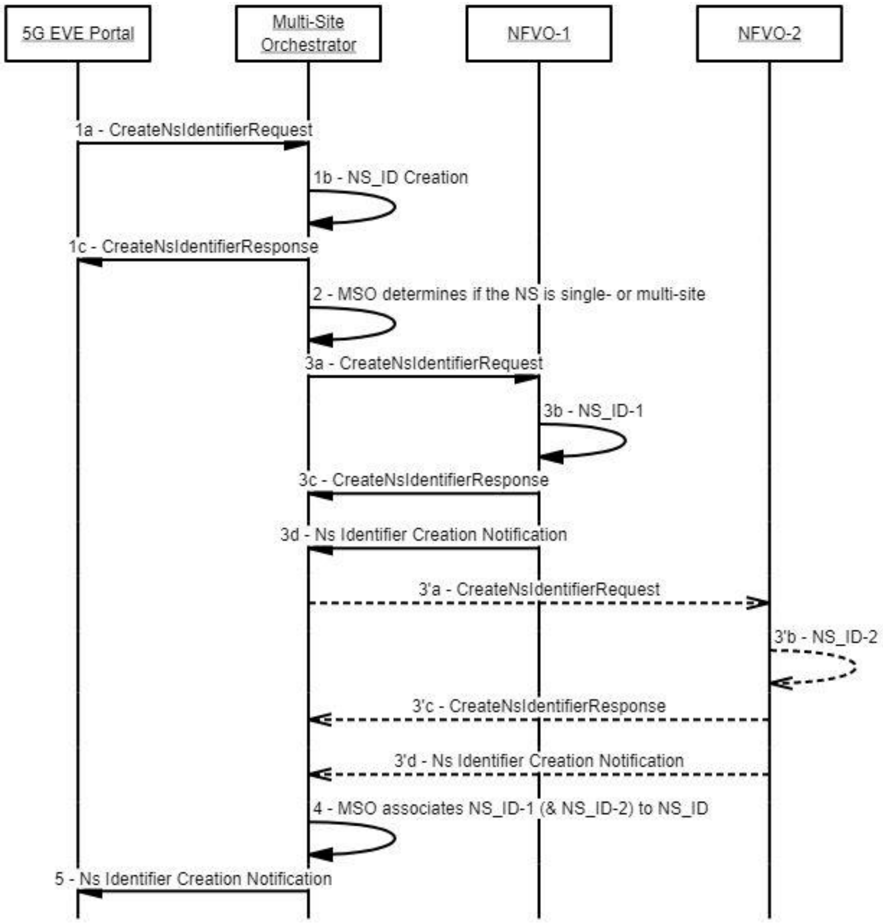

Figure 21: Workflow for the creation of a NS Identifier

The Vertical customer requests the deployment of a service through the 5G EVE Portal. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified:

Step 1. The 5G EVE Portal instructs the Multi-site Orchestrator to create the NS identifier for that service (1a). The coding of the message, and the parameters within, are those defined in the NBI. The Multi-site Orchestrator creates a new NS instance object, and the associated NS instance identifier at Orchestrator level (1b). Subsequently, the Orchestrator returns a response to the 5G EVE Portal with reference to the new NS and the corresponding identifier (1c).

Step 2. The Orchestrator will determine if the Network Service subject of the creation action expands to more than one local orchestrator[21].

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the creation of each nested NS identifier, by sending an individual request to each of the involved local orchestrators (e.g., 3a and 3'a). In case the NS is not multi-site, the request is directed only to the proper orchestrator. The different orchestrators involved will create the corresponding NS instance and its corresponding identifier (3b and 3'b), followed by a response to the Multi-site Orchestrator with reference to the new NS and the corresponding identifier (3c and 3'c). Finally, each local orchestrator will send an Ns Identifier Creation Notification to the

[21] As described in section 4.5, one site may have more than one local orchestrator; therefore, the Multi-site Orchestrator addresses "local orchestrators" not "sites".

Multi-site Orchestrator indicating the creation of the per-orchestrator NS instance resource and the associated NS instance identifier (3d and 3’d).

Step 4. Once the different creation actions are performed with success, the Orchestrator will create the association of the per-site NS identifiers (NS_ID-1 and NS_ID-2 in the exemplary workflow) with the NS Identifier of the Vertical customer (NS_ID).

Step 5. Finally, the Orchestrator sends an NS Identifier Creation Notification to the 5G EVE Portal indicating the creation of the NS instance resource and the associated NS instance identifier.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the creation request.

4.3.1.2 Delete NS Identifier

This operation (Figure 22) deletes a NS instance identifier (which is in the NOT_INSTANTIATED state) and the associated NsInfo information element.

In case of success, the NS instance identifier and the associated instance of the NsInfo information element become deleted and can no longer be used.

In case the NS instance was not in the NOT_INSTANTIATED state (i.e. terminated or not instantiated), the operation is rejected.

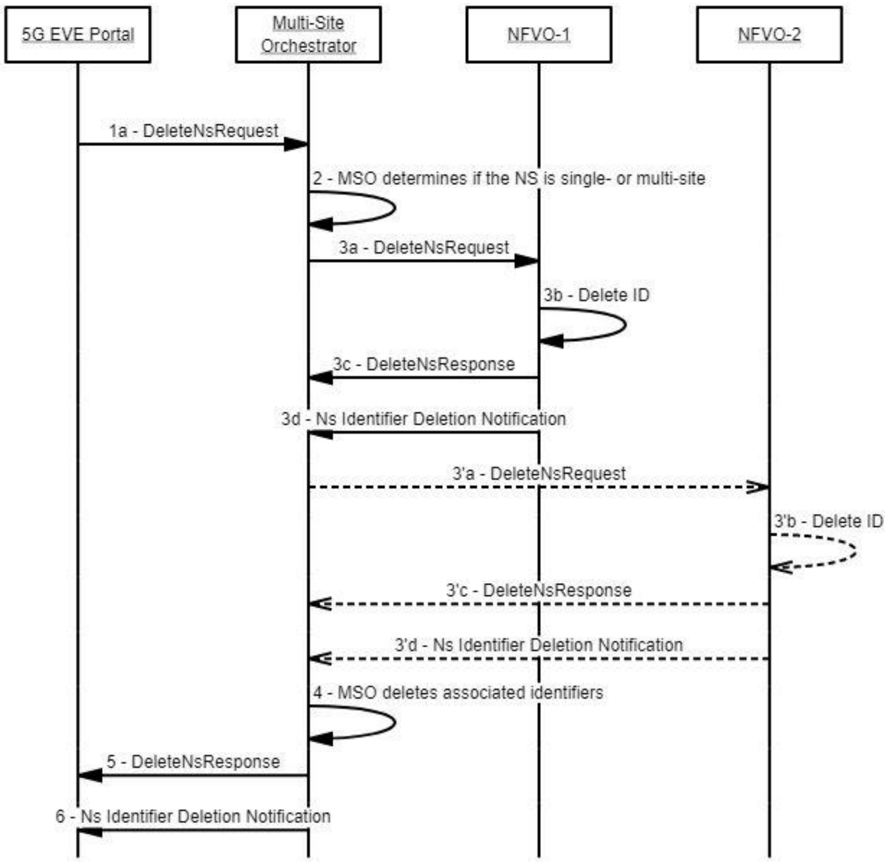

Figure 22: Workflow for the deletion of a NS Identifier

The Vertical customer requests the deletion of a service through the 5G EVE Portal. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified.

Step 1. The 5G EVE Portal instructs the Multi-site Orchestrator to delete the corresponding NS identifier of the service. The coding of the message, and the parameters within, are those defined in the NBI.

Step 2. The Orchestrator will determine if the Network Service subject of the delete action expands to more than one site.

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the deletion of the corresponding NS identifier, to the multiple local orchestrators involved in the service, by sending an individual request to each of them for the particular (nested) NS (e.g., 3a and 3'a). In case the NS is not multi-site, the request is directed only to the proper orchestrator. The different orchestrators involved will delete the corresponding NS instance and its corresponding identifier (3b and 3'b), followed by a response to the Multi-site Orchestrator (3c and 3'c). Finally, each local orchestrator will send a NS Identifier Deletion Notification to the Orchestrator indicating the deletion of its NS instance resource and the associated NS instance identifier (3d and 3'd).

Step 4. Once the different deletion actions are performed with success, the Orchestrator will delete the association of the identifiers of the per-site NSs with the NS Identifier of the Vertical customer. After that, the Orchestrator will also delete the NS Identifier previously allocated for the Vertical customer.

Step 5. The Multi-site Orchestrator returns a delete response to the 5G EVE Portal.

Step 6. The Multi-site Orchestrator sends an NS Identifier Deletion Notification to the 5G EVE Portal to indicate the deletion of the NS instance resource and the associated NS instance identifier.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the deletion request.

4.3.1.3 Instantiate NS

This operation (Figure 23) instantiates a NS in the 5G EVE facility. It can only be invoked for a NS Instance which is in the NOT_INSTANTIATED state. If the operation is successful, the NS will be instantiated. Otherwise, an error notification will be provided as "result".

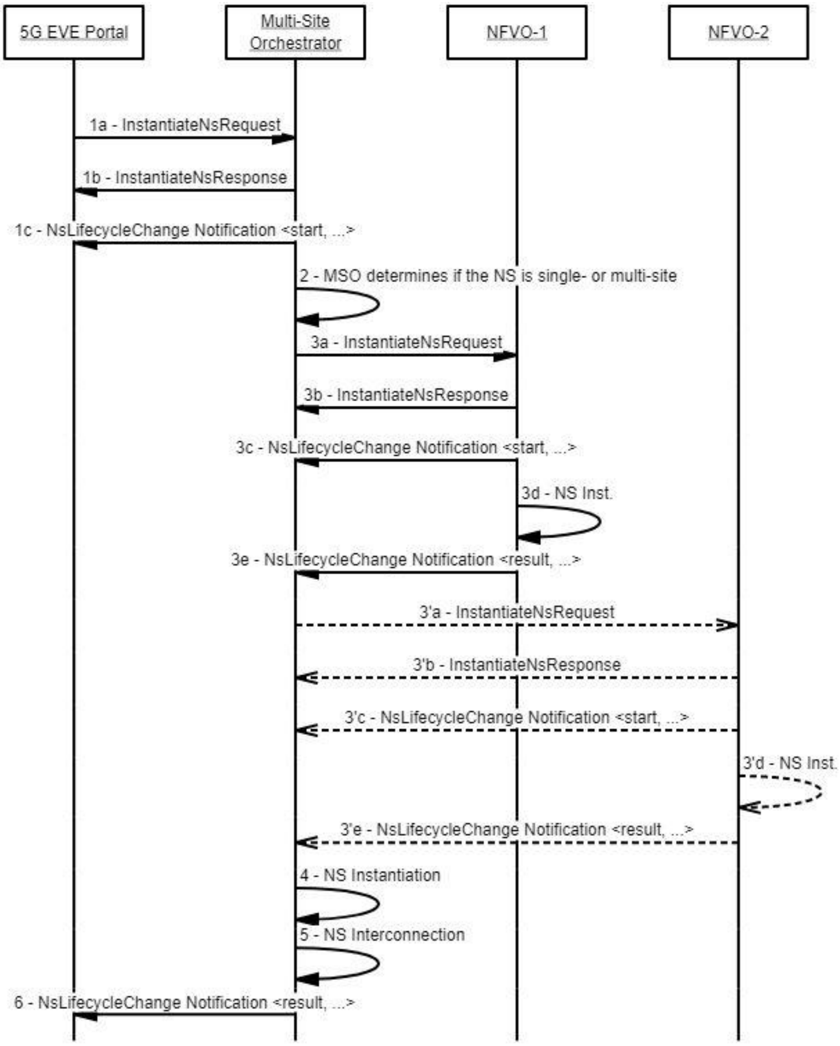

Figure 23: Workflow for the instantiation of a NS

The Vertical customer requests the instantiation of a service through the 5G EVE Portal by using a service identifier previously created. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified.

Step 1. The 5G EVE Portal instructs the Multi-site Orchestrator to instantiate a service through the corresponding NS identifier (1a). The coding of the message and the parameters within are those defined in the NBI. The Orchestrator returns a lifecycleOperationOccurrenceId in a response message that identifies the LCM operation (1b). Finally, the Orchestrator sends a NS LCM Operation Occurrence Notification indicating the "start" of the lifecycle operation (1c).

Step 2. The Orchestrator will determine if the Network Service subject of the instantiation action expands to more than one site.

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the instantiation of the corresponding NSs, to the multiple local orchestrators involved in the service, by sending an individual request to each local orchestrator for the particular (nested) NS (e.g., 3a and 3'a). In case the NS is not multi-site, the request is directed only to the proper orchestrator of the concerned site. The different orchestrators involved send a response to the Multi-site Orchestrator (3b and 3'b). Subsequently, each orchestrator sends a NS LCM Operation Occurrence Notification indicating the "start" of its corresponding lifecycle operation (3c and 3'c). The particular orchestrator in each site proceeds to instantiate the service under its responsibility (3d and 3'd) followed of a NS LCM Operation Occurrence Notification indicating the "result" of the lifecycle operation (3e and 3'e).

Step 4. Once the different local instantiation actions are performed with success, the Orchestrator completes the multi-site instantiation (if applicable) by associating the per-site NSs instances.

Step 5. The Orchestrator interconnects the NSs under control of distinct local orchestrators (in case of multi-site).

Step 6. The Orchestrator sends a NS LCM Operation Occurrence Notification indicating the "result" of the lifecycle operation to the 5G EVE Portal to indicate the complete instantiation of the NS.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the instantiation request.

4.3.1.4 Scale NS

This operation (Figure 24) will scale a Network Service instance. Scaling an NS instance can be performed by: i) explicitly adding/removing existing VNF instances to/from the NS instance, ii) by leveraging on the abstraction mechanism provided by the NS scaling aspects and NS levels information elements declared in the NSD, or iii) by scaling individual VNF instances that are part of the NS itself.

In case of success, the NS instance becomes scaled according to the request. In case of failure, appropriate error information is provided. On the successful as well as the unsuccessful completion of the operation, any local NFVO involved in the scaling procedure shall send its "result" Notification

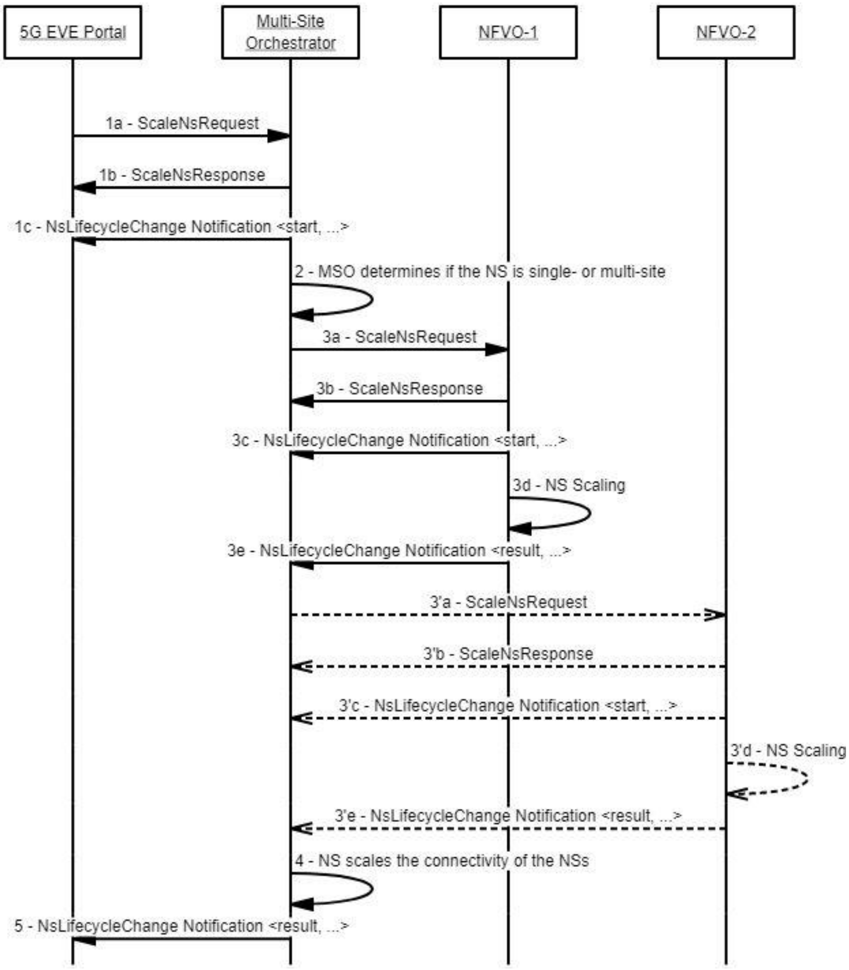

Figure 24: Workflow for the scaling of a NS

The Vertical customer requests the scaling of a service through the 5G EVE Portal by using a service identifier previously created. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified.

Step 1. The 5G EVE Portal instructs the Orchestrator to scale a service through the corresponding NS identifier (1a). The coding of the message and the parameters within are those defined in the NBI. The Orchestrator returns a lifecycleOperationOccurrenceId in a response message that identifies the LCM operation (1b). Finally, the Orchestrator sends a NS LCM Operation Occurrence Notification indicating the "start" of the lifecycle operation (1c).

Step 2. Based on the service to be scaled, the Orchestrator will determine if the NS consists of nested NSs, and if so, what and where of those NS are required to be scaled across the different sites.

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the scaling of the corresponding NSs, to the multiple local orchestrators involved in the service, by sending an individual request to each of the local orchestrators for the particular (nested) NS (e.g., 3a and 3'a). In case the NS is not multi-site, the request is directed only to the proper orchestrator. The different orchestrators involved send a response to the Multi-site Orchestrator (3b and 3'b). Subsequently, each site send a NS LCM Operation Occurrence Notification indicating the "start" of its corresponding lifecycle operation (3c and 3'c). The local orchestrator in each site proceeds to scale the service under its responsibility (3d and 3'd) followed of a NS LCM Operation Occurrence Notification indicating the "result" of the lifecycle operation (3e and 3'e).

Step 4. Once the different local scaling actions are performed with success, the Orchestrator assess if the interconnection of the NSs should be also scaled (in case of multi-site).

Step 5. The Orchestrator sends a NS LCM Operation Occurrence Notification indicating the "result" of the lifecycle operation to the 5G EVE Portal to indicate the complete scaling of the NS.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the scaling request.

4.3.1.5 Update NS

This operation (Figure 25) updates an NS Instance. In case of success, the NS results updated according to the request. In case of failure, appropriate error information is provided in the "result" Notification. On the successful as well as the unsuccessful completion of the operation, any particular NFVO involved in the update procedure in each site shall send its "result" Notification.

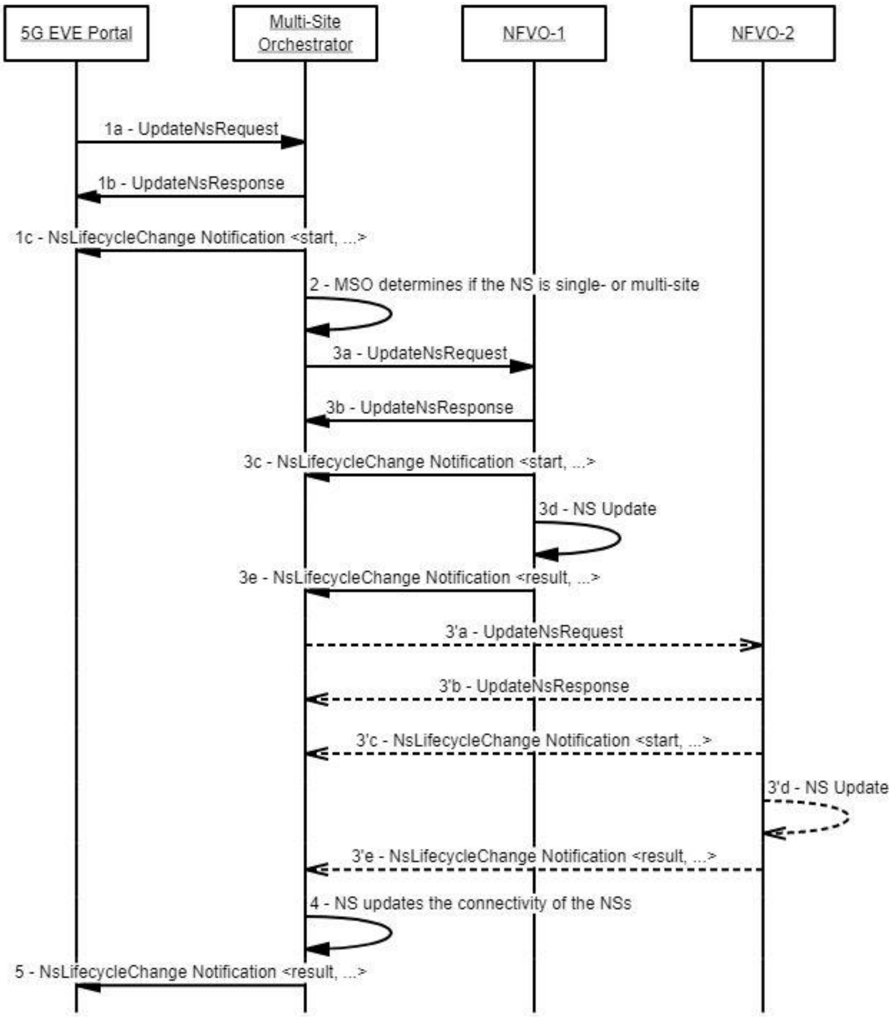

Figure 25: Workflow for the update of a NS

The Vertical customer requests the update of a service through the 5G EVE Portal by using a service identifier previously created. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified.

Step 1. The 5G EVE Portal instructs the Multi-site Orchestrator to update a service through the corresponding NS identifier (1a). The coding of the message and the parameters within are those defined in the NBI. The Orchestrator returns a lifecycleOperationOccurrenceId in a response message that identifies the LCM operation (1b). Finally, the Orchestrator sends an NS LCM Operation Occurrence Notification indicating the “start” of the lifecycle operation (1c).

Step 2. Based on the service to be scaled, the Orchestrator will determine if the NS consists of nested NSs, and if so, what and where of those NS are required to be updated across the different sites.

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the update of the corresponding NSs, to the multiple local orchestrators involved in the service, by sending an individual request to each of the local orchestrators for the particular (nested) NS (e.g., 3a and 3’a). In case the NS is not multi-site, the request is directed only to the proper orchestrator. The different orchestrators involved send a response to the Multi-site Orchestrator (3b and 3’b). Subsequently, each orchestrator sends a NS LCM Operation Occurrence Notification indicating the “start” of its corresponding lifecycle operation (3c and 3’c). The particular orchestrator in each site proceeds to update the service under its responsibility (3d and 3’d) followed of a NS LCM Operation Occurrence Notification indicating the “result” of the lifecycle operation (3e and 3’e).

Step 4. Once the different per-site update actions are performed with success, the Orchestrator assesses if the interconnection of the NSs should be also updated (in case of multi-site).

Step 5. The Orchestrator sends a NS LCM Operation Occurrence Notification indicating the "result" of the lifecycle operation to the 5G EVE Portal to indicate the complete update of the NS.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the update request.

4.3.1.6 Terminate NS

This operation will terminate an active NS. This operation can only be used with an NS instance in the INSTANTIATED state, notifying proper error message otherwise. The termination of the NS instance does not delete the NS instance identifier, nor the associated instance of the NsInfo information element, which is performed through the workflows defined in section 4.3.1.2.

In case of success, the NS becomes terminated (i.e. moved into NOT_INSTANTIATED state), and the corresponding resources used by the NS are released. In case of failure, appropriate error information is provided in the "result" NS LCM Operation Occurrence Notification.

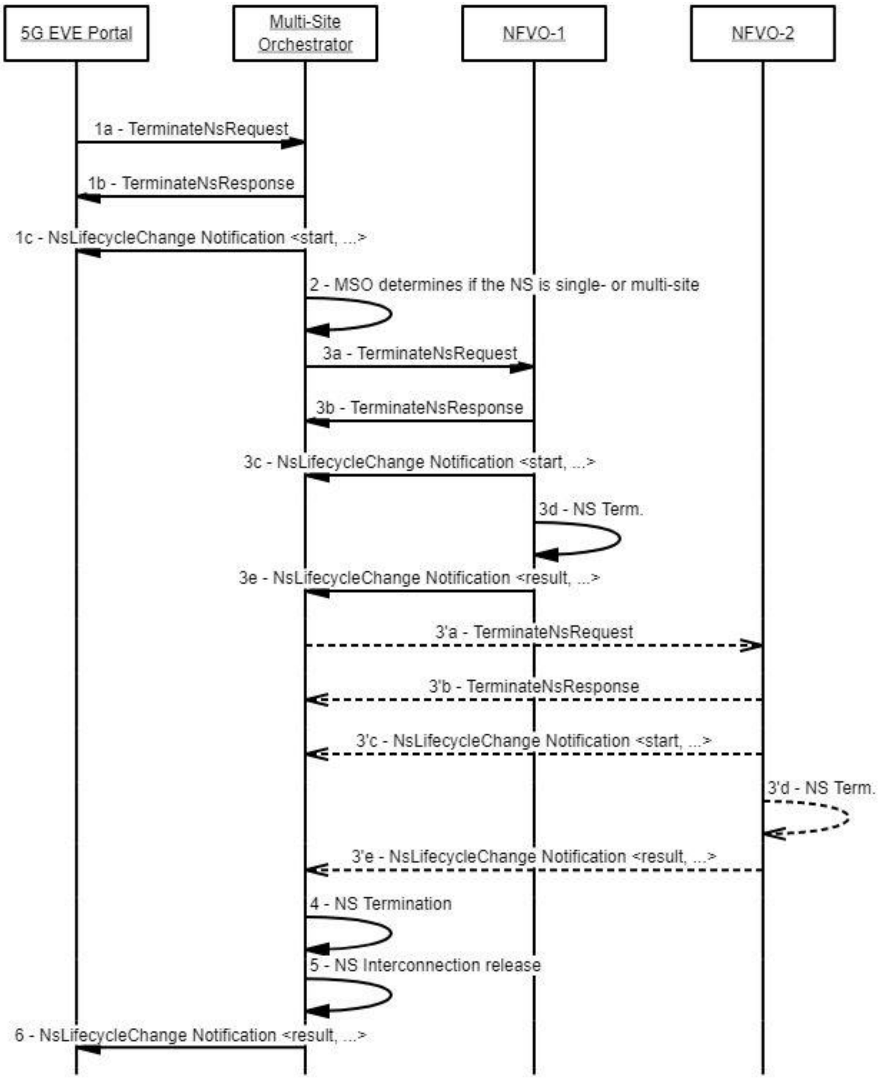

Figure 26: Workflow for the termination of a NS

The Vertical customer requests the termination of a service through the 5G EVE Portal by using a service identifier previously created. The specific method and involved parameters for such request will be defined in WP4. The following workflow is identified.

Step 1. The 5G EVE Portal instructs the Multi-site Orchestrator to terminate a service through the corresponding NS identifier (1a). The coding of the message and the parameters within are those defined in the NBI. The Orchestrator returns a lifecycleOperationOccurrenceId in a response message that identifies the LCM operation (1b). Finally, the Orchestrator sends a NS LCM Operation Occurrence Notification indicating the “start” of the lifecycle operation (1c).

Step 2. The Orchestrator will determine if the Network Service subject of the termination action expands to more than one site.

Step 3. In the case of being a multi-site service, the Orchestrator triggers the same action, i.e., the termination of the corresponding NSs, to the multiple local orchestrators involved in the service, by sending an individual request to each local orchestrator for the particular (nested) NS (e.g., 3a and 3’a). In case the NS is not multi-site, the request is directed only to the proper orchestrator. The different orchestrators involved send a response to the Multi-site Orchestrator (3b and 3’b). Subsequently, each site sends an NS LCM Operation Occurrence Notification indicating the “start” of its corresponding lifecycle operation (3c and 3’c). The particular orchestrator in each site proceeds to terminate the service under its responsibility (3d and 3’d) followed of a NS LCM Operation Occurrence Notification indicating the “result” of the lifecycle operation (3e and 3’e).

Step 4. Once the different per-site termination actions are performed with success, the Orchestrator completes the multi-site termination (if applicable) by terminating the NS service as offered to the Vertical customer.

Step 5. The Orchestrator releases the interconnection previously connecting the NSs under control of distinct local orchestrators (in case of multi-site).

Step 6. The Orchestrator sends a NS LCM Operation Occurrence Notification indicating the “result” of the lifecycle operation to the 5G EVE Portal to indicate the complete termination of the NS.

The Vertical customer will receive a notification from the 5G EVE Portal assessing the success or not of the termination request.

4.3.1.7 Other NS operations

There are other potential operations which are not included in this document since they are less relevant for the initial purpose of 5G EVE. These actions are:

- NS healing
- NS query
- Get Operation Status

In any case, the reference procedures and workflows are the ones detailed in NFV-IFA 013, NFV-IFA 028, and NFV-SOL 005, adapted to the multi-site scenario as in the other cases here described. If any of these is required in future versions of the I/W Framework, it will be documented in future WP3 Deliverables.

4.3.2 High-Level workflows

This section describes the high-level workflows involved in the service activation, service lifecycle and service decommissioning phases.

During service activation (Figure 27), the 5G EVE facility will basically assist on the creation of the Vertical customer service by first creating the NS instance (and its identifier) and by instantiating the NS. During this phase the service becomes created and stays in INSTANTIATED state.

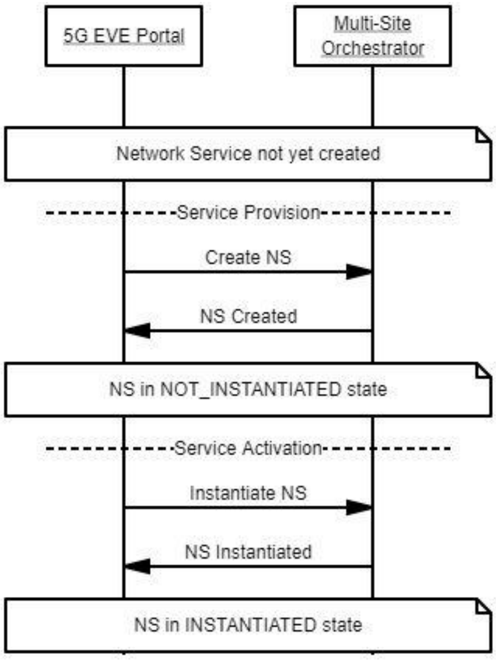

Figure 27: Service provisioning and activation workflows

Next phase is related to the service operation (Figure 28). Different lifecycle actions can take place during this phase, such as scaling, update, and healing (this one not represented in the figure below). During this phase, the service always is in INSTANTIATED state, and evolves according to the lifecycle actions triggered by the Vertical customer.

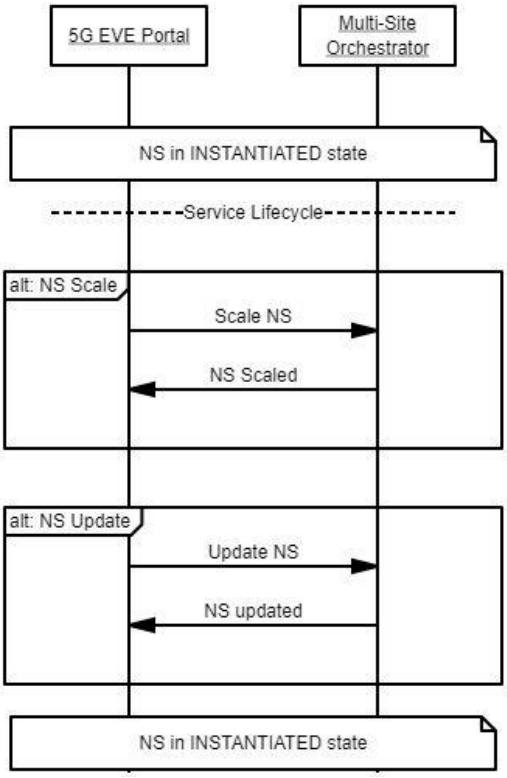

Figure 28: Service lifecycle workflows

Finally, the service decommissioning phase (Figure 29) implies first the deactivation of the service, and later the total decommissioning or release of the resources allocated for it. The final state for this service is the NOT_INSTANTIATED.

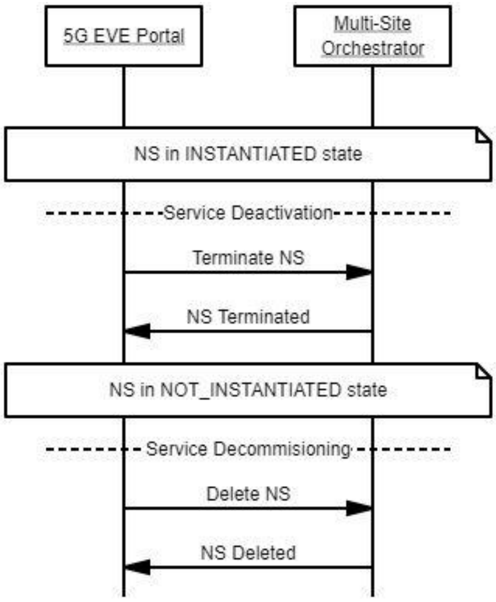

Figure 29: Service decommissioning workflows

Additional workflows can be foreseen for the overall NSs, namely:

- Service capabilities discovery: during the first stages of 5G EVE, the service discovery stage will be manually triggered, for example by declaring in the 5G EVE Portal the per-site capabilities available. Based on that, the 5G EVE service can be made available to Vertical customers.
- Service assurance: this stage is intrinsically related with the 5G EVE capabilities for testing and benchmarking the Vertical customer services. The specific entities related to the monitoring and testing capabilities are linked to the activities carried out in WP5.

These workflows are not characterized in this document. If required, this will be described in subsequent documents in WP3.

4.3.3 Specific-Purpose workflows

Having defined in the previous section the main workflows dealing with services, this section includes additional workflows for other supporting procedures which are also required as part of the 5G EVE Platform functionalities.

4.3.3.1 Catalogues synchronization

The first workflow is related to catalogues synchronization (Figure 30), in particular the Multi-site Catalogue synchronizing with the local catalogues. It has to be noted that the opposite procedure (that local catalogues synchronize with the Multi-site Catalogue) has decided to be done manually, or at least with some intervention from site managers, as described in section 4.2.6.

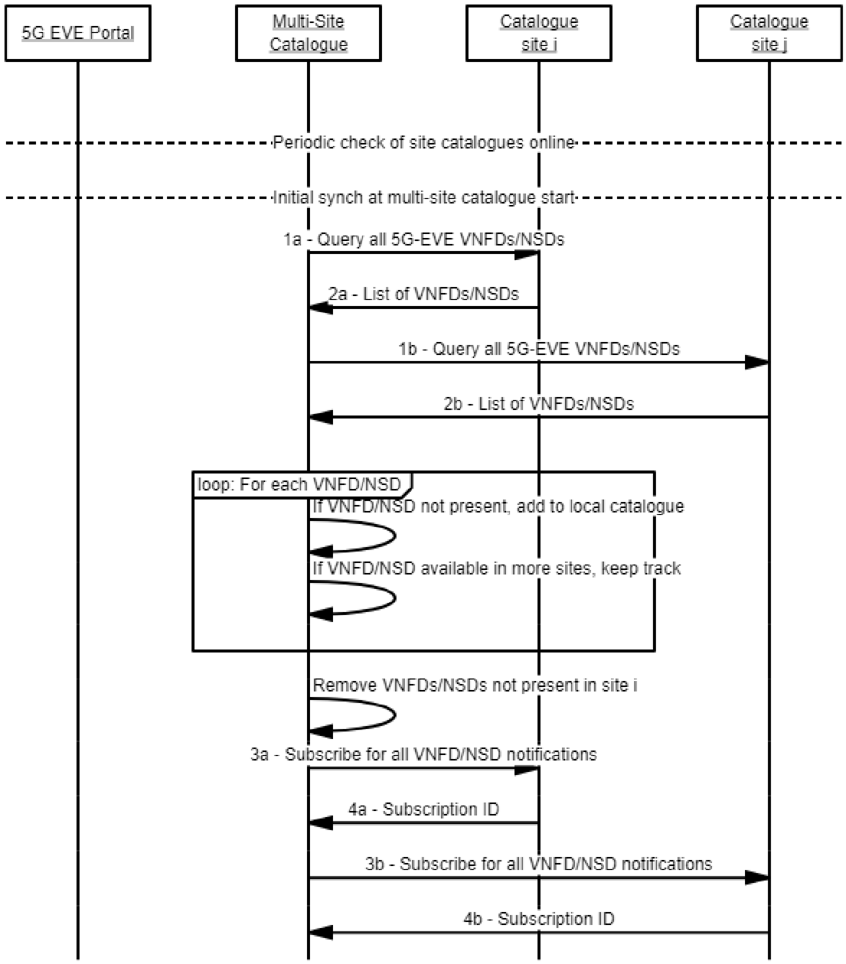

Figure 30: Catalogues synchronization workflow

Initial requirement for this workflow is that the local catalogues are known and permit external communication. Based on this, when the Multi-site Catalogue is started, the following workflow is defined.

Step 1. The Multi-site Catalogue queries all the known local catalogues for their available NS and VNF Descriptors (1a/1b). Although not depicted in the figure for simplicity, it has to be noted that all these operations among catalogues are proxied by the Multi-site NSO, using its SBI as described in 4.5.2. The response is the list of NS and VNF Descriptors from each local orchestrator catalogue (2a/2b).

Step 2. For each NSD and VNFD, the Multi-site Catalogue checks its status: if it was not present, it adds the descriptor; if it was present, and still is, it performs update operations if required; if it was present, and now it is not, it removes the descriptor.

Step 3. Finally, the Multi-site Catalogue subscribes to the notification service in the local catalogues with regards to the new NS and VNF Descriptors (3a/3b). The response is a subscription ID (4a/4b). Although not depicted, the Multi-site Catalogue also unsubscribes from the notifications regarding the Descriptors which have been removed.

4.3.3.2 VNF on-boarding

The following (Figure 31) is the workflow for VNF on-boarding in the Multi-site Catalogue. This process is independent on how the VNF has been on-boarded in the local catalogue; suffice to say that the site operational team (or site manager, or any other figure with appropriate permission) does this on-boarding.

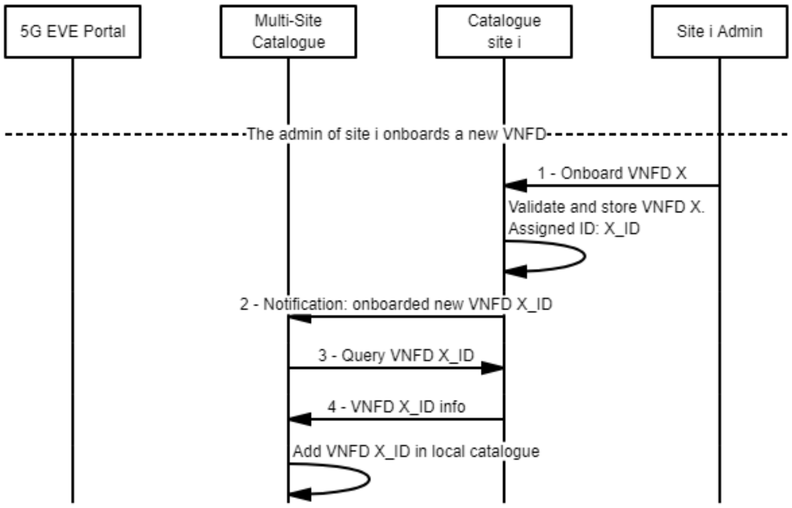

Figure 31: VNF on-boarding workflow

Step 1. Once the new VNF is on-boarded and validated onto the local catalogue (1), a new ID is assigned and a notification is sent upwards (2). As in the previous process, it has to be noted that all these communications between catalogues are proxied by the Multi-site NSO.

Step 2. Once the notification reaches the Multi-site Catalogue, it requests the related information (3), obtaining a response (4) which includes all the attributes required for the appropriate storage.

4.3.3.3 VNF removal

The following (Figure 32) is the workflow for VNF removal from the Multi-site Catalogue. Again, this process is independent on how the VNF has been removed from the local catalogue; suffice to say that the site operational team (or site manager, or any other figure with appropriate permission) does this removal.

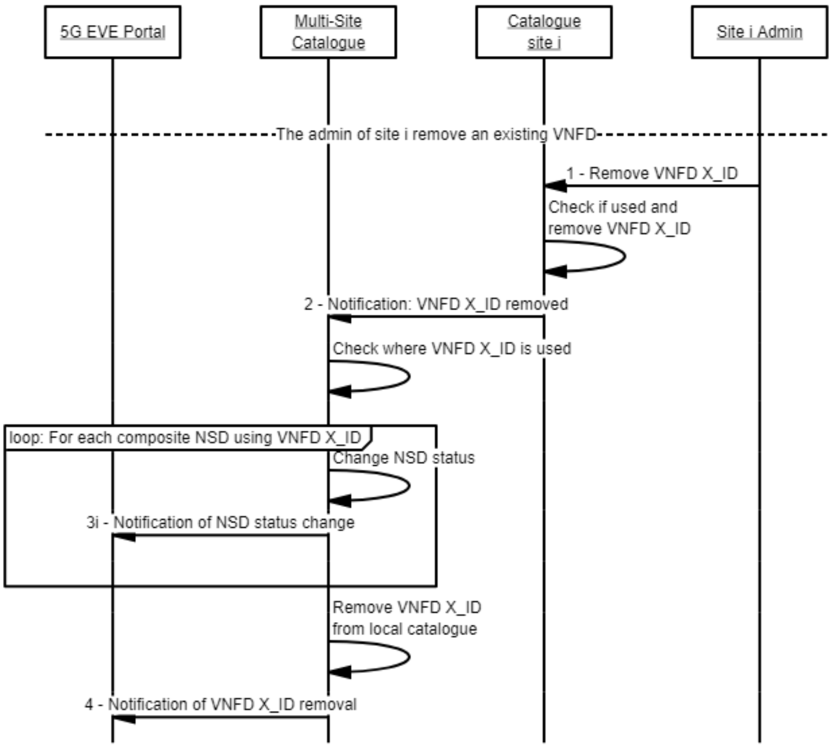

Figure 32: VNF removal workflow

Step 1. Once the new VNF is removed from the local catalogue (1), a notification message is sent upwards containing the VNFD ID (2). As in the previous processes, all these communications between catalogues are proxied by the Multi-site NSO.

Step 2. Once the notification reaches the Multi-site Catalogue, and before actually removing the VNF, the Multi-site Catalogue needs to check if that VNF was part of any available NSD. If so, it changes the status of each NSD including that VNF (3); if not, or when all NSD status have been changed, the Multi-site Catalogue removes the VNF, also dispatching the required notifications (4).

4.3.3.4 NS on-boarding

Due to the complexity of this procedure, three workflows have been designed depending on from where the NSD on-boarding is launched.

The first option is that the NSD is on-boarded from a given site towards the I/W Layer (Figure 33). The procedure is totally equivalent to that of the VNF on-boarding (described in 4.3.3.2), so it will not be repeated here.

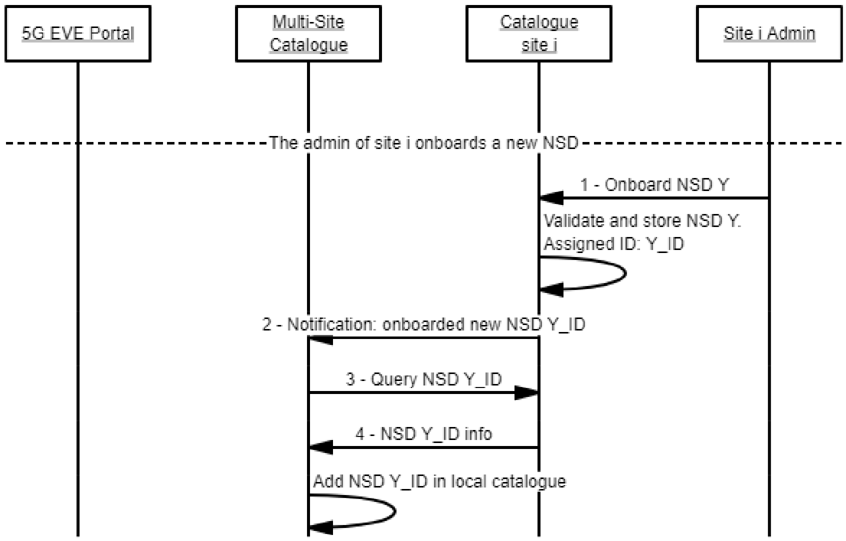

Figure 33: NS on-boarding from a given site workflow

More complex is the procedure in case the NS is on-boarded from the 5G EVE Portal. For a correct understanding of all the implications, the use case has been divided in two:

- The experiment spans across a single site
- The experiment spans across multiple sites

Figure 34 shows the workflow when the experiment only spans across one site. To account for the most general case, the NSD being on-boarded is a composite NS; in this case however, all the VNFDs and nested NSDs are from the same site.

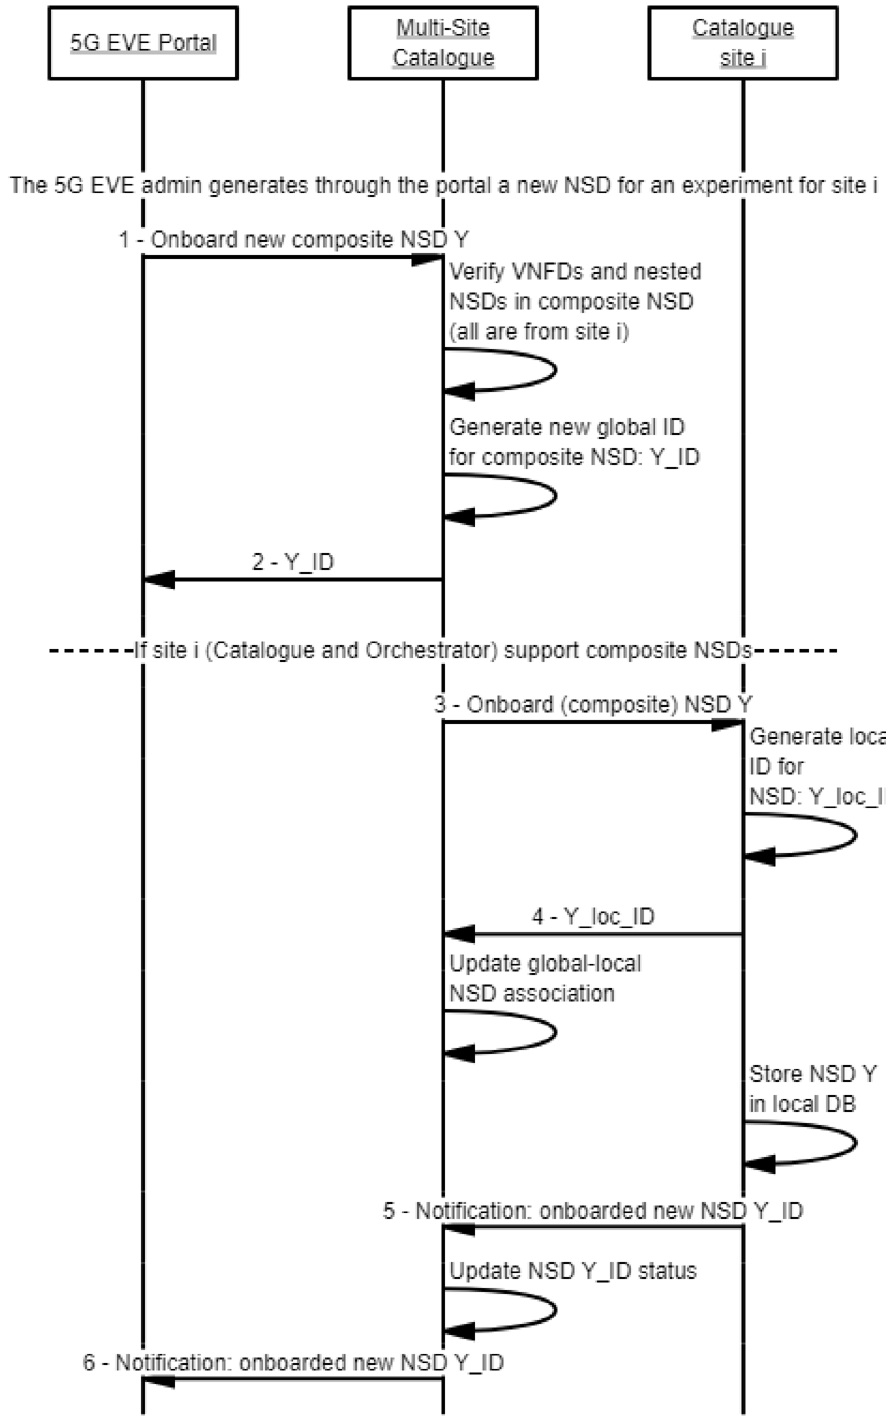

Figure 34: Single-site NS Portal on-boarding workflow

Step 1. Once the new composite NSD is generated at the 5G EVE Portal, its on-boarding in the Multi-site Catalogue is triggered (1). Once all the VNFDs and nested NSD are verified and validated, a new ID is assigned to the composite NSD and a notification is sent back (2).

Step 2. The NSD on-boarded in the Multi-site Catalogue needs also to be on-boarded onto the local catalogue of the site where the experiment is going to be executed. A request is therefore sent (3) which in turn triggers the generation of a local ID in the local catalogue for that NSD. The response (4), which includes such local ID, is used by the Multi-site Catalogue to update the association of global and local IDs. As in the previous process, all these communications between catalogues are proxied by the Multi-site NSO.

Step 3. Once the NSD is stored in the local catalogue, this catalogue sends a notification to the Multi-site Catalogue, advertising the effective on-boarding (5). At this stage, the information about the NSD can be updated in the Multi-site Catalogue, which in turn provides a notification to the Experimentation Portal (6).

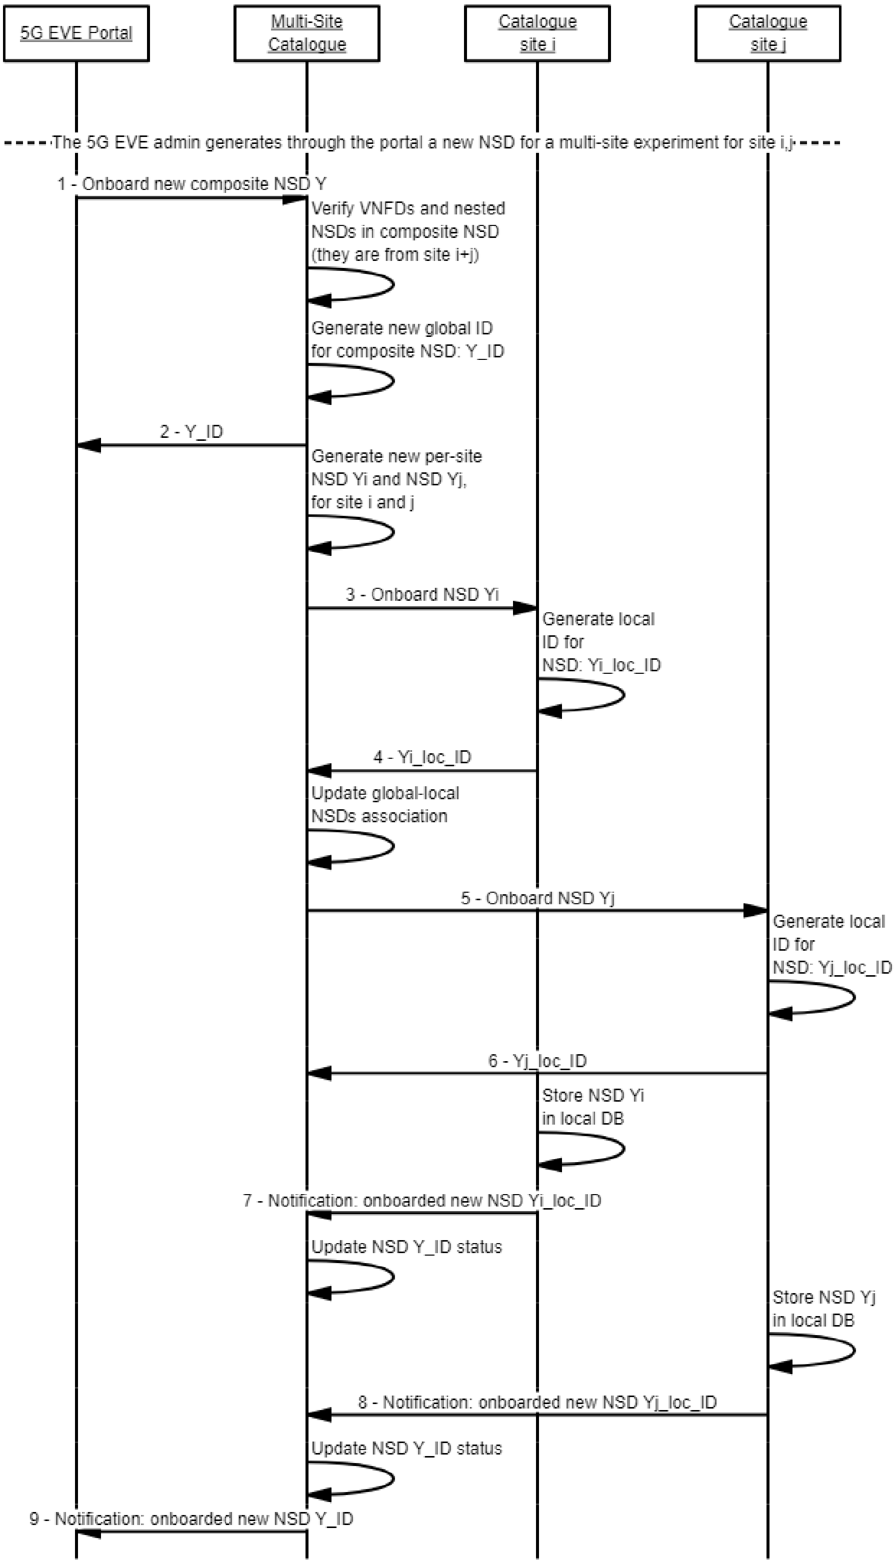

Figure 35: Multi-site NS Portal on-boarding workflow

Finally, Figure 35 above shows the workflow when the experiment spans across multiple sites. Again, to account for the most general case, the NSD being on-boarded is a composite one. In this case however, the VNFs and nested NSDs belong to different sites, which affects the workflow: for each of the participating sites, a NSD is on-boarded as per the workflow above (Figure 34). In other words, the workflow for a multi-site NS is a replication as many times as needed of the process for single-site NSs.

4.3.3.5 NSD removal

Three workflows are also possible for the NSD removal, totally equivalent to those of the on-boarding procedure. To simplify the description, however, and since the multi-site removal from the Portal is a concatenation of the single-site removal process (as it happened in the on-boarding workflow), the former has not been included in this document.

Again, the first option is that the NSD is removed from the site which had previously on-boarded it (Figure 36). The procedure is totally equivalent to that of the VNFD removal (described in 4.3.3.3), so it will not be described.

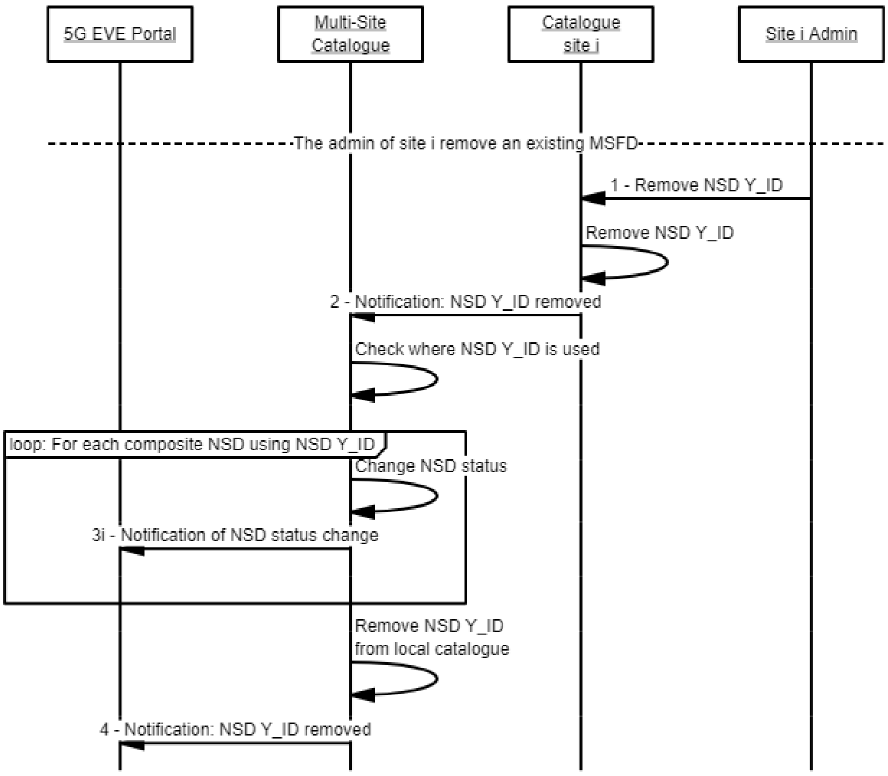

Figure 36: NS removal from a given site workflow

Figure 37 shows the workflow when an experiment spanning across one site is no longer available. To account for the most general case, the NSD being removed is a composite NSD, in this case with all the VNFDs and nested NSDs from the same site.

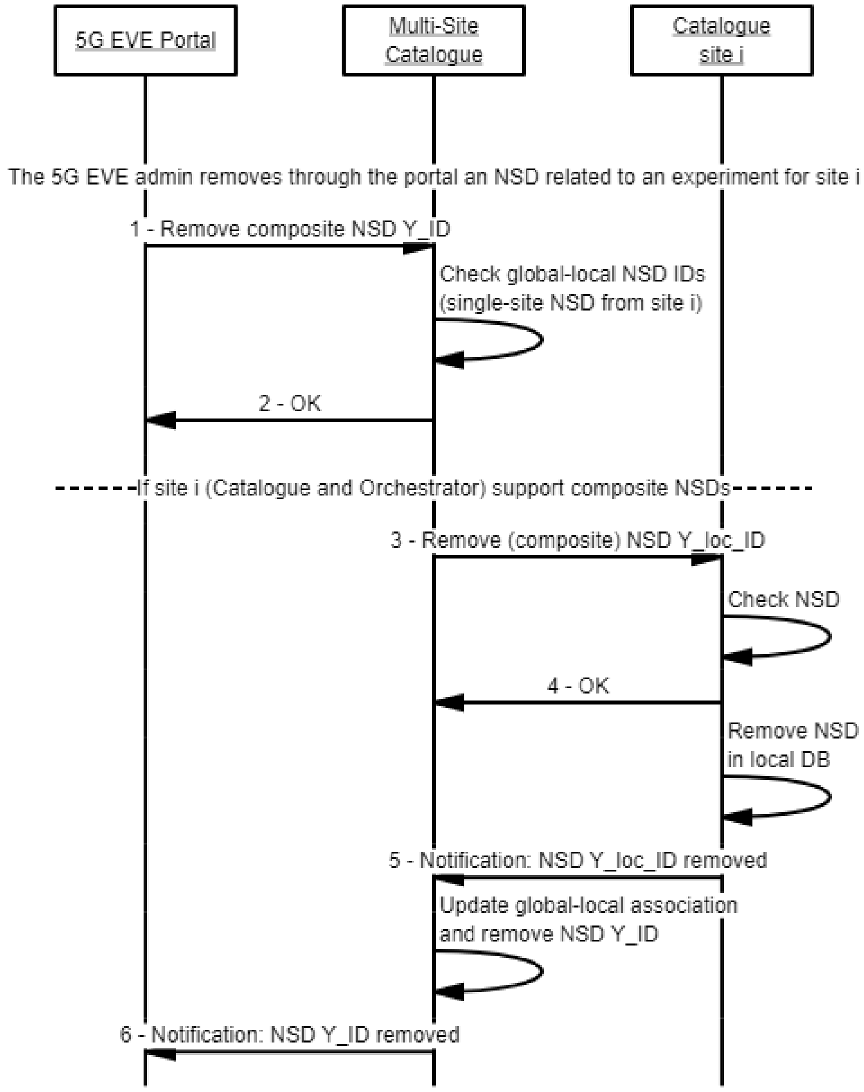

Figure 37: Single-site NS Portal removal workflow

Step 1. Once the composite NSD for an experiment wants to be removed from the 5G EVE Portal, the removal request to the Multi-site Catalogue is triggered (1). Once all the required verifications regarding the nested NSD ID are done, a notification is sent back to the Portal (2).

Step 2. The NSD removal in the Multi-site Catalogue needs to be extended to the local catalogue of the site where the experiment was being executed. A request is therefore sent (3) which in turn triggers the local status checking of the nested NSD ID and a response to the Multi-site Catalogue (4). As in the previous processes, all these communications between catalogues are proxied by the Multi-site NSO.

Step 3. Once the NSD is removed from the local catalogue, this catalogue sends a notification to the Multi-site Catalogue, advertising the effective removal (5). At this stage, the information about the NSD can be updated in the Multi-site Catalogue, which in turn provides a notification to the Experimentation Portal (6).

4.3.3.6 Monitoring exchange

The following diagram (Figure 38) represents the initial workflow with regards to the monitoring operations.

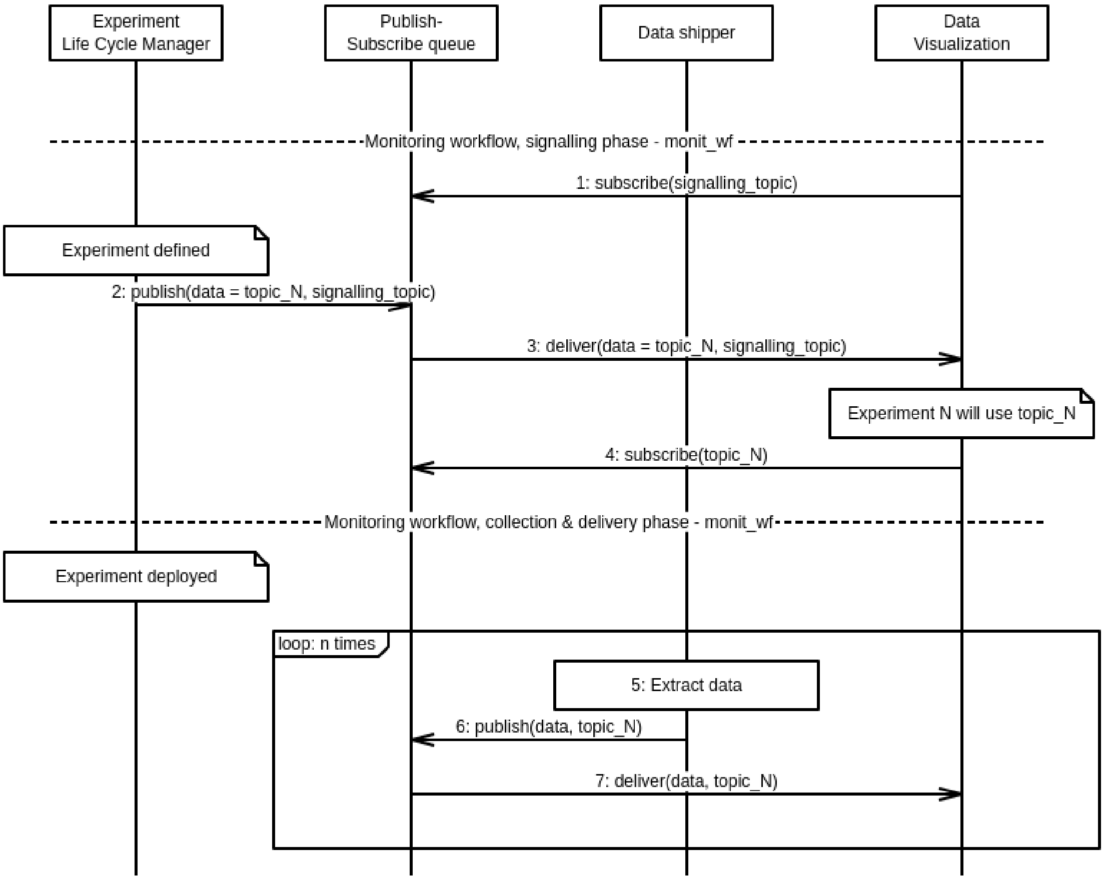

Figure 38: Workflow for the monitoring data exchange

In the workflow it is shown how, before defining an experiment, the components that are going to play the subscriber role (e.g. data visualization components at the Portal) send a subscription message to the publish/subscribe queue (located in the I/W Framework). This is done using the “signalling_topic” topic, which will be used as a sort of signalling channel for notifications to be transmitted to the subscribers in the publish/subscribe queue.

After the definition of the experiment, the Experiment Life Cycle Management from the Portal will publish data in the “signalling_topic” topic, notifying that a new experiment has been defined and will use the “topic_N” topic. Automatically, this information is delivered to the components subscribed to “signalling_topic” (in this case, the data visualization module), making them aware of the relationship between the experiment N and the usage of the “topic_N” in the publish/subscribe queue. As a result, the subscribers will be subscribed to the “topic_N” topic.

When the experiment is deployed, the traffic shippers in the site facilities where experiment N is being executed will extract the monitored data, publishing it to the publish/subscribe system and delivering that data to the “topic_N” subscribers.

4.3.3.7 Runtime configuration

In the following diagram (Figure 39), the experiment definition and execution process in detail is depicted, with all the components involved: the Portal (green boxes for service definition; red boxes for service execution) and the I/W Layer (in yellow).

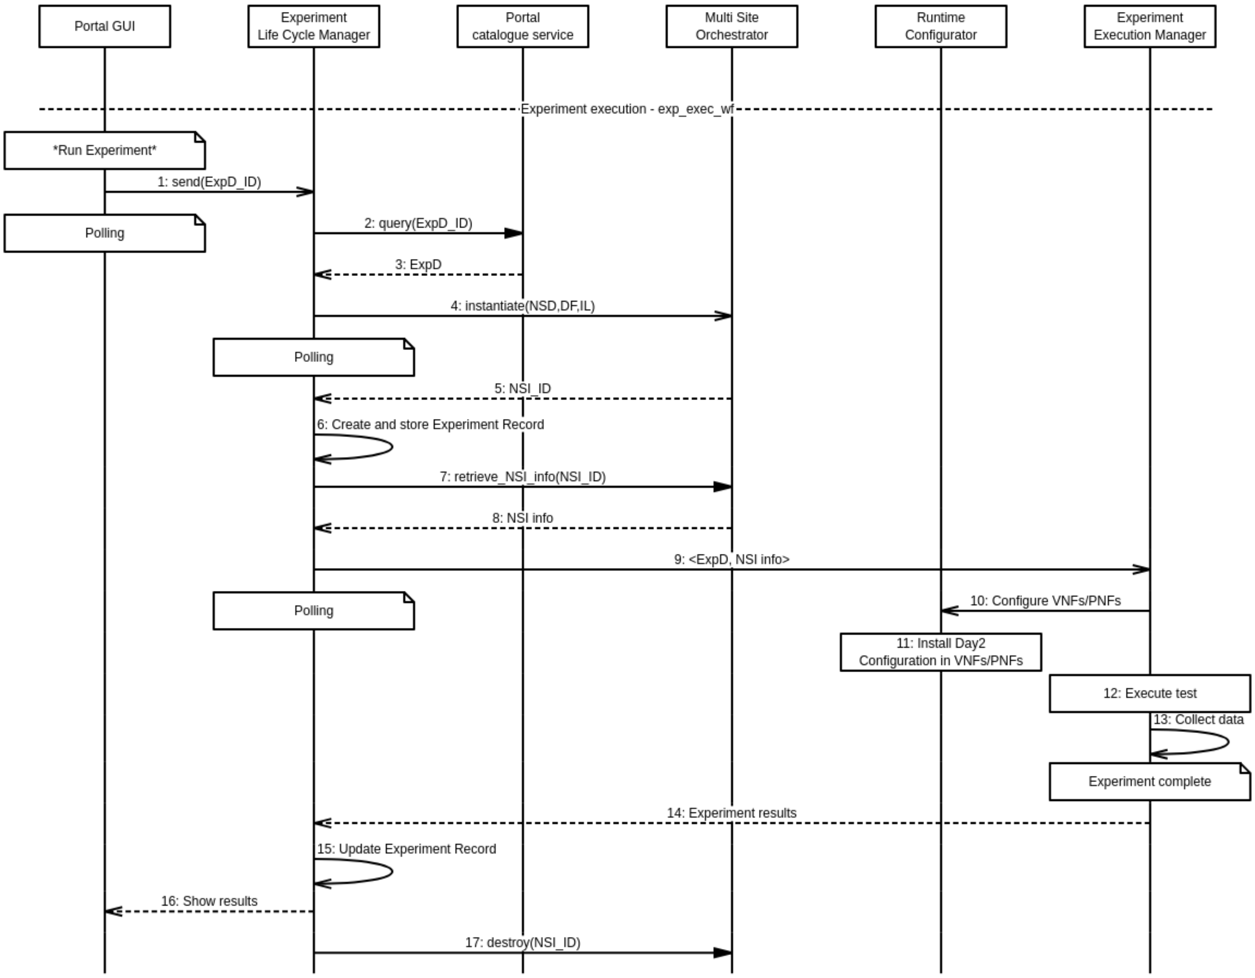

Figure 39: Workflow for the runtime configuration

Regarding the Runtime Configurator, the main steps start at flow number 9, where the Experiment Execution Manager receives all the information related to the experiment. This will be used to apply the Day2 configuration to the different VNFs and PNFs in flow number 10, which in turn will cause the installation of the provided configurations in all the VNFs and PNFs in parallel (11). After this process is finished, the experiment can be executed, process which also involves the Data Collection mechanism commented in the previous section.

4.4 Interworking Framework NBI: Interworking APIs

The 5G EVE Interworking Framework exposes at its northbound reference point a wide set of APIs for the integration with the upper 5G EVE components participating in the deployment and execution of multi-site experiments (mainly the 5G EVE Portal and the KPI Collection & Validation Frameworks). While the 5G EVE Portal governs the Vertical experiment definition and execution, the I/W Framework is responsible for the subsequent provisioning in the site facilities of those NS and slices that support the single- and multi-site experiments.

The Interworking APIs are therefore the main enablers for the fulfilment of the whole Vertical experiment lifecycle. The following subsections provide an overview of the operations and information models supported by the I/W Framework’s functional components at their NBI, in particular: i) the Multi-site Catalogue, ii) the Multi-site Inventory, iii) the Multi-site NSO, iv) the Runtime Configurator, v) the Data Collection Manager.

4.4.1 Multi-site Catalogue operations and information models

The 5G EVE Multi-site Catalogue maintains the relevant set of descriptors and deployment information required by the I/W Framework (in particular by the Multi-site NSO) for managing the lifecycle of the multi-site services and slices in support of the multi-site Vertical experiments. The Multi-site Catalogue NBI is the main entry point for on-boarding new multi-site Network Service Descriptors and querying available VNF Descriptors for enabling the 5G EVE Portal to compose and define multi-site experiments.

As a general principle, the interfaces, operations and information models supported by the Multi-site Catalogue at its northbound reference point leverage on existing definitions and specifications from relevant standards. In particular, the ETSI NFV SOL005 [10] APIs and data models for Network Service Descriptors and VNF Package Management are taken as reference for the various on-boarding related operations, while the ETSI NFV SOL001 [4] specification is used as reference for the packaging of the different descriptors.

In terms of operations, it is relevant to highlight that the on-boarding of VNFs is not supported at the Multi-site Catalogue northbound interface, as VNFs in 5G EVE are expected to be on-boarded directly on each local catalogue and retrieved from there by the Multi-site Catalogue. On the other hand, the Multi-site Catalogue allows through its northbound APIs to retrieve and collect information about the available Network Service and VNF Descriptors.

Table 7 below summarizes the Multi-site Catalogue northbound APIs, in terms of exposed operations, information models and exchanged information.

Table 7: Multi-site Catalogue Interworking APIs

Multi-Site Catalogue APIs		
Network Service Descriptor Management		
Description	The Network Service Descriptor (NSD) Management API allows the 5G EVE portal to manage NSDs that model the Vertical experiments. It is envisaged that NSDs can be either on-boarded by the 5G EVE portal or retrieved by the Multi-site Catalogue from the 5G EVE sites.	
Reference Standards	ETSI NFV SOL005, ETSI NFV SOL001	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve NSD	This operation allows to retrieve an existing NSD from the Multi-site Catalogue. The main information exchanged is: Input: • NSD identifier(s) Output: • NS type, identifying the type of service/experiment implemented by the given NSD • Operational status • List of VNF Descriptors composing the NSD • List of PNF Descriptors composing the NSD • List of monitoring metrics that can be collected from the NS • NS Forwarding Graph, describing how VNFs and PNFs interconnect • Deployment constraints, in case the NS is available for deployment in specific locations (i.e. 5G EVE sites) • List of runtime configurations that can be applied on the NS	ETSI NFV SOL001, extended with: • Testing tools information • Runtime configurations exposed • Deployment constraints, only if the NSD was retrieved from one of the 5G EVE site catalogues

On-board new NSD	This operation allows to register a new NSD in the Multi-site Catalogue, representing a new experiment description. The main information exchanged is: Input: • The new NSD to be on-boarded, including the information described in the output of the "Retrieve NSD" operation Output: • NSD identifier	ETSI NFV SOL001, extended as described for the "Retrieve NSD" operation:
Modify NSD	This operation allows to update the operational status of an existing NSD, as a way to temporarily enable or disable its usage. The main information exchanged is: Input: • NSD identifier • New NSD operational status Output • None	ETSI NFV SOL005
Delete NSD	This operation allows to delete an existing NSD from the Multi-site Catalogue. The main information exchanged is: Input: • NSD identifier(s) Output: • None	ETSI NFV SOL005
Create NSD Subscription	This operation allows to subscribe to NSD related notifications generated by the Multi-site Catalogue. The main information exchanged is: Input: • NSD identifier(s) for which the subscription is valid • Notification filter, to select which types of notifications to subscribe to (e.g. NSD status change) Output: • NSD subscription identifier	ETSI NFV SOL005
VNF Descriptor Management		
Description	The VNF Descriptor (VNFD) Management API allows the 5G EVE Portal to retrieve information about VNFs available in the 5G EVE sites. It is not envisaged to support on-boarding of new VNFs through this NBI API. This API is also used to retrieve information about Testing Tools that are available in the 5G EVE sites as VNFs, e.g. traffic generators, UE emulators.	
Reference Standards	ETSI NFV SOL005, ETSI NFV SOL001	
Operations Exposed	***Information Exchanged***	***Information Model***

<table>
<tr><td>Retrieve VNFD</td><td>This operation allows to retrieve an existing VNFD from the Multi-site Catalogue. The VNFDs are intended to be collected by the Multi-site Catalogue from each 5G EVE site catalogue. The main information exchanged is:
Input:
• VNFD identifier(s)
Output:
• VNF type, identifying the type of application provided by the VNF
• Operational status
• List of monitoring metrics that can be collected from the VNF
• Deployment constraints, in case the VNF is available for deployment in specific locations (i.e. 5G EVE sites)
• List of runtime configurations that can be applied on the VNF</td><td>ETSI NFV SOL001, extended with:
• Runtime configurations exposed
• Deployment constraints</td></tr>
<tr><td>Create VNFD Subscription</td><td>This operation allows to subscribe to VNFD related notifications generated by the Multi-site Catalogue. The main information exchanged is:
Input:
• VNFD identifier(s) for which the subscription is valid
• Notification filter, to select which types of notifications to subscribe for (e.g. VNFD status change)
Output:
• VNFD subscription identifier</td><td>ETSI NFV SOL005</td></tr>
<tr><td colspan="3">PNF Descriptor Management</td></tr>
<tr><td>Description</td><td colspan="2">The PNF Descriptor (PNFD) Management API allows the 5G EVE Portal to retrieve information about PNFs available in the 5G EVE sites. It is not envisaged to support on-boarding of new PNFs through this NBI API. This API is also used to retrieve information about Testing Tools that are available in the 5G EVE sites as PNFs, e.g. traffic generators, UE emulators</td></tr>
<tr><td>Reference Standards</td><td colspan="2">ETSI NFV SOL005, ETSI NFV SOL001</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve PNFD</td><td>This operation allows to retrieve an existing PNFD from the Multi-site Catalogue. The PNFDs are intended to be collected by the Multi-site Catalogue from each 5G EVE site catalogue. The main information exchanged is:
Input:
• PNFD identifier(s)
Output:
• PNF type, identifying the type of application provided by the PNF
• Operational status
• List of monitoring metrics that can be collected from the PNF
• Location constraints, to describe where the PNF is available for use (i.e. 5G EVE sites)
• List of runtime configurations that can be applied on the PNF</td><td>ETSI NFV SOL001, extended with:
• Runtime configurations exposed
• Location constraints</td></tr>
</table>

Create PNFD Subscription	This operation allows to subscribe to PNFD related notifications generated by the Multi-site Catalogue. The main information exchanged is: Input: • PNFD identifier(s) for which the subscription is valid • Notification filter, to select which types of notifications to subscribe for (e.g. PNFD status change) Output: • PNFD subscription identifier	ETSI NFV SOL005
UE and SIMs Management		
Description	The UE and SIM Management API allows the 5G EVE Portal to retrieve information about UEs and subscriptions that are available in the 5G EVE sites to run the experiments. This information is intended to be manually configured through admin operations in the Multi-site Catalogue. Indeed, each 5G EVE site is responsible for the management of its set of UEs and subscriptions, and the related data is not expected to be automatically fed into the I/W Framework	
Reference Standard	None	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve UE	This operation allows to retrieve the available UEs from the Multi-site Catalogue. The main information exchanged is: • The type of UE • The type of radio interface (e.g. 5G/4G support) • Operational status • Location constraints, to describe where the UE is available for use (i.e. 5G EVE sites) • List of subscriptions that can be associated with the UE	No standard available
Retrieve SIM	This operation allows to retrieve the available subscriptions and SIMs from the Multi-site Catalogue. The main information exchanged is: • The type of subscription (e.g. 5G/4G) • Location constraints, to describe where the subscription is available for use (i.e. 5G EVE sites, or location within a site) • List of profiles that can be activated for the subscription (e.g. related to eMBB, mMTC, uRLLC capabilities)	No standard available

4.4.2 Multi-site Inventory operations and information models

The Multi-site Inventory exposes at its northbound reference point a set of APIs for retrieving runtime information related to deployed vertical experiments. In particular, Network Service and VNF instances provisioned in the site facilities in support of single- and multi-site Vertical experiments are maintained in the Multi-site Inventory. They are also offered to the Portal for visualization purposes, as well as for enabling re-use and sharing of existing services and slices among different experiments.

Table 8 below summarizes the Multi-site Inventory northbound APIs, in terms of exposed operations, information models and exchanged information.

Table 8: Multi-site Inventory Interworking APIs

<table>
<tr><th colspan="3">Multi-Site Inventory APIs</th></tr>
<tr><td colspan="3">Network Service Instance Management</td></tr>
<tr><td>Description</td><td colspan="2">The Network Service (NS) Instance Management API allows the 5G EVE Portal to retrieve information about deployed NSs in support of vertical experiments. A NS Instance can either relate to the full vertical experiment (i.e. including all its components) or part of it (e.g. only the vertical application components)</td></tr>
<tr><td>Reference Standards</td><td colspan="2">ETSI NFV SOL005</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve NS Instance</td><td>This operation allows to retrieve information related to a running NS Instance from the Multi-site Inventory. The main information exchanged is:
Input:
• NS Instance identifier(s)
Output:
• NS type, identifying the type of service/experiment implemented
• Operational status
• List of VNF Instances part of the NS Instance (identifiers)
• List of PNF Instances part of the NS Instance (identifiers)
• NS Forwarding Graph, describing how VNFs and PNFs interconnect
• List of nested NS Instances</td><td>ETSI NFV SOL005, extended with:
• Testing tools information</td></tr>
<tr><td colspan="3">VNF Instance Management</td></tr>
<tr><td>Description</td><td colspan="2">The VNF Instance Management API allows the 5G EVE Portal to retrieve information about VNFs deployed in the 5G EVE end-to-end infrastructure and running in one or more 5G EVE sites.</td></tr>
<tr><td>Reference Standards</td><td colspan="2">ETSI NFV SOL005</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve VNF Instance</td><td>This operation allows to retrieve information related to a deployed VNF instance from the Multi-site Inventory. The main information exchanged is:
Input:
• VNF Instance identifier(s)
Output:
• VNF type, identifying the type of application provided by the VNF
• Operational status</td><td>ETSI NFV SOL005</td></tr>
<tr><td colspan="3">PNF Instance Management</td></tr>
<tr><td>Description</td><td colspan="2">The PNF Instance Management API allows the 5G EVE Portal to retrieve information about PNFs deployed and configured in the 5G EVE end-to-end infrastructure and used as part of Vertical experiments (i.e. in a Network Service Instance)</td></tr>
<tr><td>Reference Standard</td><td colspan="2">ETSI NFV SOL005</td></tr>
</table>

Operations Exposed	*Information Exchanged*	*Information Model*
Retrieve PNF Instance	This operation allows to retrieve information related to a PNF instance (deployed, configured and part of a NS) from the Multi-site Inventory. The main information exchanged is: Input: • PNF Instance identifier(s) Output: • PNF type, identifying the type of application provided by the PNF • Operational status	ETSI NFV SOL005

4.4.3 Experiment Execution operations and information models

Two are the main I/W Framework modules involved in the deployment and execution of experiments: the Multi-site Network Service Orchestrator (NSO) and the Runtime Configurator.

4.4.3.1 Multi-site Network Service Orchestrator

The Multi-site Network Service Orchestrator exposes a northbound interface for the definition and deployment of network services across the 5G EVE Sites. The interface is based on ETSI NFV-SOL 005 definition, as detailed in Table 9, although it could be extended during the development phase in order to include specific 5G EVE parameters in the NS instance information structure

Table 9: Multi-site Network Service Orchestrator APIs

Multi-Site Network Service Orchestrator APIs		
NS Identifier management		
Description	The NS Identifier API allows the 5G EVE Portal to create and delete Network Services	
Reference Standard	ETSI NFV SOL 005	
Operations Exposed	***Information Exchanged***	***Information Model***
Create NS Identifier	This operation allows to create a Network Service Identifier in the Multi-Site Orchestrator. Input Parameters: • Reference to the NSD that describes the NS • NS name and description Output • Result of the operation • NS information	ETSI NFV SOL005
Delete NS Identifier	This operation allows to delete a Network Service Identifier. Only valid for NS in NOT_INSTANTIATED state Input: • NS instance ID Output • Result of the operation	ETSI NFV SOL005

NS management		
Description	The NS management API allows the 5G EVE Portal to manage Network Services, by instantiating, terminating, scaling and updating the NS	
Reference Standard	ETSI NFV SOL 005	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve NS information	This operation allows to retrieve the information description of a particular NS or for all NS. Input • NS instance ID(s) • Filters for defining the output Output • One or more NS information fields	ETSI NFV SOL005
Instantiate NS	The instantiate NS operation allows to instantiate a network service. Input • Flavour • PNF and VNF Information Output • Immediate result code (final result will be communicated through a notification)	ETSI NFV SOL005
Scale NS	The Scale NS allows to scale an existing NS by adding VNF instances or by scaling a particular VNF Input • The necessary information to scale the referenced NS or VNF Output • Immediate result code (final result will be communicated through a notification)	ETSI NFV SOL005
Update NS	The update NS allows to modify an existing Network Service Input • List of PNFs and VNFs to be updated • NS Forwarding Graph, describing how VNFs and PNFs interconnect • List of nested NS Instances Output • Immediate result code (final result will be communicated through a notification)	ETSI NFV SOL005
Terminate NS	The Terminate NS allows to terminate an existing NS Input • When (time) Output • Immediate result code (final result will be communicated through a notification)	ETSI NFV SOL005

NS notifications		
Description	The NS notifications API allows the Multi-site Network Service Orchestrator to notify 5G EVE Portal regarding subscribed changes in NS status and lifecycle	
Reference Standard	ETSI NFV SOL 005	
Operations Exposed	***Information Exchanged***	***Information Model***
NS Identifier Creation Notification	Multi-site Network Service Orchestrator notification of the creation of a NS. Information: • NS instance identifier of the created NS	ETSI NFV SOL005
NS Identifier Deletion Notification	Multi-site Network Service Orchestrator notification of the deletion of a NS. Information: • NS instance identifier of the deleted NS	ETSI NFV SOL005
NS LifeCycleChange Notification	Multi-site Network Service Orchestrator notification of a lifecycle change in a deployed Network Service. Information: • NS Instance ID • LCM operation • Status (START, RESULT) • Information of NS, VNF, PNF and connectivity affected	ETSI NFV SOL005

4.4.3.2 Runtime Configurator

The Runtime Configurator northbound API is mainly focused on the reception of experiment-specific and Vertical-oriented Day2 configurations for the Network Services and VNFs/PNFs already deployed in the underlying infrastructure, according to the data provided by the 5G EVE Portal. Therefore, the operations considered for this module are related with the experiment definition, involving specific changes in the deployed components at the application level. Examples of these changes are enabling/disabling functionalities/configurations or executing specific updates/upgrades of certain functionalities/configurations.

Other complex operations closer to the service definition or lifecycle have been contemplated in the Multi-site Network Service Orchestrator (previous section), as they are out of the scope of the Day2 configuration itself.

As already stated in section 4.2.4, the data and information models to be used in the Runtime Configurator communications depend on the PNFs that sites expose to the I/W Layer, and the type of deployed NSs and VNFs. Therefore, the amount of possibilities is extremely large, and depends on the number of Vertical Use Cases that want to be supported. In that sense, the API specification shown in Table 10 below is a very high-level one. Details will be provided in future Deliverables, as low-level features are implemented and tested.

Table 10: Runtime Configurator Interworking API

Runtime Configurator API	
Update of current instances	
Description	Modification of specific configurations in a running experiment according to the information provided by the Vertical in the experiment definition.
Reference Standards	Agentless, push paradigm

Operations Exposed	*Information Exchanged*	*Information Model*
Enable a functionality or configuration	• VNF/PNF ID(s) • Functionality/configuration to enable • Specify if it is a disabled functionality/configuration or a new one to be installed	Mapping of the information provided by Verticals in configuration templates to be executed by a configurator tool
Disable a functionality or configuration	• VNF/PNF ID(s) • Functionality/configuration to be disabled	Mapping of the information provided by Verticals in configuration templates to be executed by a configurator tool
Update/Upgrade of a functionality or configuration	• VNF/PNF ID(s) • Functionality/configuration to be updated/upgraded.	Mapping of the information provided by Verticals in configuration templates to be executed by a configurator tool

4.4.4 Experiment Results operations and information models

The Data Collection Manager is the only module in the I/W Framework involved in the gathering of results from executed experiments.

4.4.4.1 Data Collection Manager

The Data Collection Manager northbound API is the interface with the 5G EVE Experimentation Portal in charge of handling experiment results from the underlying Vertical and network functions. This northbound interface is initially related to the execution of tests, and not designed for other purposes, like infrastructure monitoring for example. This way, we can control the type of information experimenters have access to. However, an equivalent interface could be perfectly well designed for monitoring purposes, with the main difference of data not being sent to the Experimentation Portal (instead, maybe, it could be sent to an Operation Portal).

The main operations included in Table 11 below follow the publish/subscribe paradigm already described in section 4.2.5, in which there are two main entities: **publishers** (they publish data related to a certain topic) and **subscribers** (they subscribe to a topic).

In this project, subscribers are placed in the 5G EVE Portal mainly (e.g. data visualization tools or experiment execution tools), but there could be other components in the I/W Layer that could also interested in specific topics (e.g. the orchestrator lifecycle manager). The access mechanisms for these would be equivalent to those for the Portal.

On the other hand, the main publishers come from the different services (VNFs and PNFs) deployed in each site facility, which are the main sources of the experimentation data. Similarly, there could be other components that could also publish meaningful information (for example, the signalling function of the Experiment Life Cycle Manager). Again, the procedures for these would be equivalent to those for VNFs/PNFs.

It is interesting to note that the Data Collection Manager’s northbound and southbound interfaces are practically the same thanks to the publish/subscribe paradigm, simplifying the required operations in both interfaces: when a high-level subscriber is subscribed to a topic, the publish/subscribe system will permit the automatic delivery of the published data directly from the low-level publishers.

As for the Runtime Configurator, the range of potential publishers in 5G EVE is extremely large, which prevents from defining a detailed specification of this interface at this stage. It will be built as required by the project Verticals, and details will be provided in future Deliverables, as low-level features are implemented and tested.

Table 11: Data Collection Manager Interworking API

Data Collection Manager API		
Access to experimentation data		
Description	Application-independent access to monitored experimentation data obtained during the execution of use cases. Monitored data can be related to network conditions in general, or Network/Service KPIs. The former will be valid to ensure that conditions are adequate for the execution of the test. The later will also be used for validation purposes.	
Reference Standards	Interface based on Publish/Subscribe model. Regarding network monitoring: RMON (RFC2819, extensions and updates), Syslog (RFC5424, extensions and updates), SNMP (RFC1157, extensions and updates). Regarding testing: NGMN Alliance "Definition of the testing framework for the NGMN 5G pre-commercial networks trials"	
Operations Exposed	***Information Exchanged***	***Information Model***
Topic subscription	• Topic(s) to be subscribed (which are related to different parameters/metrics from a specific use case and component(s) that are monitored).	Publish/subscribe mechanism, based on topics. When publishers send data to a certain topic, it is automatically delivered to the interested subscribers in a unified format (JSON chain).
Topic un-subscription	• Topic(s) to be unsubscribed.	This operation would be the un-subscription to a topic.
Publish signalling data	• Signalling topic. • Signalling information to be provided to the subscribers.	This operation is related to the monitoring workflow described in section 4.3.3.6.

4.5 Interworking Framework SBI: Adaptation Layer

The Adaptation Layer at the SBI provides a common access interface to the individual site facilities' services and APIs for the orchestration. Therefore, it exposes to the components of the I/W Framework a set of common internal APIs and models for accessing per-site management, control, orchestration and monitoring services, and translates them into the site-specific APIs and models. Each component of the I/W Framework, based on its requirements, can implement internally specific functions of the Adaptation Layer.

4.5.1 Interworking architectures and relations with ETSI NFV-IFA

In the context of 5G EVE, the Multi-site NSO uses the resources exposed by the participating sites via the SBI to deploy the NS Instances that comprise a Network Slice Instance. Note that the laid out resources may consist of virtualized infrastructure resources or network functions, which are provided respectively by a local VIM in mode of NFVI as a service (NFVIaaS) and by a local orchestrator NFVO in the mode of NS as a service (NSaaS). To enable the Multi-site NSO to control the sites exposed resources, the SBI shall encompass resource management functions, such as discovery function, selection function, allocation function and monitoring function, dedicated to each resource type.

In addition, the internal APIs of the SBI should be abstract and generic regardless of the technology used by each site to expose its resources, so the Multi-site NSO can be agnostic of the site specificity. To achieve that goal, the SBI relies on the Adaptation Layer to implement the abstract interfaces using the site features.

As aforementioned, resources controlled by the Multi-site NSO can be:

- A virtualized infrastructure resource: in this case, a local manager (VIM) provides to the orchestrator some amounts of virtual compute capacities (virtual CPU, virtual RAM) and virtual storage space. Moreover, a local SDN controller, if present, may also provide virtual network resources over the WAN connecting two datacentres.
- A virtual network service: in this case, the site orchestrator exposes its catalogue of Network Services.

The Adaptation Layer of the SBI exposes the Data Models and Information Models collected by the drivers as a homogeneous data structure for the different components of the Multi-site NSO. The SBI through the drivers exposes CRUD operations for services Life Cycle Management (LCM) across multiple administrative domains, in order to deliver mainly two uses cases:

- NFVI as a service, in which a client uses the Multi-site NSO to deploy VNFs (available in the Multi-site Catalogue) on the sites.
- NS as a service, in which a client deploys advertised NSs on the sites.

This section details how different ETSI NFV-IFA documents can be applied to the context of the I/W Framework's SBI. The selected documents covering such communications are: NFV-IFA 028 [32], NFV-IFA 005 [33], NFV-IFA 006 [34], NFV-IFA 013 [9] and NFV-IFA 030 [12], as shown in Figure 40.

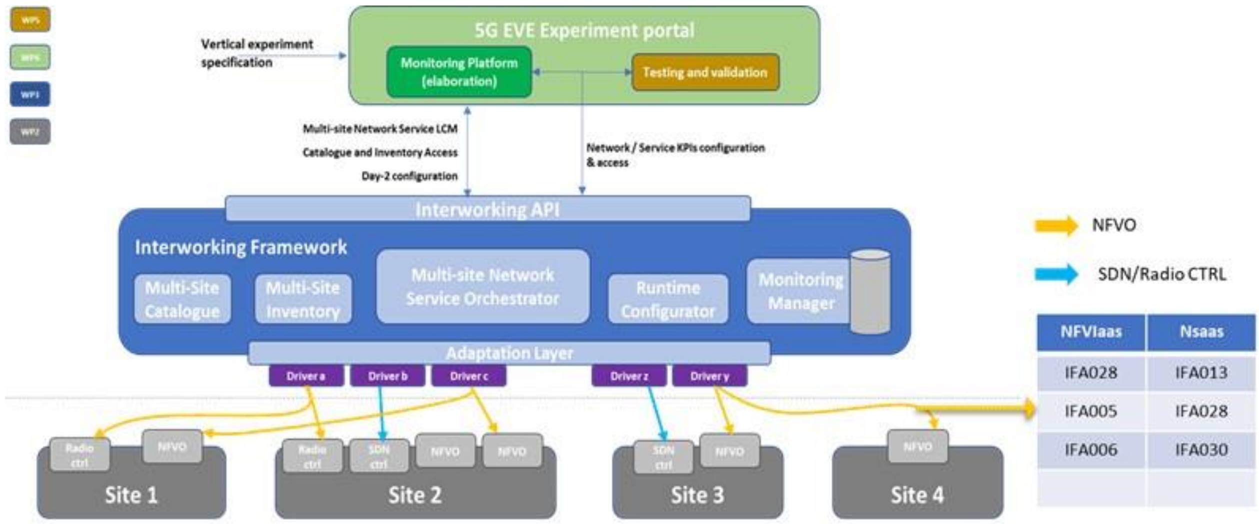

Figure 40: ETSI NFV-IFA as a reference for the SBI

4.5.1.1 NFVI as a Service (NFVIaaS)

In this use case, a parallelism can be made with the description at NFV-IFA 028. The sites offering NFVI services are called NFVI providers and the Multi-site NSO is the NFVI consumer. The Multi-site NSO leverages via the SBI the computing, storage and network resources offered by NFVI providers (sites) to deploy VNFs within a single or multiple sites Figure 41.

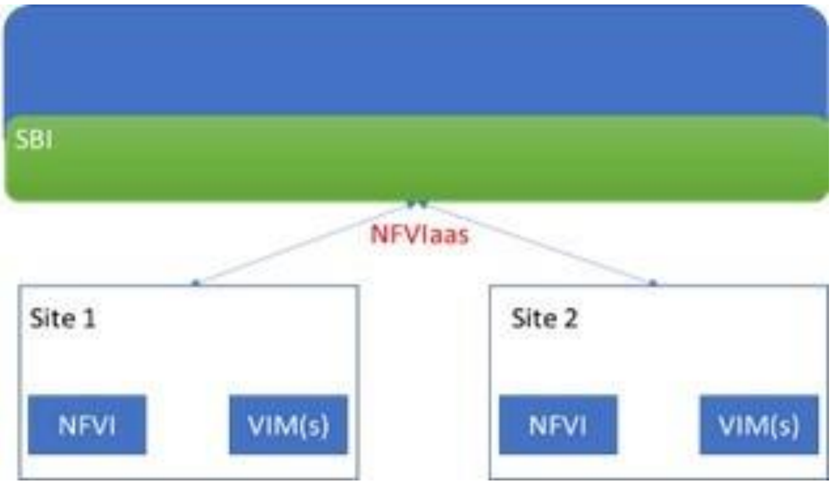

Figure 41: NFVI as a service

The main split of responsibilities for NFVIaaS is as follows:

- Each NFVIaaS provider (site) controls its own resources, which means that the Multi-site NSO has no control over the sites underlying infrastructures.
- The Multi-site NSO, as said, plays the role of the NFVIaaS consumer and therefore its responsibilities are adapted in accordance with the NFV-IFA 028 specification. In the context of 5G-EVE, the Multi-site NSO leverages the SBI drivers to request resources from the sites to run VNFs and connect them to NSs.

The definition of Multi-site NSO as consumer and sites as providers allows for two architectural options.

- ➢ A SLPOC (Single Logical Point of Contact) is integrated into the NFVO as shown in Figure 42, where the multiple VIMs in the NFVIaaS provider are hidden from the consumer. The Multi-site NSO interfaces directly with the SLPOC via the SBI implementations of NFV-IFA 006, NFV-IFA 005 and NFV-IFA 013.

  The SLPOC can also be implemented in scenarios with a single VIM. In that case, the SLPOC can be implemented at the VIM itself, instead of at the NFVO.

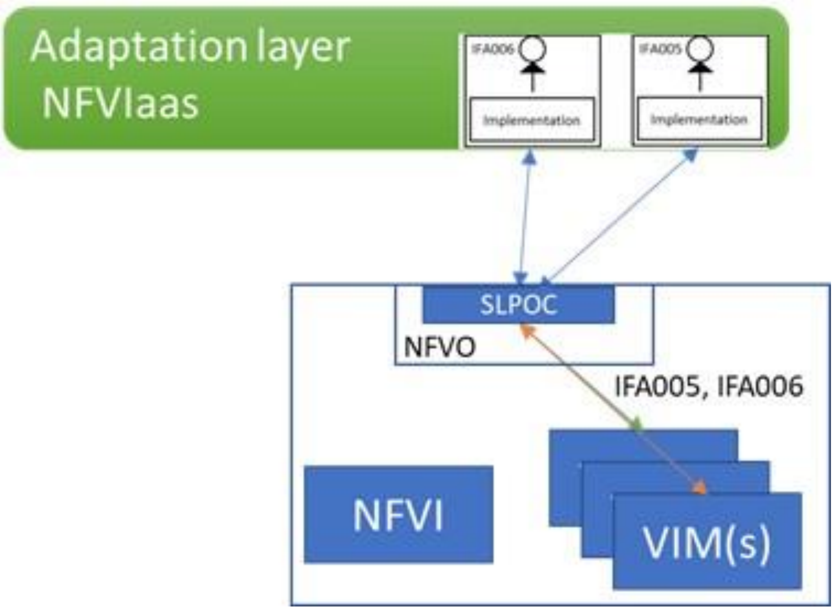

Figure 42: NFVIaaS single point of contact implemented into the NFVO

- It is also possible that a site chooses to use a MLPOC (Multiple Logical Point of Contact). In this case the NFVIaaS provider allows direct access to its VIMs, as shown in Figure 43.

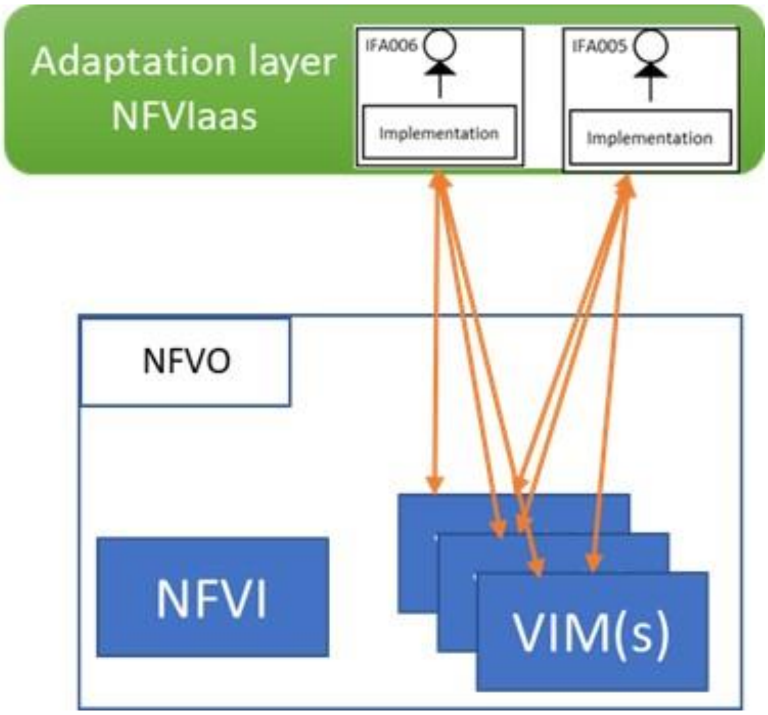

Figure 43: NFVIaaS multiple points of contact

There are implementation implications that should be considered when choosing between SLPOC or MLPOC architectural options. Those implications are discussed in NFV-IFA 028. In 5G EVE, we will focus on the SBI with a local NFVO that implements a SLPOC, as shown in Figure 44.

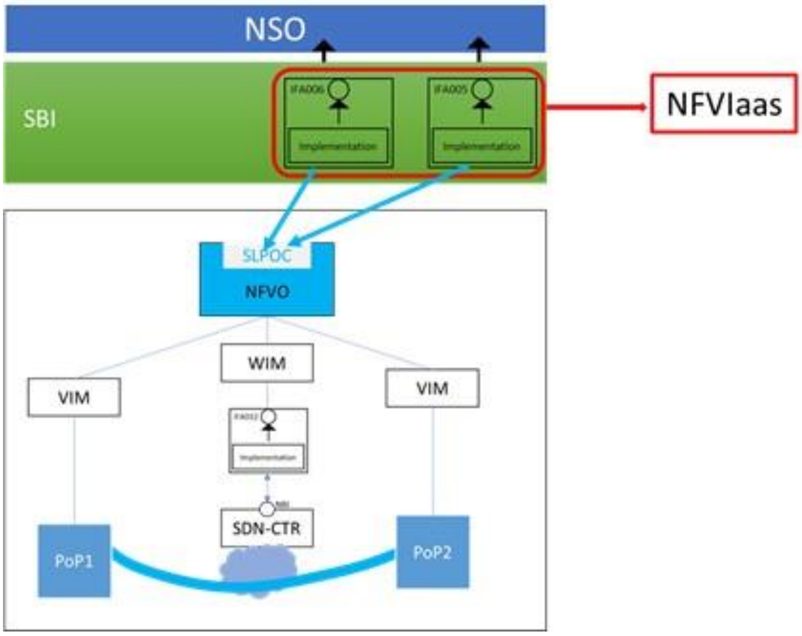

Figure 44: NFVIaaS SBI with NFVO SLPOC

More broadly speaking, the goal of 5G EVE to interconnect sites that have heterogeneous stacks of NFVOs and VIMs (Figure 45) demands specific agreed implementations of the northbound interfaces of each NFVO. These specifications will be covered in a later section of the document (4.5.2).

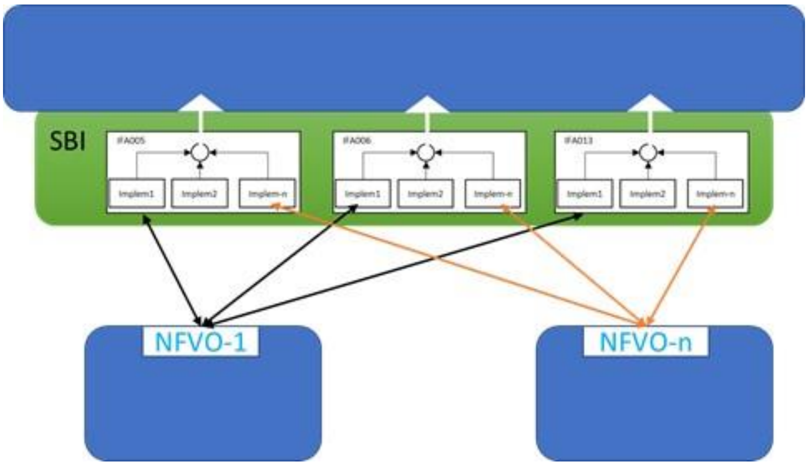

Figure 45: SBI global architecture

4.5.1.2 NS as a Service (NSaaS)

In this use case, the Multi-site NSO uses to the SBI in order to deliver NS lifecycle management across multiple sites. To achieve that, we rely on NFV-IFA 013 and NFV-IFA 030.

As in the NFVIaaS use case, SLPOC and MLPOC are architectural options that need to be considered for NSaaS. In fact, depending on the selected option the call flows change, as detailed in NFV-IFA 028. Following the example of NS resource management, the call flows depending on the architectural option are:

- MLPOC: the Multi-site NSO issues network resource management operations towards the local VIM(s).
- SLPOC: the Multi-site NSO issues network resource management operations towards the SLPOC, which in turn forwards the operations to the relevant VIMs.

As for the previous case, for NSaaS we will focus on a SBI with NFVOs implementing a SLPOC (Figure 46).

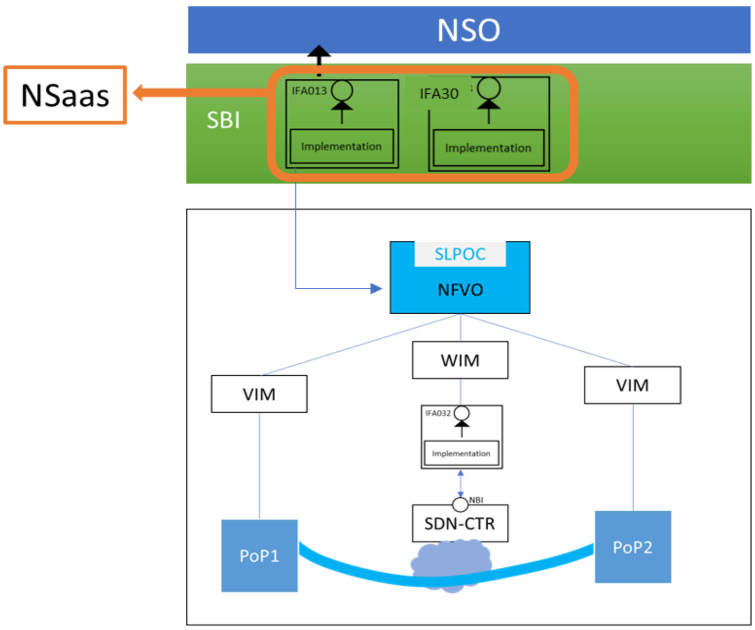

Figure 46: NSaaS SBI with NFVO SLPOC

Finally, because in 5G EVE we seek to orchestrate services across multiple sites, the Multi-site NSO can use/create nested NS at each site to offer composite NS spanning multiple local sites. In this use case, the Multi-site NSO uses the SBI for NS lifecycle management on each site, in addition to the composite NS. In the example shown in Figure 47, the Multi-site NSO uses the nested NSs exposed by each site (NS-A and NS-B) to offer a composite NS-C.

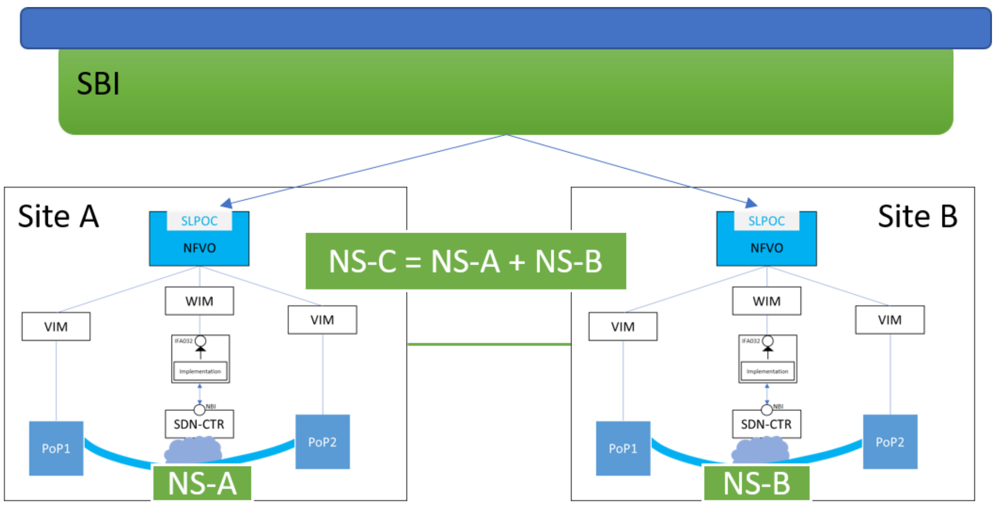

Figure 47: Composite NS using nested NSs

4.5.2 Multi-Site NSO to local Orchestrators interface

Since each site can have more than one orchestrator (see [1], Section 4.3.5), and as already outlined in section 4.2.6, in the context of the SBI we have decided to avoid the term "site" in favour of "local orchestrator", as the Adaptation Layer will actually interact with these entities. The adjective local is useful to distinguish the orchestrator operating in a site (NFVO) from the one described above and referred to as Multi-site Network Service Orchestrator (NSO).

The main features offered by the Adaptation Layer API to the I/W Framework are the following:

- Retrieving information about managed local NFVOs and their capabilities
- NSD on-boarding to local NFVOs' catalogues
- NSD management on local NFVOs' catalogues
- VNFD and PNFD retrieval from local NFVOs' catalogues
- Subscription to receive notifications about local NFVOs' catalogues updates
- Lifecycle Management of Network Service Instances (NSIs)
- Subscription to receive notifications about local NFVOs' running NSI updates

The Adaptation Layer is not meant to deal with multi-domain (multi-orchestrator) service deployments that can be potentially realized with nested NSDs as defined in ETSI NFV-IFA 014 [3] and ETSI NFV-SOL 001 [4]. Instead, and as commented in the previous section, the Interworking Framework will send requests targeting one and only one local NFVO at a time. The Multi-site NSO is expected to deal with composite NSDs by decomposing them in several connected nested NSD. The resulting single-NFVO NSDs are then sent to the Adaptation Layer for deployment on local NFVOs.

Given the previous set of high-level features, the following components are expected to be provided by the Adaptation Layer:

- A **storage** element to keep local NFVOs information and to track subscriptions to catalogue notifications
- A set of **translation** drivers to translate NSD elements[22] from the ETSI NFV-SOL 001 standard to the specific format of each local NFVO.
- A set of APIs to offer the **operations** to the I/W Framework and to route requests to the appropriate local NFVO.
- A **subscription and notification** system to propagate updates from the local NFVOs to the I/W Framework.

The reference standard for the Adaptation Layer API is ETSI NFV-IFA 013 [9] as it already defines the operations and exchanged information elements to manage NSDs, NSIs (NS lifecycle), VNFDs, and PNFDs. Furthermore, ETSI NFV-IFA 030 [12] is also useful as it defines the Or-Or (Orchestrator to Orchestrator) interface to allow the communications of two NFVOs from different domains. In fact, the interaction between the I/W Framework and the local NFVOs can be seen as a multi-domain operation.

Not all the operations specified by the standards are needed to be implemented in 5G EVE. However, some extensions will be needed to support the management of the information about local NFVOs (e.g., name, location, radio coverage, capabilities, etc.) and to fully define the subscription and notification system (partially defined by the standards).

4.5.2.1 API exposed by the Adaptation Layer to the Multi-site NSO

The following tables report the detail of the APIs of the Adaptation Layer interfacing the Multi-site NSO with a number of NFVOs. Tables are organized by topics and each one reports operations, information exchanged and reference standard (if available) about a specific set of entities managed by the API.

Table 12 provides operations for retrieving information about the identities and some high-level capabilities of the available local NFVOs, together with the relevant location field. We found no standard to support this kind of operations and information exchange.

[22] VNFD and PNFD (and related management operations) translation is not necessary as the onboarding and modifications of these elements from the I/W Framework to local NFVOs is not supported.

Table 12: Adaptation Layer exposed APIs – Local NFVO information management

Local NFVO information management		
Reference Standards	No standard available	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve list of local NFVO	This operation allows to retrieve the list of local orchestrators managed by the Adaptation Layer Input: - Output: • List of local orchestrators	No standard available
Retrieve information about local NFVO	This operation allows to retrieve all the information about a local orchestrator Input: • Local orchestrator identifier Output: • Local orchestrator information	No standard available. The request shall return local orchestrator information including: ID, name, geographic area, type of NFVO…

Table 13 shows the operations the Multi-site NSO can leverage to subscribe to the notifications about changes in the local NFVO. This is especially useful to support the synchronization of the Interworking catalogue and inventory with the status of local NFVOs. The referenced standard provides only a preliminary definition of these kind of operations.

Table 13: Adaptation Layer exposed APIs – Subscriptions and Notifications management

Subscriptions and notifications management		
Reference Standards	NFV-IFA 013	
Operations Exposed	***Information Exchanged***	***Information Model***
Create subscription to notifications related to NSD management changes in a local orchestrator	This operation allows to subscribe to notification changes on a local orchestrator catalogue related to NSDs. Input: • Local orchestrator identifier • Filter for selecting the NSD(s) and related change notifications Output: • Subscription identifier	NFV-IFA 013, extended with: • Local orchestrator identifier
Create subscription to notifications related to VNF Packages(s) changes in a local orchestrator	This operation allows to subscribe to notification changes on a local orchestrator catalogue related to VNFs. Input: • Local orchestrator identifier • Filter for selecting the VNF Package(s) and related change notifications Output: • Subscription identifier	NFV-IFA 013, extended with: • Local orchestrator identifier

<table>
<tr><td>Create subscription to notifications related to NSI management changes in a local orchestrator</td><td>This operation allows to subscribe to notification changes on a local orchestrator regarding NSIs.
Input:
• Local orchestrator identifier
• Filter for selecting the NSI(s) and related change notifications
Output:
• Subscription identifier</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
<tr><td>Terminate subscription</td><td>This operation allows to terminate the referenced subscription.
Input:
• Subscription identifier
Output:
• Success/Error result</td><td>NFV-IFA 013</td></tr>
</table>

Table 14, Table 15 and Table 16 deal with details about the APIs the Multi-site NSO can leverage to manipulate NSD, VNFD, and PNFD elements managed by local NFVO catalogue. The reference standard declares all of the following operations, but extensions to the information model are needed to allow the Multi-site NSO to select the specific local NFVO that it wants to interact with.

Table 14: Adaptation Layer exposed APIs – NSD management

<table>
<tr><th colspan="3">Network Service Descriptor Management</th></tr>
<tr><td>Reference Standards</td><td colspan="2">NFV-IFA 013</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve NSD</td><td>This operation allows to retrieve an existing NSD from a specific local orchestrator catalogue.
Input:
• NSD identifier
• Local orchestrator identifier
Output:
• NSD</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
<tr><td>Onboard NSD</td><td>This operation allows to onboard an NSD to one or more local orchestrator catalogues.
Input:
• NSD
• List of local orchestrator identifiers
Output:
• NSD identifier for each local orchestrator</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
</table>

Modify NSD	This operation allows to modify an already on-boarded NSD. Input: • NSD identifier • Local orchestrator identifier • The modified content of the NSD Output: • NSD updated	NFV-IFA 013, extended with: • Local orchestrator identifier
Delete NSD	This operation allows to delete an existing NSD from the local orchestrator catalogue. Input: • NSD identifier • Local orchestrator identifier Output: • Outcome of the operation	NFV-IFA 013, extended with: • Local orchestrator identifier
Retrieve list of NSDs	This operation allows to retrieve the list of on-boarded NSD from a specific local orchestrator catalogue Input: • Local orchestrator identifier Output: • List of NSDs	NFV-IFA 013, extended with: • Local orchestrator identifier

Table 15: Adaptation Layer exposed APIs – VNFD management

VNF Descriptor Management		
Reference Standards	NFV-IFA 013	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve VNFD	This operation allows to retrieve an existing VNFD from a local orchestrator catalogue. Input: • VNFD identifier • Local orchestrator identifier Output: • VNFD	NFV-IFA 013, extended with: • Local orchestrator identifier
Retrieve list of VNFD	This operation allows to retrieve the list of on-boarded VNFD from the local orchestrator catalogue Input: • Local orchestrator identifier Output: • List of VNFDs	NFV-IFA 013, extended with: • Local orchestrator identifier

Table 16: Adaptation Layer exposed APIs – PNFD management

<table>
<tr><th colspan="3">PNF Descriptor Management</th></tr>
<tr><td>Reference Standards</td><td colspan="2">NFV-IFA 013</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve PNFD</td><td>This operation allows to retrieve an existing PNFD from the local orchestrator catalogue.
Input:
• PNFD identifier
• Local orchestrator identifier
Output:
• PNFD in YAML/JSON format</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
<tr><td>Retrieve list of PNFD</td><td>This operation allows to retrieve the list of on-boarded PNFD from the local orchestrator catalogue
Input:
• Local orchestrator identifier
Output:
• List of PNFDs in YAML/JSON format</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
</table>

Finally, Table 17, Table 18 and Table 19 illustrate the general operations that the Multi-site NSO can leverage to manage instances running in the NFVI managed by the local NFVO.

Table 17: Adaptation Layer exposed APIs – NSI management

<table>
<tr><th colspan="3">Network Service Instance Management</th></tr>
<tr><td>Reference Standards</td><td colspan="2">NFV-IFA 013</td></tr>
<tr><td>Operations Exposed</td><td>Information Exchanged</td><td>Information Model</td></tr>
<tr><td>Retrieve NSI</td><td>This operation allows to retrieve an existing NSI from a local orchestrator.
Input:
• NSI identifier
• Local orchestrator identifier
Output:
• NSI in YAML/JSON format</td><td>NFV-IFA 013, extended with:
• Local orchestrator identifier</td></tr>
</table>

Instantiate NSI	This operation allows to instantiate an NSI in a local orchestrator. Input: • NSD identifier • NSI name • VIM Account Id (already registered) • SSH Keys • Configuration parameters • Local orchestrator identifier Output: • NSI identifier	NFV-IFA 013, extended with: • Local orchestrator identifier
Terminate NSI	This operation allows to terminate a running NSI in a local orchestrator. Input: • NSI identifier • Local orchestrator identifier Output: • Outcome of the operation	NFV-IFA 013, extended with: • Local orchestrator identifier
Scale NSI	This operation allows to perform an NSI scaling action. Input: • NSI identifier • Local orchestrator identifier • Scale group Output: • Outcome of the operation	NFV-IFA 013, extended with: • Local orchestrator identifier
Retrieve list of NSI	This operation allows to retrieve the list of NSI from a local orchestrator. Input: • Local orchestrator identifier Output: • List of NSIs in YAML/JSON format	NFV-IFA 013, extended with: • Local orchestrator identifier

Table 18: Adaptation Layer exposed APIs – VNFI management

VNF Instances Management		
Reference Standards	NFV-IFA 013	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve VNFI	This operation allows to retrieve an existing VNFI from a local orchestrator. Input: • VNFI identifier • Local orchestrator identifier Output: • VNFI YAML/JSON format	NFV-IFA 013, extended with: • Local orchestrator identifier
Retrieve list of VNFI	This operation allows to retrieve the list of VNFI from a local orchestrator. Input: • Local orchestrator identifier Output: • List of VNFIs in YAML/JSON format	NFV-IFA 013, extended with: • Local orchestrator identifier

Table 19: Adaptation Layer exposed APIs – PNFI management

PNF Instances Management		
Reference Standards	NFV-IFA 013	
Operations Exposed	***Information Exchanged***	***Information Model***
Retrieve PNFI	This operation allows to retrieve an existing PNFI from a local orchestrator. Input: • PNFI identifier • Local orchestrator identifier Output: • PNFI YAML/JSON format	NFV-IFA 013, extended with: • Local orchestrator identifier
Retrieve list of PNFI	This operation allows to retrieve the list of PNFI from a local orchestrator. Input: • Local orchestrator identifier Output: • List of PNFIs in YAML/JSON format	NFV-IFA 013, extended with: • Local orchestrator identifier

In the next two sections, the mapping of the API defined above to the more common NFVO technologies in 5G EVE (OSM and ONAP) is presented.

4.5.2.2 Multi-site NSO Adaptation Layer towards OSM

This section reports the mapping of the Adaptation Layer's API to OSM, with the goal to show reusable operations and the missing ones that need to be implemented. In that sense, all tables below relate directly to the ones in the previous section.

The OSM Northbound Interface is a RESTful API loosely following ETSI NFV-SOL 005 [10] standard. It admits both YAML and JSON formats. A description of the OSM NBI and of the related information model can be found at [35] and [36], respectively.

Table 20: Available methods in OSM for Subscriptions and Notifications management

Subscriptions and notifications management	
Adaptation Layer Operation	***OSM equivalent Operation***
Create subscription to notifications related to NSD management changes in a local orchestrator	Not present. It will be implemented in the Adaptation Layer.
Create subscription to notifications related to VNF Packages(s) changes in a local orchestrator	Not present. It will be implemented in the Adaptation Layer.
Create subscription to notifications related to NSI management changes in a local orchestrator	Not present. It will be implemented in the Adaptation Layer.
Terminate subscription	Not present. It will be implemented in the Adaptation Layer.

Table 21: Available methods in OSM for NSD management

Network Service Descriptor Management	
Adaptation Layer Operation	***OSM equivalent Operation***
Retrieve NSD	GET /nsd/v1/ns_descriptors/<nsd_id>/nsd_content
Onboard NSD	POST /nsd/v1/ns_descriptors
Modify NSD	PUT /nsd/v1/ns_descriptors/<nsd_id>/nsd_content
Delete NSD	DELETE /nsd/v1/ns_descriptors/<nsd_id>/
Retrieve list of NSDs	GET /nsd/v1/ns_descriptors

Table 22: Available methods in OSM for VNFD management

VNF Descriptor Management	
Adaptation Layer Operation	***OSM equivalent Operation***
Retrieve VNFD	GET /vnfpkgm/v1/vnf_packages/<vnfPkgId>/ package_content
Retrieve list of VNFD	GET /vnfpkgm/v1/vnf_packages

Table 23: Available methods in OSM for PNFD management

PNF Descriptor Management	
Adaptation Layer Operation	***OSM equivalent Operation***
Retrieve PNFD	GET /nsd/v1/pnf_descriptors/<pnfdId>/pnfd_content
Retrieve list of PNFD	GET /nsd/v1/pnf_descriptors

Table 24: Available methods in OSM for NSI management

Network Service Instance Management	
Adaptation Layer Operation	***OSM equivalent Operation***
Retrieve NSI	GET /nslcm/v1/ns_instances /<nsInstanceId>
Instantiate NSI	POST /nslcm/v1/ns_instances
Terminate NSI	DELETE /nslcm/v1/ns_instances /<nsInstanceId>
Scale NSI	POST /nslcm/v1/ns_instances/<nsInstanceId>/scale
Retrieve list of NSI	GET /nslcm/v1/ns_instances

Table 25: Available methods in OSM for VNFI management

VNF Instances Management	
Adaptation Layer Operation	***OSM equivalent Operation***
Retrieve VNFI	GET /nslcm/v1/vnf_instances/<vnfInstanceId>
Retrieve list of VNFI	GET /nslcm/v1/vnf_instances

Table 26: Available methods in OSM for PNFI management

PNF Instances Management	
Adaptation Layer Operation	***OSM equivalent Operation***[23]
Retrieve PNFI	GET /nslcm/v1/vnf_instances/<vnfInstanceId>
Retrieve list of PNFI	GET /nslcm/v1/vnf_instances

4.5.2.3 Multi-site NSO Adaptation Layer towards ONAP

The ONAP Northbound Interface brings a set of APIs that can be used by external systems, as BSS for example. These APIs are based on the TM Forum API specification [37]. The description of ONAP NBI can be found at [38].

Figure 48 below shows the main ONAP APIs exposed to components such as: Service Catalogue, Service Orchestrator and Service Inventory. Most of the APIs are already implemented, while some operations presented in the following tables are still under development. As of today, however, external APIs are not implemented yet; we are using internal APIs.

[23] As seen, in OSM there are no fundamental differences between a VNF and a PNF

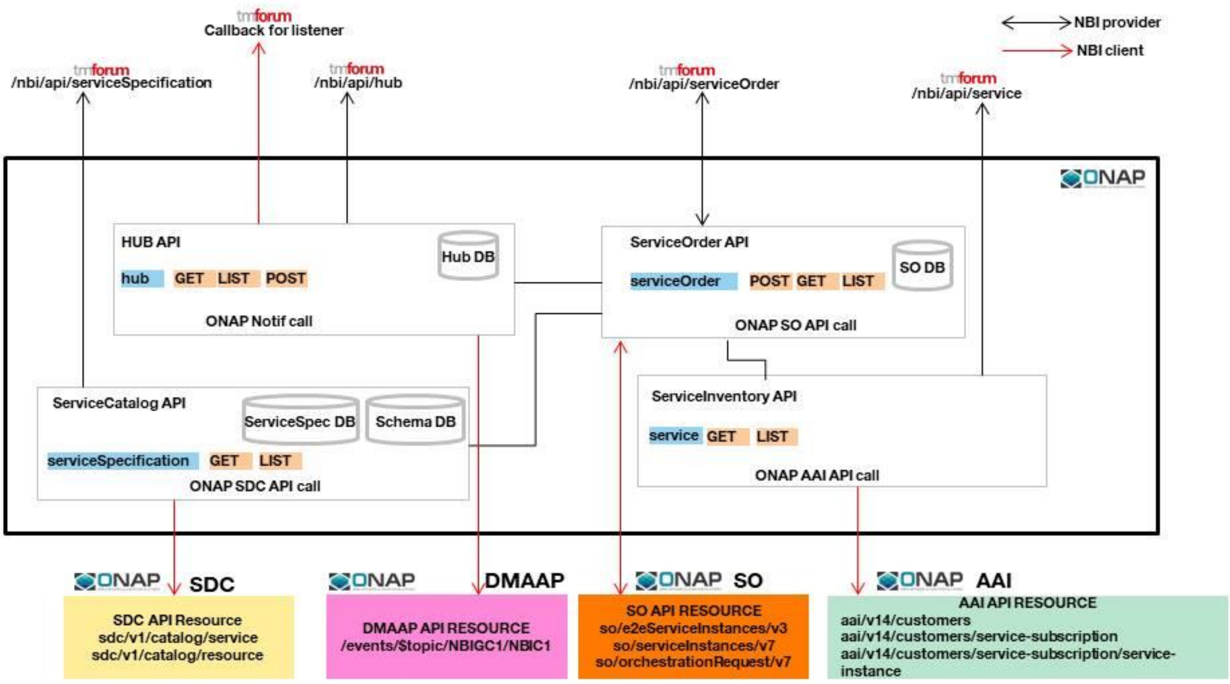

Figure 48: NSaaS SBI with NFVO SLPOC

Below APIs are implemented in the SDC module, especially in the catalogue.

Table 27: Available methods in ONAP for Subscriptions and Notifications management

Subscriptions and notifications management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Create subscription to notifications related to NSD management changes in a local orchestrator	Implemented in ONAP Dublin release; to be validated.
Create subscription to notifications related to VNF Packages(s) changes in a local orchestrator	Implemented in ONAP Dublin release; to be validated.
Create subscription to notifications related to NSI management changes in a local orchestrator	Implemented in ONAP Dublin release; to be validated.
Terminate subscription	Implemented in ONAP Dublin release; to be validated.

Table 28: Available methods in ONAP for NSD management

Network Service Descriptor Management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Retrieve NSD	GET /sdc/v1/catalog/service/<id>
Onboard NSD	POST /sdc/v1/catalog/service/<ns_descriptors>
Modify NSD	PUT /sdc/v1/catalog/service/<id>/<ns_descriptors>
Delete NSD	DELETE /sdc/v1/catalog/service/<id>
Retrieve NSD List	GET /sdc/v1/catalog/service/

Table 29: Available methods in ONAP for VNFD management

VNF Descriptor Management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Retrieve VNFD	GET /sdc/v1/catalog/resource/<vnf_id>
Retrieve VNFD List	GET /sdc/v1/catalog/resource/

Table 30: Available methods in ONAP for PNFD management

PNF Descriptor Management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Retrieve PNFD	GET /sdc/v1/catalog/resource/<vnf_id>
Retrieve PNFD List	GET /sdc/v1/catalog/resource/

Table 31: Available methods in ONAP for NSI management

Network Service Instance Management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Retrieve NSI	GET /nbi/api/v4/service/<nsInstanceId>
Instantiate NSI	POST /nbi/api/v4/service/<ns_instances>
Terminate NSI	DELETE /nbi/api/v4/service/<nsInstanceId>
Scale NSI	Op. already implemented, but not yet available in the external API.
Retrieve list of NSI	GET /nbi/api/v4/service/<ns_instances>

Table 32: Available methods in ONAP for VNFI management

VNF Instances Management	
Adaptation Layer Operation	***ONAP equivalent Operation***
Retrieve VNFI	GET /network/generic-vnfs/generic-vnf/{vnf-id}
Retrieve list of VNFI	GET /network/generic-vnfs/generic-vnf/

Table 33: Available methods in ONAP for PNFI management

PNF Instances Management	
Adaptation Layer Operation	***ONAP equivalent Operation***[24]
Retrieve PNFI	GET /network/generic-vnfs/generic-vnf/{vnf-id}
Retrieve list of PNFI	GET /network/generic-vnfs/generic-vnf/{vnf-id}

[24] Again, in ONAP there are no fundamental differences between a VNF and a PNF

4.5.3 Multi-site NSO to local SDN Controllers interface

The SDN controller is used to request for connectivity services through the Nf-Vi reference point. Typically, an SDN controller interfaces the VIM or WIM with the NFVI. However, in some cases the SDN controller is part of the VIM itself, controlling directly the virtual entities including switches and network functions inside the NFVI-PoP. Figure 49 below depicts how VIMs are interacting with the SDN controllers and the infrastructure via the Nf-Vi reference points. It also shows how to establish an end-to-end connectivity across a WAN under the control of the WIM.

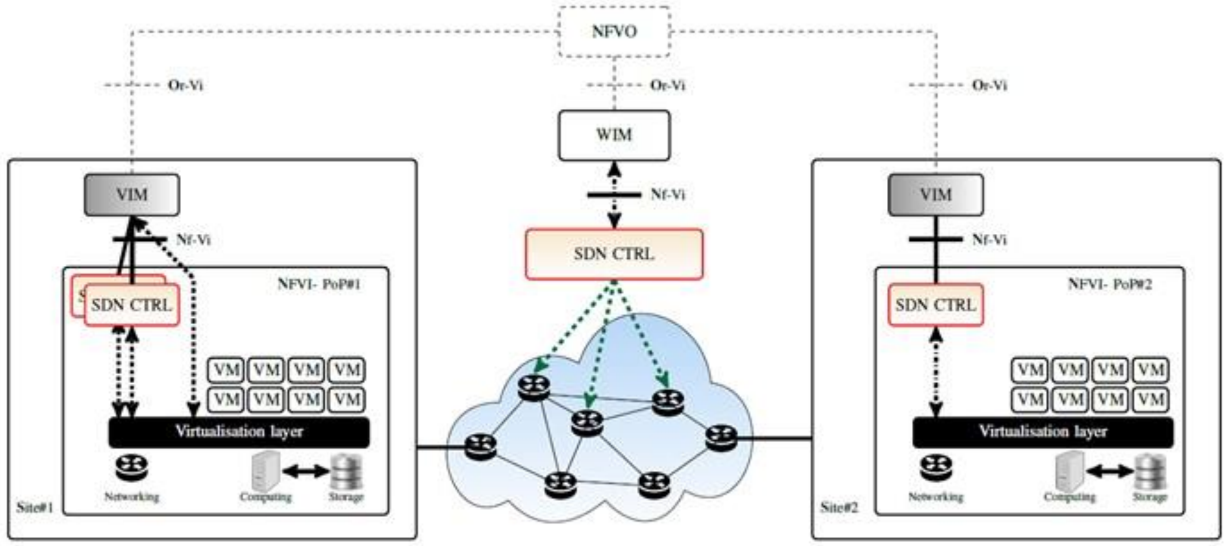

Figure 49: Local SDN Controllers location and interfaces

In the 5G-EVE project, there are four participating sites: Spanish, French, Italian and Greek. Some of these include an underlying SDN controller interfacing with the VIM of the related NFVI-PoP. Therefore, it is the responsibility of the site manager to define and manage the interfaces between the VIM and the NFVI. Some others, however, do not use an underlying SDN controller; its functionality should then be integrated within the VIM.

In the SBI, the design will be limited to the interconnection between the sites, although the considerations are equally valid for interconnections between VIMs at the same site. For that purpose, a WAN Infrastructure manager (WIM) that is responsible for establishing the connectivity between endpoints in different NFVI-PoPs is needed. The proposal in 5G EVE for the interface between WIM and SDN Controller is to follow ETSI NFV-IFA 032 [39], which defines the interfaces and Information Model specification for Multi-site Connectivity Services.

4.5.4 Runtime Configurator SBI

When Verticals deploy and instantiate experiments on top of the 5G EVE infrastructure, it has already been described how to include the Day0/1 configurations, based on the Vertical's preferences expressed in the Portal. However, during experimentation, it will also be possible to modify certain configuration parameters without executing a decommissioning and subsequent instantiation of the service. This is the key functionality for the Runtime Configurator.

The potential capabilities required of this component would be extremely large if we wanted to globalize its specification: the number of components (VNFs, PNFs), the type of Vertical services, the differences among

Verticals' use cases, etc., would make it impossible to model all the required operations. Instead, we will start building in the project on a per Use Case basis, trying to reuse as much as possible the features among them.

A very simple example of the complexity of this component is the traffic generation tools. If we wanted to start a traffic flow using open source tools like Iperf, the Adaptation Layer would have to connect (e.g., via SSH) to the compute(s)/VM(s) hosting the Iperf client&server applications and send the CMD commands that specify the traffic flows and start the traffic. The same operation towards a commercial traffic generator would derive in a totally different set of operations, as it would have to use the available APIs of the tool to load for example the desired topology and traffic patterns, and to start the uploaded scenario. Yet another simple example is the configuration of a VLAN in a switch, to provide connectivity in the data plane: depending on the switch manufacturer, the sequence of required commands will vary, so scripts would not be directly reusable.

Having said this, there are mechanisms to try to harmonize as much as possible the behaviour of the deployed VNFs/PNFs. The most common one for network connectivity, which use will be promoted in 5G EVE when possible, is the modelling of the required configurations via YANG. There are many defined YANG models, depending on the organization who is making the definition. There are even proprietary YANG models. In 5G EVE we will rely on the supported YANG models by each network component, trying to unify them in one of the two more common specifications: IETF or Openconfig.

Other alternatives like SNMP will normally be avoided for the network configuration, as these are protocols which are currently not being so much used for provisioning (in production networks today they are normally required only for monitoring).

Of course, not all devices support YANG models, or there are not YANG models for all of the required functionalities. This demands from the project a very flexible solution for the SBI of the Runtime Configurator. Whenever possible, this solution will be Ansible, although scripting (possibly using templates) could also be used.

The definition of the interface itself does not vary from that of the NBI and is adapted in Table 34 below.

Table 34: Runtime Configurator Interworking API

Runtime Configurator API		
Update of current instances		
Reference Standards	Openconfig/IETF Yang models, Netconf, Ansible	
Operations Exposed	***Information Exchanged***	***Information Model***
Enable a functionality or configuration	• VNF/PNF management IP addresses • Functionality/configuration to enable	YANG model or configuration template depending on the required configuration
Disable a functionality or configuration	• VNF/PNF management IP addresses • Functionality/configuration to be disabled	YANG model or configuration template depending on the required configuration
Update/Upgrade of a functionality or configuration	• VNF/PNF management IP addresses • Functionality/configuration to be updated/upgraded.	YANG model or configuration template depending on the required configuration

4.5.5 Data Collection Manager SBI

As it was mentioned in the Data Collection Manager NBI interface, the operations for the SBI interface are focused on handling the interaction between the publishers (placed in the site facilities) and the subscribers from upper layers. Particularly, the only operation handled by the SBI interface is the publishing of monitored data, as the possible topics are already defined in the overall architecture and the subscribers are already using them. For that reason, when a use case is deployed, the data shippers (related to these certain topics) are already collecting data and publishing them into the queue, and when the publish/subscribe system detects that there are new data, it is automatically delivered to the subscribers. Moreover, if there were new subscribers to these topics in the system (as a result of the "topic subscription" operation in the NBI), data is automatically delivered to them too.

Table 35: Data Collection Manager SBI

Experiment results and monitoring APIs		
Collection of monitored data		
Description	Data collection from the different components deployed in the site facilities for a given experiment.	
Reference Standards	Interface based on Publish/Subscribe model. Regarding network monitoring: RMON (RFC2819, extensions and updates), Syslog (RFC5424, extensions and updates), SNMP (RFC1157, extensions and updates). Regarding testing: NGMN Alliance "Definition of the testing framework for the NGMN 5G pre-commercial networks trials"	
Operations Exposed	***Information Exchanged***	***Information Model***
Publish data	• Topic associated to the experiment/component. • Monitored data.	Operation of publishing monitored data to the publish/subscribe queue. Data is provided in a unified format (JSON chain, typically).

5 Interworking Planes

This section addresses the requirements from WP3 to the site facilities interconnection solution, to enable the following goals:

- Connect the site facilities with the Interworking Framework in a secured way (the I/W Framework will be hosted by a single site)
- Comply with connectivity requirements identified in D3.1 [1]

All the aspects related to this interconnection will be covered in the following sub-sections, together with an implementation proposal which, at the time of writing of this Deliverable, is still under evaluation by WP2:

- ✓ Potential Technologies
- ✓ Topology
- ✓ Main site identification (i.e., the one hosting the I/W Framework)
- ✓ Addressing plan
- ✓ Previous examples

Apart from the above, there exist some additional requirements complementing the connectivity solution, like the need for experimenters to access their deployed VNFs for management, the capability of the sites to provide Internet access to VNFs (for example, for patch downloading), or the need for a global Monitoring and Support scheme that ensures correct operation during experimentation. All these requirements, derived from previous experiences, will be formalized in section 6.

Finally, and as stated above, the final implementation has not yet been decided in 5G EVE; it shall be discussed and implemented in the coming months, however, to cope with the project planning which fixes the first delivery of the I/W layer software by Month 16.

5.1 Requirements

The following table summarizes the requirements that where identified in D3.1 regarding the sites interconnection, with one correction[25]. Please, note that a definition of each of the interconnection planes was already given in the Terminology section of this document (section 1.1.2).

Table 36: Interconnection requirements from D3.1

Requirement	Target I/W Capability[26]	Description
Orchestration Plane interconnection	Connectivity	Connectivity with the 5G EVE Interworking Framework shall be established with a minimum availability of 99.9%
Control Plane interconnection	Connectivity	Connectivity with other 5G EVE sites shall be established with an availability of at least 99.9% and 20 Mbps of guaranteed bandwidth

[25] Table in D3.1 read “99.9% reliability”. The correct requirement is now included.

[26] As per definition in section 3.1

Data Plane interconnection	Connectivity	Connectivity with other 5G EVE sites with an availability of 99,9%, 200 Mbps of guaranteed bandwidth (at least between a pair of sites) and the following maximum latency, per pair of sites[27]: • Italy(Turin)-Greece(Athens) interconnection: 160 ms • Italy(Turin)-Spain(Madrid) interconnection: 110 ms • Italy(Turin)-France interconnection: 60 ms • Greece(Athens)-Spain(Madrid) interconnection: 240 ms • Greece(Athens)-France interconnection: 210 ms • Spain(Madrid)-France interconnection: 110 ms

5.2 Technologies

There are several technologies that can be used to interconnect the site facilities among them, and with the Interworking Framework. The choice of a given technology will be mainly based on its capacity to meet the following criteria:

- Cope with the performance requirements summarized above.
- Provide an appropriate level of security for the data flowing through the connection, and with regards to the exposition of the site facility.
- Consider the capacity scaling in case new site facilities are added during the project's lifetime, based for example on the ICT-19 proposals that will be instantiated in the 5G EVE end to end site facility.

The following alternatives could be considered to connect the different interconnection Planes:

- **Public Internet connection**. This solution is the easiest to setup and does not present scaling issues based on the number of supported tunnels. Since it does not encapsulate the traffic, it does not affect the performance either. However, it has two main withdraws for the site facility interconnection: the security and the number of public IP addresses required to access each platform's components.
- **Internet connection with tunnelling (e.g., VxLAN, GRE)**. This solution allows to preserve most of the platform's privacy but it does not add security for the data exchanged in the tunnels. It is still required to setup specific equipment/configurations to build the tunnels. However, these tunnelling mechanisms should not generate a big impact on the overall performance. Scale complexity will depend on the chosen tunnel technology.
- **IPSec Site2Site**. This a variant of the tunnel interconnection allowing to greatly improve the security of the exchanged traffic. However it also increments the setup complexity and may impact the overall performance, depending on the equipment used to build the tunnels (IPSec is a very demanding feature). While the interconnection of the four site facilities could be achieved easily, scaling-out the solution may become challenging quickly, especially in full-mesh topologies.
- **Operator VPN**. Offers the advantages of a dedicated private network across the site facilities with minor setup on the premises. This solution allows a better scaling than other tunnelling solutions and provides another security level based on the capacity to define the path(s) used by the tunnel. The main withdraws are the price and the capacity to build the VPN across different operator domains.
    - One identified solution is the multi domain L2/L3VPN service provided by GÉANT[28], but this is only available for National Research & Education Networks.

[27] For the French facility, which is composed of four sites, the value refers to the furthest location from Turin, Athens and Madrid.

[28] https://www.geant.org/

A special attention is required for inter-site facility routing when choosing any of the tunnelling solutions. In fact, the private networks available in IPv4 are often extensively used by companies to implement the different parts of their private networks. In order to reduce the number of IP addresses that need to be routable in 5G EVE (which reduces the complexity of identifying a common IP addressing plan), **it is recommended that each site facility implements a VRF dedicated to this interconnection**.

5.3 Topologies

For any of the tunnelling technologies discussed in the previous section, mainly two basic topologies should be considered:

- Star: a central site acts a hub for the rest of the site facilities.
- Mesh: the site facilities are connected to each other either in full mesh or partial mesh.

The topology used may be the same or different across the different connectivity Planes. The choice will be driven by the performances that are obtained out of the implemented interconnection solution, to which the choice about the I/W Framework hosting site also affects. Another point to keep in mind is the possibility to include new site facilities in the future.

5.4 Interconnection example

In this section, an interconnection example is provided based on what has been already performed to connect all the testbed nodes belonging to the French site facility (shown in Figure 50):

- Orange (Paris Châtillon)
- Nokia France (Paris Saclay)
- B-com (Rennes)
- Eurecom (Sophia Antipolis)

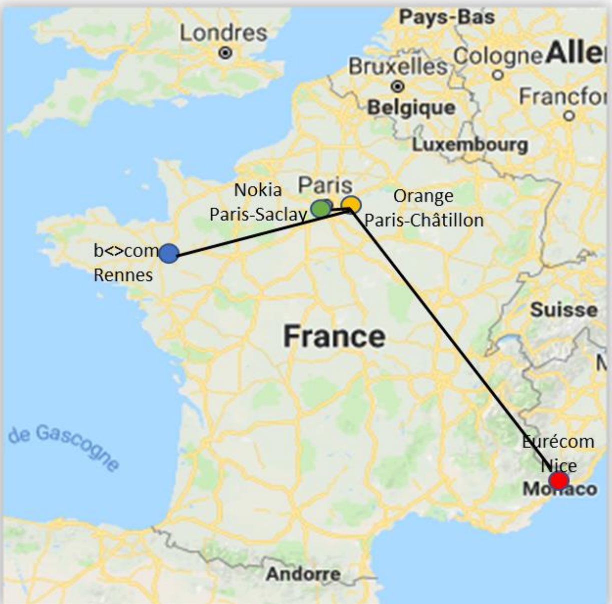

Figure 50: French site facility testbed nodes

All the sites are connected by Site2Site IPSec VPN using a star topology, with Orange Chatillon as the main site facility, as depicted in the following schema:

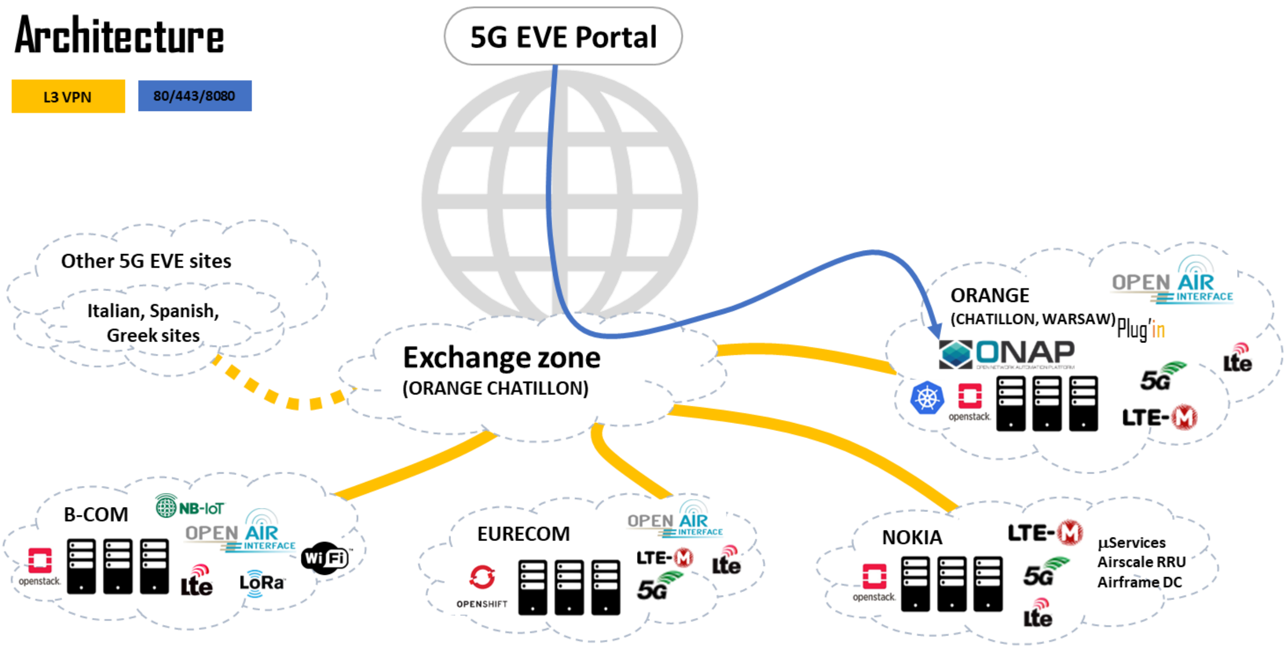

Figure 51: French site facility VPN interconnection

As shown in Figure 51 each testbed node is connected to Orange main site by the means of an IPSec VPN (yellow traces), using their own existing VPN concentrators.

No choice has already been made for other items like the Verticals' remote access to their VNFs, or the VNF Internet access, as this is expected to be done in accordance with the scheme globally adopted by the Project.

The monitoring of the interconnection status, however, will be done in two phases:

- First, a basic "ping" test will allow to check the tunnels status ("alive", "not alive".
- Later, it is under study the use of "Iperf" [30] based tools, like Trafic [41] from MAMI[29] project, to collect performance metrics.

5.5 Interconnection proposal

As conclusion to this chapter, an interconnection proposal is given based on all the information already provided. This proposal is currently being discussed, mainly in the scope of WP2.

The following assumptions are made based on experiences from previous projects and the setup of the French site facility internal interconnection:

- The site facilities interconnection shall be performed by the means of IPSec VPN tunnels. In a first run, a star topology with a main site facility should be sufficient to build up the site facility interconnection.
- A routable private addressing plan for the interconnection shall be identified, agreed by all sites.
- The main site facility shall be designed to host centralized services (e.g., monitoring of the orchestration components).
- The Portal and the Interworking Framework shall be collocated in one of the site facilities, with preference of the main one. This is not a strict requirement but allows the Interworking Framework to take advantage of the deployed star topology, and the Portal to interact locally with the Framework.

[29] https://www.mami-project.eu

- The verticals' remote access shall be performed in a centralized way through the main site facility in order to provide a uniform access to all the experiments regardless of the location of their VNFs.
- The Internet access for the VNFs shall be provided locally by each site facility to simplify the routing management. However, an agreement should be done to provide the same access policies on all the site facilities.
- Experiments are supposed to happen across site facilities, so direct access to VNFs from Internet shall be forbidden by default, to avoid security risks and complex access setups.

Based on the geographical locations of the different site facilities:

- French: Paris
- Greek: Athens
- Italian: Turin
- Spanish: Madrid

the Italian site facility appears in a centralized position, as depicted in Figure 52:

Figure 52: Sites facilities interconnection phase 1

Based on this placement, it would be interesting to explore the possibility to set up the main site at Turin.

Once this setup is up and running, it will be important to have the right monitoring tools to collect the interconnection metrics which will allow to check whether the setup copes with the requirements identified in D3.1, and summarized in section 5.1.

In case the performance requirements are not met with the above solution, an evolution on the topology could be envisioned (normally for the data plane, as the required paths are not optimized in a star, and the required performance is higher than that of the control plane). A full or a partial mesh IPSec VPN can be implemented to avoid potential bottlenecks and to reduce the data plane latency, as depicted in Figure 53:

Figure 53: Site facilities interconnection phase 2

This setup, however, will increase the exploitation overhead and should be motivated.

6 5G EVE required facility Services and Functionalities

This section describes, on the one hand, global services valuable for Verticals with regards to the experiments deployment and execution and, on the other, additional functional considerations for the different sites.

The mentioned services, for example a ticketing tool, could actually be handled on a per-site basis. However, due to their importance and project-wide scope, it has been decided to define them globally, associated with the Interworking Reference Model. This way, the derived recommendations will lead to a unique solution to be adopted by the whole project.

On the contrary, services with a thinner scope, like DNS, DHCP or NTP, will not be part of this specification. There is no technical reason why they should be provided globally instead of locally. Furthermore, 5G EVE sites do not only host 5G EVE experiments, but many others which should not depend on 5G EVE provided tools for this type of services.

Regarding the functionalities, these are not features to be offered to Verticals. Instead, and based on previous experiences, they provide support to global considerations like security, access control and authentication, etc.

6.1 Infrastructure Hardening

When working with shared networks, security considerations are critical to protect the key infrastructure. As already stated, 5G EVE is deploying its own infrastructure (e.g., the I/W Framework or the Web Portal) interacting with infrastructure from the different sites, which will be normally also shared.

In this scenario, even if the orchestration or control planes are implemented over isolated networks, it is important to define security rules that protect each component on its own. This is typically called hardening, so each node reduces the number of vulnerabilities by limiting the accepted communications to the minimum required.

In particular, we will consider that all the components of the I/W Framework are deployed in the same site (as proposed in section 5.5), and share the same LAN. They may even share the same local LAN with the local orchestration modules, in which case a firewall would be very beneficial. Instead, in the more general case the Web Portal (WP4) and the Experimentation Engine modules (WP5) may be deployed at other sites, which might, or might not, coincide with the rest of the 5G EVE facility sites, as depicted in Figure 54.

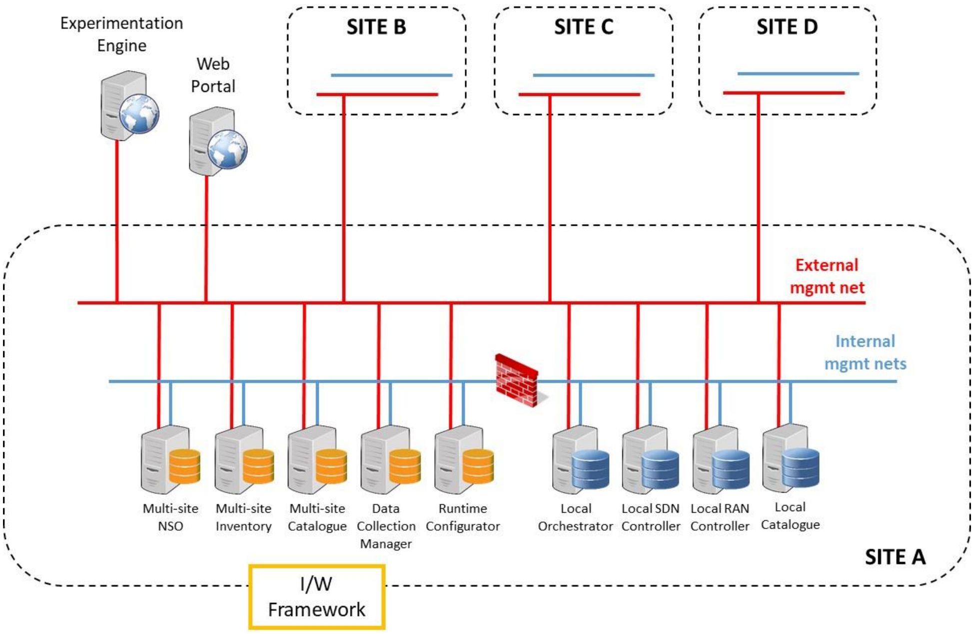

Figure 54: Example of key infrastructure deployment

In scenarios of this type it is easy to imagine the type of attacks that could be possible from outside. First, it is mandatory that the external connectivity solution is secure itself, as commented in section 5.2. VPN encryption methods can sort this requirement out, but there are other vulnerabilities apart from the inter-site connectivity.

A very critical access point to the whole infrastructure is, for example, the Web Portal, since it has an interface which is exposed to the Internet. If someone was able to compromise the integrity of the Web Portal, and no additional security measures were implemented, or they were implemented only on a few of the available modules, the whole infrastructure would then be compromised. Somebody could, for example, jump from the Web Portal to the Data Collection Manager via the external network, and from here to the Local Catalogue of Site A, which could be shared with several different projects, which might get all of their already validated VNFs corrupted.

This type of attacks cannot be controlled by hardening only the components which are shared or externally exposed; all components have to implement hardening measures themselves to ensure the maximum resistance to attacks.

In 5G EVE, it is not yet possible to define the exact rules that would harden each and every deployed function. These rules also depend on how the different APIs are implemented, since different implementations employ, for example, different TCP/UDP ports. Therefore, these rules will be fully specified during the development and (especially) deployment phases. However, a few considerations for the I/W Framework include:

- **Multi-site NSO**: the main incoming external flows in the NBI will be originated at the Web Portal, while in the SBI it will have to interface the local Orchestrators and, depending on the use case, the local SDN Controllers. Internally, it will have to interact mainly with the Multi-site Inventory and the Multi-site Catalogue, and permit local management to operators.
- **Data Collection Manager and Runtime Configurator**: their main interactions with external modules in the NBI are with the WP5 components dealing with the KPI collection and the execution of experiments. In the SBI, the list could be long and complex to define: active and passive probes, VNFs, PNFs… In the end, any device capable of providing monitoring information should be able to contact the Data Collection Manager, and any configurable device should permit the interaction with the Runtime Configurator. On the other hand, the main objective of the internal interfaces should only be the local management by the operators.
- **Multi-site Inventory and Multi-site Catalogue**: in the NBI it will be the Portal who will communicate with these two modules, while in the SBI the Multi-site Inventory may have to interact with the sites equipment, and the Multi-site Catalogue with the local Catalogues. Internally, as commented above, the main contact point will be the Multi-site NSO (and, again, the local management).

Finally, it is very important when defining hardening rules to consider where the communications are originated. To avoid complexity, allowing the "established" connections (i.e., those originated in the node itself) results in not needing to include the target modules (with their target ports) in the list of permitted incoming flows. This "incoming flows" list is therefore limited to the external nodes that need to establish incoming communications over specific ports.

6.2 Authentication in the Orchestration Plane

As described in Section 4, the Interworking Framework includes various components that need to communicate with each other and interfaces (e.g. Adaptation Layer) to communicate with external components. Furthermore, it acts as a bridge between the 5G EVE Portal and the local sites resources. Beyond the hardening measures contemplated in the previous section, authentication tools and mechanisms are required to ensure that only trusted users or trusted software components can access the resources managed by the Interworking Framework. There are two options to realize authentication:

1. Each component in the Interworking Framework implements its own authentication mechanism. The advantage of this solution is that each component is fully independent in its operation. Furthermore, different levels of authentication can be implemented e.g., we expect the NBI to have a strong

authentication mechanism as it is an interface to the external world, while communications between internal components can have simpler authentication methods. The disadvantage is a heterogeneous and scattered environment and the problem for some clients to store a large set of different credentials.

2. Authentication is provided by a single, dedicated service for all the components of the interworking framework. The advantage of this solution is having a homogenous environment in terms of interactions. This would ease the implementation of clients and would leverage each component from the burden of implementing its own solution. The disadvantage is that the operation of the entire Interworking Framework is dependent on the authentication service, which becomes a single point of failure (SPOF).

Because implementing custom authentication mechanisms can be hard and prone to the introduction of vulnerabilities, we prefer to adopt the second option. Authentication services with consolidated procedures and a successful history in real world deployments exist and can be easily applied to our use case. We can cite for example Openstack Keystone [42] that provides an identity service for client authentication. The service generates authentication tokens that can be used in the standard X-Auth-Token header of REST API requests. Keystone also implements a role-based access control to define policies for fine grained resource access control. Another available tool is Apache Shiro [43], a Java security framework that offers developers an intuitive yet comprehensive solution to authentication, authorization, cryptography, and session management. This component has been used successfully in the previously mentioned 5GinFIRE project (Section 2.2.2). Apache Shiro can be used to secure simple applications or complex, clustered environments with REST components interacting with each other.

The SPOF disadvantage can be easily overcome by providing high availability for the authentication service. This can be achieved through redundancy possibly on sites located in different geographical areas; since 5G EVE is a project involving different sites across Europe, we believe this is a feasible solution.

6.3 Verticals' remote access to deployed VNFs

Verticals deploying experiments on top of 5G EVE infrastructure should be capable of accessing their VNFs during the experimentation life cycle. Providing such access requires the definition of an access strategy and the identification of a technical solution, compliant with all sites' security policies, which should be widely accessible for Verticals to use.

There are two alternatives when it comes to the access strategy:

- Each site facility provides access to their local resources which allows each site owner to implement its own security policies and use their own access solution. On the other hand, the principal withdraw is that this may be messy from a Vertical point of view, considering that it will require to know the location of a given VNF to apply the access procedure of the given site.
- A centralized access is implemented on one site facility, providing a single entry-point for all experiments. This solution requires an agreement on the remote access policies across the site facilities, though.

The potential solutions could be the utilization of a Dialup VPN like OpenVPN [44] or a Bastion server that allows protocols like SSH and SSL to establish tunnels through a secured equipment.

In general terms, the recommendation will be to implement the lighter possible solution for experimenters, as it is not always easy to install additional components on personal computers, nor is permitted the use of specific protocols (e.g., IPSec) depending on the corporate security policies.

6.4 Internet access for deployed VNFs

Functions deployed by experiments may eventually require accessing Internet, or to be accessed from Internet. Both cases need to be addressed but the proposed solution may differ. While allowing the functions to have "limited" or "controlled" Internet access may be quite straightforward and may not be a big security issue, allowing Internet access to the functions will be more challenging both in terms of security concerns, and in terms of the configuration required to expose the functions.

In any case, the Internet access may be:

- Centralized: managed on one of the site facilities which requires in turn to agree on a common access policy.
- Distributed: each site facility manages the Internet access for their hosted functions and provides the means to access them. It also requires some agreement on the policies specification, as potential attackers could jump from one site to another via any of the interconnected Planes.
- Forbidden: it may be established that a given type of access is forbidden, which will normally be the default policy for incoming connections.

6.5 Security

It has already been mentioned the special attention that must be placed on the security aspects at all levels. Hardening and authentication have already been tackled in previous sections. Hereunder are some hints on how security can be applied to each of the connectivity Planes based on the current connectivity proposal and taking into account the previous requirements about remote access.

6.5.1 Orchestration Plane

The Orchestration Plane is used exclusively by 5G EVE internal components, like the Interworking Framework and the Orchestrators of each site facility. This relaxes the security constraints but it must not be forgotten that site components may be reused in many activities beyond 5G EVE. Therefore, security solutions like firewalling should be considered on each site facility, ensuring that only the correct components interact with each other using the expected protocols and ports.

A very important consideration, if we assume that one of the four sites will host the I/W Framework, is the internal interaction between the I/W Framework itself and the local Orchestrator, which should not be relaxed in terms of security based on "confidence" relationships (e.g., all components are internal).

6.5.2 Management Plane

The Management Plane does not expand across multiple sites, so it can be locally implemented by each site at its own will. However, security concerns apply the same to this Plane as to any other, as it may also be interconnecting shared devices (e.g., maintenance Portals, SNMP servers, etc.).

Policies in the Management Plane have a lot to do with the implemented solution, i.e., whether the Management Plane is in-band or out-band with regards to the Data Plane. If it is out of band, then similar considerations apply as those of the Orchestration Plane: all components in the Management Plane should implement hardening rules, and security policies should be applied on each site facility, separating security zones via firewalling.

If the Management Plane is implemented in-band, that is, using the same physical infrastructure as the Data Plane, then security rules i) are more difficult to establish, as the default action in the Data Plane must be "permit", and ii) must be applied to more interfaces, as the Data Plane expands through all the interconnected interfaces. Normally, some "isolation" mechanism is implemented when the Management Plane is in-band, for example, using a separate VLAN; still, in-band management is not normally used unless in environments which are physically difficult to reach to. In that sense, in-band management is highly discouraged in 5G EVE.

6.5.3 Control Plane

What was only a possibility for the Management Plane (to be carried in-band with the Data Plane), is normally a must for the Control Plane: network nodes exchange control information like that of routing protocols through the same physical connections as the Data Plane. This implies that specifications are by default complex. Fortunately, this Control Plane will eventually benefit from the "slicing" concept, reducing the impact of potential crashes.

Experience will drive the implementation of security measures in this Plane. To name just one interesting consideration, limiting the permitted bandwidth on a per protocol basis is a good measure to avoid that control processes crash due to DoS attacks.

6.5.4 Data Plane

This should be by far the most complex setup, where it might be necessary to trade-off between the security enforcement and the capacity to drive the experimentations. Hereunder are some tips applicable to the current proposal:

- Internet access for VNFs should be limited and monitored, and special requirements should be submitted for approval. Basic connectivity could include ICMP, NTP, DNS, Web Browsing and SSL. This would allow to avoid providing customs per site facility services (NTP, DNS), to check connectivity (ICMP), and to allow software installation (Web browsing).
- Policies must allow that components instantiated on different site facilities and participating of a given experiment exchange data. The following considerations should be taken into account:
    - Only traffic from the same experiment is allowed between site facilities.
    - Only traffic identified for the experiment is allowed within the policies.
    - The access policies are set dynamically at experiment instantiation.
- VNFs will be allowed to be accessed by the experimenter using the remote access method[30]. In this way, it should be considered that:
    - The experimenter will log into the remote access solution and should be able from there to connect to any of its experimentation VNFs.
    - Only methods like SSH and SSL should be allowed for this connection.
    - Isolation must be provided, in the sense that only the experimenter can access the VNFs of its experiment, and also in the sense that it cannot access any other component.

6.6 Monitoring and Support

In order to check if the interconnection reliability and performance meets the requirements it is needed to collect a certain amount of metrics:

- Link status (reliability): up, down
- Link performance: throughput, latency, jitter

With regards to those metrics it should be possible both to check their status at a specific point in time, and to graph them through a period of time. It is a possibility to be discussed if this information should be available through the project's Portal to experimenters; project partners, on the other hand, should have access to it.

The measures collected should be able to help identifying anomalous behaviours linked to the interconnection during the experimentation phases. Moreover, certain measures should trigger support actions, like:

- Link down detection
- High variations in throughput, latency, jitter

The way to trigger the support is yet to be defined but it could be done by the automatic creation of a helpdesk ticket addressed to right operations team.

[30] Such remote access method may not be implemented in the Data Plane. If so, those considerations apply to the selected Plane only.

6.7 Ticketing

One of the key characteristics of 5G EVE is the degree of automation that it will offer to the Verticals willing to execute their tests on top of our infrastructure. This means that (in most cases) they will be able to be very independent with regards to their activities, not needing the continuous help or monitoring of 5G EVE operational personnel. While this is a very positive feature, problems may appear and it may not be obvious for 5G EVE to be aware of them.

Furthermore, there are two important ideas regarding automation that need to be understood:

- Automation is a growing process which does not happen from zero to one hundred instantaneously. 5G EVE’s offered services and functionalities will be automated progressively.
- Full automation is not possible. Normally, features appear and are used even before they are fully automated, so continuous moves towards full automation usually coincide with the appearance of other new features that are typically started to be consumed manually.

For all these reasons, Verticals will need a ticketing tool that permits them to report bugs or malfunctioning, and to request required manual operations to be executed. Making this a global tool will help the project and the operational teams in particular, as all information will be available at the same place. Furthermore, the automated ticketing mentioned in the previous section could also leverage of the same tool.

With regards to WP3, in the final stages of the project WP3 will not be the main “target” of tickets: Verticals will mainly interact with the Web Portal, all the deployment and execution of experiments will be done automatically, and even if failures occurred, it would be impossible for Verticals to really understand which 5G EVE component failed. However, it is very likely that in the initial phases Verticals need to request for specific actions happening at some of the I/W Framework modules. This, only, already justifies the requirement from WP3 to develop a ticketing tool within 5G EVE.

7 Required Capabilities from sites

The final section in this document addresses the capabilities that are required from sites participating in the 5G EVE infrastructure. Such capabilities will allow the offering of a multi-site platform for Verticals' testing, capable of inter-site experimentation. This way, this section effectively concludes the Interworking Framework specification, and it serves as a direct input to drive discussions with WP2.

7.1 Required capabilities to host the Interworking Framework

As already stated, the Interworking Framework is a unique component in the 5G EVE Project, part of the Orchestration Plane. In section 5 we have discussed the proposal from WP3 for connectivity on this Plane, summarized in two key ideas:

- The star-like topology as basic distribution, with the I/W Framework hosted at the centre of the star
- The proposal for Turin site to be the central site, due to delay considerations

Beyond the obvious need about having available the required capacity to install the I/W Framework components, the following are additional requirements associated to the central site:

1. The central site must be able to act as hub for Orchestration Plane communications of all the sites in 5G EVE. Depending on the selected technologies (see section 5), this could have a performance impact as implementing encrypted tunnelling is a very resource-consuming functionality.
2. The central site should be capable of running the desired Interworking Connectivity solution monitoring, i.e., it should be able to track the operational status of all the components in the Orchestration Layer. Again, depending on the periodicity of the defined measurements, this could be a very resource-consuming feature.
3. Although desirable, the Experimentation Portal does not necessarily need to be installed at the same site as the I/W Framework. Therefore, the central site must also provide the correct connectivity towards the Portal (normally, extending the connectivity solution in the Orchestration Plane to the required "new site").
4. The Management Plane in the central site must be remotely accessible for the I/W Framework developers and for the support team, which is not necessarily fully composed of local personnel. This is a requirement for troubleshooting and resolution of bugs, but also because updates of the I/W Framework will normally be installed by the development teams.
5. Physical impairments on any of the components of the I/W Framework (e.g., a broken cable) will have to be repaired by the local support team. The rapid reaction to such impairments is critical due to the uniqueness of the component.
6. It is very beneficial that experimenters have at their disposal some infrastructure that permits them to execute simple tests in advance of the final experimentation. Not all features can be replicated, of course, but having an I/W Framework deployment in parallel of the one in production will permit experimenters not to impact any of the working tests. Again, it is not a strict requirement that the central site hosts this parallel deployment, but it helps mainly in making sure that both deployments are exactly equivalent (e.g., same versions on all components).

7.2 Required capabilities for the rest of the sites

Sites not hosting the Interworking Framework, i.e., the spoke sites at the star topology, have some relaxed requirements, mainly related to supporting the globally defined facility features (for example, those of section 6). As a summary, the following can be cited:

1. Spoke sites need to implement the defined Interworking Connectivity solution at all planes, plus the associated authentication and security procedures when demanded.
2. Spoke sites are responsible for their own hardening policies, considering the rest of filters defined in the Project.

3. Spoke sites are responsible for implementing at least a local solution to permit access by Verticals' to their deployed experiments.
4. Spoke sites are responsible for implementing a controlled and secured mechanism to allow access by VNFs to Internet.
5. An IP addressing scheme must be agreed among all the sites (hub and spokes).
6. Spoke sites are responsible for deploying the monitoring solution which is globally defined in the Project

It is worth mentioning that the main responsible for the compatibility of the Multi-site NSO with the local NFVOs is not WP2, but WP3, which is the WP having the resources for implementation and adaptation of such interface.

7. Spokes sites are responsible for locally deploying the upgrades/updates in the local NFVOs that will permit the multi-site interaction.

7.3 Addition of new sites to 5G EVE infrastructure

Although extending the 5G EVE facility with new sites is not yet envisioned, it is something that partners are not willing to forget about. Considerations in the previous section are still valid, with mainly two modifications:

- ✓ 5G EVE partners cannot ensure the correct interoperability of the I/W Framework with every single version of every single NFVO, commercial or not. Therefore, it will be a responsibility of the new site to integrate its local NFVO with the implementation of the Or-Or interface that is agreed between WP3 and WP2.
- ✓ Site managers cannot be changing the IP addressing scheme each time a new site joins 5G EVE. Therefore, it shall be a responsibility of the new site to adapt its test bed to the proposal that is initially agreed by sites.

This list, together with the previous ones, may be extended in the future, as part of the implementation and deployment work done within WP3.

References

[1] 5G EVE Deliverable D3.1 “Interworking Capability Definition and Gap Analysis Document” – ICT-17-2018 5G EVE Project – 21/12/2018

[2] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/011/02.05.01_60/gs_NFV-IFA011v020501p.pdf

[3] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/014/02.05.01_60/gs_NFV-IFA014v020501p.pdf

[4] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-SOL/001_099/001/02.05.01_60/gs_NFV-SOL001v020501p.pdf

[5] [Online]. Available at: http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.2/csd01/TOSCA-Simple-Profile-YAML-v1.2-csd01.html

[6] [Online]. Available at: https://forge.etsi.org/gitlab/nfv/SOL001

[7] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-SOL/001_099/007/02.05.01_60/gs_NFV-SOL007v020501p.pdf

[8] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-SOL/001_099/004/02.05.01_60/gs_NFV-SOL004v020501p.pdf

[9] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/013/02.05.01_60/gs_NFV-IFA013v020501p.pdf

[10] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-SOL/001_099/005/02.05.01_60/gs_NFV-SOL005v020501p.pdf

[11] [Online]. Available at: https://nfvwiki.etsi.org/index.php?title=API_specifications#OpenAPIs

[12] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/030/03.01.01_60/gs_NFV-IFA030v030101p.pdf

[13] [Online]. Available at: https://www.etsi.org/deliver/etsi_gr/NFV-EVE/001_099/012/03.01.01_60/gr_NFV-EVE012v030101p.pdf

[14] [Online]. Available at: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3416

[15] [Online]. Available at: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3091

[16] [Online]. Available at: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3273

[17] [Online]. Available at: https://www.gsma.com/newsroom/wp-content/uploads//NG.116-v1.0-3.pdf

[18] [Online]. Available at: https://www.ngmn.org/fileadmin/user_upload/170307_5G_Network_and_Service_Management__including_Orchestration_2.12.7.pdf

[19] Introducing ONOS - a SDN network operating system for Services Providers – http://onosproject.org/wp-content/uploads/2014/11/Whitepaper-ONOS-final.pdf

[20] TAPI Wiki – https://wiki.opennetworking.org/display/OTCC/TAPI

[21] [Online]. Available at: https://tools.ietf.org/html/rfc6241

[22] [Online]. Available at: https://tools.ietf.org/html/rfc6020

[23] 5G-TRANSFORMER Deliverable D3.1 “Definition of vertical service descriptors and SO NBI” – 5G-TRANSFORMER Project – March 2018

[24] 5G-TRANSFORMER Deliverable D4.1 “Definition of service orchestration and federation algorithms, service monitoring algorithms” – 5G-TRANSFORMER Project – March 2018

[25] 5G-TRANSFORMER Deliverable D2.1 “Definition of the Mobile Transport and Computing Platform” – 5G-TRANSFORMER Project – March 2018

[26] 5GinFIRE Deliverable D4.2 “Intermediate Report on the MANO Platform” – 5GinFIRE Project – 23/01/2019

[27] A. Sgambelluri, A. Milani, J. Czentye, J. Melian, W.Y. Poe, F. Tusa, O.G. de Dios, B. Sonkoly, M. Gharbaoui, F. Paolucci, E. Maini, G.Giuliani, A. Ramos, P. Monti, L.M. Contreras-Murillo, I. Vaishnavi, C.J. Bernardos Cano, R. Szabo, “A Multi-Operator Network Service Orchestration Prototype: the 5G Exchange”, IEEE OFC 2017, Mar. 2017.

[28] I. Afolabi, A. Ksentini, M. Bagaa, T. Taleb, M. Corici, and A. Nakao, "Towards 5g Network Slicing over Multiple-Domains," IEICE Transactions on Communications, vol. E100.B, no. 11, pp. 1992–2006, 2017.

[29] L. Xu, H. Assem, I. G. B. Yahia, T. S. Buda, A. Martin, D. Gallico, M. Biancani, A. Pastor, P. A. Aranda, M. Smirnov, D. Raz, O. Uryupina, A. Mozo, B. Ordozgoiti, M.-I. Corici, P. O'Sullivan, and R. Mullins, "CogNet: A network management architecture featuring cognitive capabilities". IEEE, Jun. 2016, pp. 325–329.

[30] [Online]. Available at: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144

[31] 5G EVE Deliverable D1.1 "Requirements definition and analysis from participant vertical industries" – ICT-17-2018 5G EVE Project – 31/10/2018

[32] [Online]. Available at: https://www.etsi.org/deliver/etsi_gr/NFV-IFA/001_099/028/03.01.01_60/gr_NFV-IFA028v030101p.pdf

[33] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/005/02.01.01_60/gs_NFV-IFA005v020101p.pdf

[34] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/006/02.01.01_60/gs_NFV-IFA006v020101p.pdf

[35] [Online]. Available at: https://osm.etsi.org/wikipub/index.php/NBI_API_Description

[36] [Online]. Available at: https://osm.etsi.org/wikipub/index.php/OSM_Information_Model

[37] [Online]. Available at: https://www.tmforum.org/open-apis/

[38] [Online]. Available at: https://docs.onap.org/en/casablanca/submodules/externalapi/nbi.git/docs/offeredapis/offeredapis.html

[39] [Online]. Available at: https://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/032/03.02.01_60/gs_NFV-IFA032v030201p.pdf

[40] [Online]. Available at: https://software.es.net/iperf/

[41] [Online]. Available at: https://github.com/mami-project/trafic

[42] [Online]. Available at: https://docs.openstack.org/keystone/latest/

[43] [Online]. Available at: https://shiro.apache.org/

[44] [Online]. Available at: https://openvpn.net/