Deliverable D4.1 - Operational MANO Platform

Editor:	Diego R. López, TID Iván Vidal, UC3M	
Deliverable nature:	Report (R)	
Dissemination level:	Public (PU)	
Date: planned | actual	31 December 2017	22 January 2018
Version | No. of pages	1.0	43
Keywords:	5G, NFV, MANO, NFVI, service orchestration, NS, VNF, VVF, VxF, Multi-site	

Abstract

This document describes how the 5GINFIRE MANO platform has been deployed, and its main technical characteristics. According to the NFV architecture, the MANO platform is in charge of the management and orchestration of the different network functions and the services built by connecting them. The 5GINFIRE MANO is able to support multi-site deployments, and it is based on Open Source MANO (OSM), an open-source implementation of the NFV MANO stack hosted by ETSI.

Disclaimer

This document contains material, which is the copyright of certain 5GINFIRE consortium parties, and may not be reproduced or copied without permission.

All 5GINFIRE consortium parties have agreed to full publication of this document.

Neither the 5GINFIRE consortium as a whole, nor a certain part of the 5GINFIRE consortium, warrant that the information contained in this document is capable of use, nor that use of the information is free from risk, accepting no liability for loss or damage suffered by any person using this information.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 732497. This publication reflects only the author's view and the European Commission is not responsible for any use that may be made of the information it contains.

Impressum

Full project title: Evolving FIRE into a 5G-Oriented Experimental Playground for Vertical Industries

Short project title: 5GINFIRE

Number and title of work-package: WP4 Core MANO Service Management and Orchestration

Number and title of task:

- Task 4.1: Orchestration platform adaptation and integration
- Task 4.2: Orchestration platform deployment

Document title: Deliverable D4.1 - Operational MANO Platform

Editor: Diego R. López, Telefónica Investigación y Desarrollo SA; Iván Vidal, Universidad Carlos III de Madrid.

Work-package leader: Diego R. López, company: Telefónica Investigación y Desarrollo SA

Copyright notice

Executive summary

In the NFV architecture, the MANO framework is in charge of the management and orchestration of the network functions, supporting their lifecycle, and the composition of these functions to build Network Services. Management and orchestration are applied following a set of requirements, usually expressed as policy statements. The MANO framework reads and processes these policies. The 5GINFIRE project [1] has selected one of the most advanced open-source MANO platforms currently available, Open Source MANO (OSM) [2] as the base for its MANO framework.

The planned 5GINFIRE environment poses a series of particular challenges for the MANO platform to be used, especially related to the high diversity of functions to be considered and the multi-domain nature of the infrastructure to be managed. The technical solution adopted by the project considers the utilization of a single orchestration domain, where an NFV orchestrator manages and coordinates the creation of network services on three participating infrastructures:

- A site at the 5TONIC lab, made available by TID and UC3M.
- A site located at ITAv.
- A site made available through a collaborative agreement by UNIVBRIS and BIO.

These participating experimental facilities are at the core of the 5GINFIRE MANO deployment, each one with its independent resource management element (VIM). Besides this, at the time of writing, an additional experimental infrastructure is being created at UFU.

In such a distributed infrastructure, there is a strong need for a clear definition of the mechanisms supporting inter-site communications that will also be used to support the integration of additional infrastructures into the 5GinFIRE MANO platform, such as those that are expected to be incorporated through the Open Call process of the project. From the point of view of the MANO platform, we can consider inter-site control-plane communications (within the MANO stack, and between this MANO stack and the deployed functions), and inter-site data-plane communications (among the deployed functions). Mechanisms for both kind of communications are described in this document.

The WP4 5GINFIRE team has addressed the orchestration challenges in the project, and produced a MANO platform suitable for them, considering the aspects related to a multi-site and multi-user environment. The MANO platform has been validated, and its ability to deploy and execute cross-site services demonstrated. A set of valuable documentation on these tests, usable as future reference by 5GINFIRE experimenters, has been produced. In this document, we cover all these aspects, presenting a detailed description of the first stable version of the 5TONIC MANO platform, along with a summary of the tests done.

In order to properly address all the requirements for the orchestration platform of the 5GINFIRE environment, the team has started to work on several enhancements to the current platform elements. The document describes the already initiated work on these enhancements, which have a different degree of maturity. All of them will be contributed to the relevant upstream communities, and a continuous integration environment will be set in place to guarantee a seamless evolution of the framework.

List of authors

Company	Author
TID	Diego R. López, Juan Rodríguez Martínez
UC3M	Borja Nogales Dorado, Iván Vidal Fernández
UNIVBRIS	Aloizio P. Silva
ITAv	Diogo Gomes, Eduardo Sousa
UFU	Flávio de Oliveira Silva

Table of Contents

Abbreviations

MANO: Management and Orchestration

NFV: Network Function Virtualization

NFVI: NFV Infrastructure

NFVO: NFV Orchestrator

NS: Network Service

SO: Service Orchestrator

RO: Resource Orchestrator

VCA: VNF Configuration and Abstraction

SDN: Software Defined Networks

VIM: Virtual Infrastructure Management

VNF: Virtualized network function

VNFM: VNF Manager

VVF: Virtualized Vertical Function

VxF: Virtualized (Network or Vertical) Function

1 Introduction

In the NFV architecture, the MANO framework is in charge of the *management and orchestration* of the Virtual Network Functions (VNFs), supporting the lifecycle of such functions, and the composition of these functions to build Network Services (NSs).

This lifecycle management includes all events related to the execution of a VNF, since its initial incorporation to the cloud environment where it will run (the so-called *onboarding*) up to an eventual *decommission* of such a function. And naturally encompasses essential events such as particular *instantiations* and *activations*, and those related to the properties of a cloud-based functions, like *scaling* events. Given the final goal of NFV is to provide network services, these events at the VNF level have to be coordinated at the NS level as well, requiring the MANO framework to *orchestrate* individual VNF-related events to ensure the proper behaviour of the services under its control.

Management and orchestration are applied following a set of requirements, usually expressed as policy statements within metadata structures (the *descriptors*) that define the intended properties a service and its component functions have to comply to. The MANO framework reads and processes these descriptors, managing the identified components, and interpreting how to satisfy the policy statements that define their intended behaviour.

The 5GINFIRE project has selected one of the most advanced open-source MANO platforms currently available, Open Source MANO (OSM) [2] due to: 1) its high degree of maturity; 2) its proven support for some of the required characteristics of the 5GINFIRE management and orchestration environment, especially the support for a distributed and diverse infrastructure in different administrative domains; 3) the cloud-native mechanisms for function control and management, suitable to address the high variety of potential VxFs; and 4) the ability to apply *continuous integration* techniques to the platform.

This document describes the initial deployment of the 5GINFIRE MANO platform, describing the process followed for it, and analysing the main aspects related to the management of virtual functions in general. The 5GINFIRE approach implies a common NFV orchestrator, interacting at its *northbound* interface with the portal, and controlling a set of *Virtualized Infrastructure Managers* (VIMs), one at each participating site [3].

The following sections addresses this description, starting in Section 2 with some initial considerations on the design of the 5GINFIRE MANO solution, describing the requirements, the architecture proposed to address them, and the planning followed to provide a MANO platform implementing it. Section 3 considers the main characteristics of the management stack at each testbed infrastructure, together with a description of their interconnection with the common orchestrator, and a discussion of how connectivity is provided among functions (and experimenters) in the distributed 5GINFIRE infrastructure. Section 4 discusses the functional validation of the MANO platform, including the different phases the project has followed for it. Finally, the future enhancements planned for the MANO platform are introduced, including how they will be applied following a continuous integration approach.

2 Design of the 5GINFIRE MANO platform

The planned 5GINFIRE environment poses a series of particular challenges for the MANO platform to be used, especially related to the high diversity of functions to be considered and the multi-domain nature of the infrastructure to be managed. This section describes the main requirements, and introduces the 5GINFIRE MANO architecture and the stages for its deployment.

2.1 Requirements

The Virtualized Network Functions (VNFs) and Virtualized Vertical Functions (VVF) to be deployed and managed by the 5GinFIRE MANO platform, denoted as VxFs in this document, are intended to span a very wide area of the network function domain, and can thus represent, from a service provider's and developer's perspective, a set of significant use cases at different vertical domains. The requirements imposed to the NFV experimentation environment of 5GinFIRE, and particularly to its MANO platform, include:

1. The realization of a multi-site NFV ecosystem through the definition of an architecture composed of:
    a. A distributed NFV infrastructure (NFVI), where the hardware and software resources that support a cloud virtualized network, computational and storage resources will be under different administrative domains
    b. Generalized VxFs running over the distributed NFVI
    c. The management and orchestration of the lifecycle of VxFs.
2. The ability to specify/model the Experimental Vertical Instances (EVIs) and their internal structure as a graph of VxFs along with the corresponding resource requirements and policies expressed with domain specification languages (DSLs), which are then processed by MANO functional elements for their instantiation.
3. Adopting and being interoperable with current cloud/SDN/NFV/MEC standards.
4. Being interoperable with FIRE standards and facilities.
5. Using open standards and integrating technologically mature, and widespread open source tool sets.
6. Enabling experimentation, making effortless for experimenters to deploy experimentation scenarios.
7. Enabling experimentation in multiple levels, either for applications, services or VxFs on top of 5G-enabled experimentation infrastructure that also supports specialized infrastructure for verticals e.g. automotive.
8. Focusing on state-of-the-art EVI 5G automotive testbed resources while still being generic to enable as many EVIs as possible from other verticals (ie Manufacturing, Entertainment, IoT, eHealth, etc) or even cross- vertical EVIs.
9. The validation of the MANO deployment in a real use case, along with the fulfillment of its corresponding defined requirements.

2.2 Architectural design

As already commented, the main goal of the 5GinFIRE MANO platform is to enable the management and orchestration of network services (NS) across the experimental infrastructures provided by 5GinFIRE partners. To fulfil this objective, the technical solution adopted by the project (see Figure 1) considers the utilization of a single orchestration domain, where an NFV orchestrator, implemented with Open Source MANO (OSM) [5], manages and coordinates the creation of network services, being a network service generally defined as a composition of VxFs. Each of these VxFs may be in turn deployed at any of the experimental infrastructures made available by 5GinFIRE partners. At the time of writing, three of these infrastructures have been enabled for experimentation activities through the 5GinFIRE MANO platform: 1) an infrastructure at the global 5G Telefonica Open Innovation Laboratory (5TONIC) [4], made available by TID (founding member of 5TONIC) and UC3M (member of 5TONIC); 2) an infrastructure located at ITAv; and 3) an infrastructure made available through a collaborative agreement by UNIVBRIS and BIO. An additional experimental infrastructure is under development at UFU.

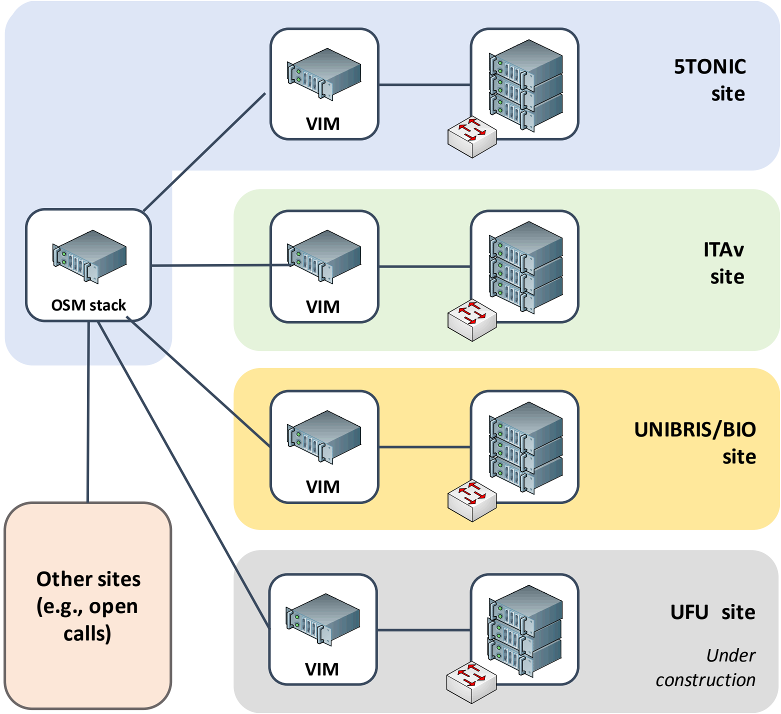

Figure 1: Technical solution adopted by 5GinFIRE

Each partner running an experimental infrastructure (hereafter referred to as testbed provider) is in charge of the deployment and maintenance of a Virtualized Infrastructure Manager (VIM), compliant with the OSM software stack. On top of that, the NFV orchestrator of 5GinFIRE, which according to the agreements made by the project consortium will be deployed at 5TONIC, will interact with the VIMs of the testbed providers involved in a service deployment, coordinating the allocation and setup of the computing, storage and network resources which are necessary for the instantiation and interconnection of the VxFs that compose the network service.

Additional sites (e.g., coming from the Open Call process of 5GinFIRE), with heterogeneous infrastructure and equipment, can be flexibly incorporated as needed, as long as they support a compliant VIM [5] and they set up the inter-site connection mechanisms defined in this document (see section 3.2).

Figure 2 presents an overview of the architectural design of the 5GinFIRE MANO platform, and its relationship with the components of the NFV reference architectural framework defined by ETSI [6]. The current orchestration service of 5GinIFIRE is based on Open Source MANO release TWO[1] [5], which provides a Service Orchestrator (**SO**), a Resource Orchestrator (**RO**) and a VNF Configuration and Abstraction (**VCA**) module.

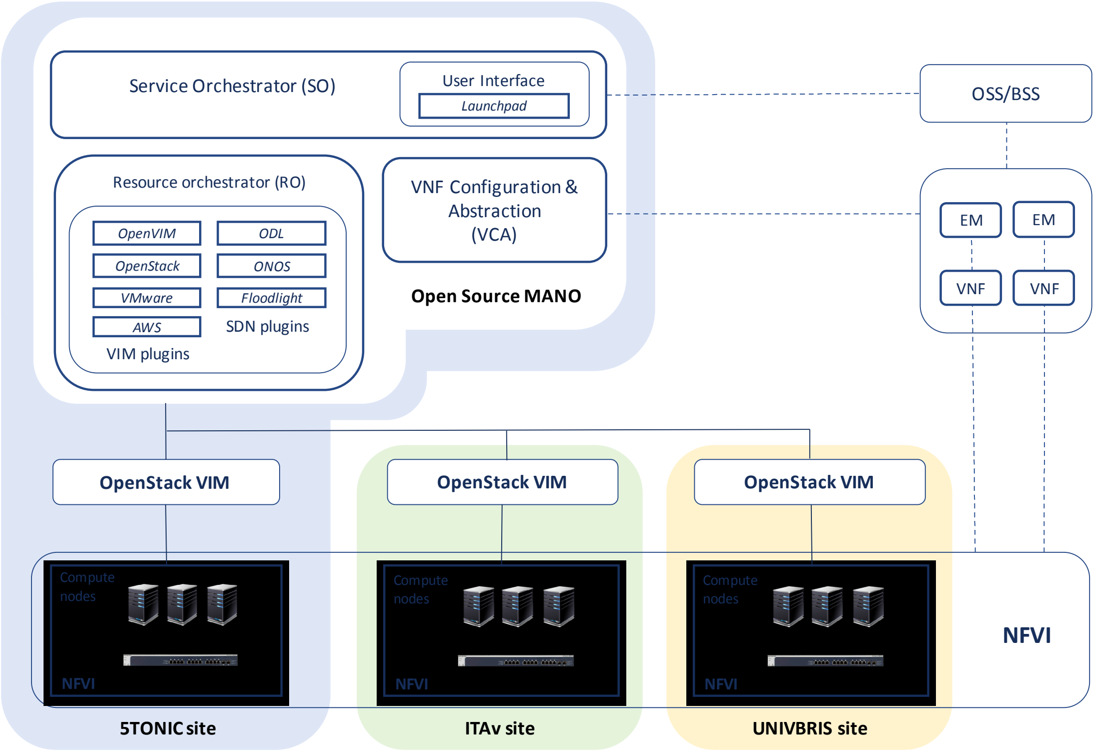

Figure 2: Architectural design of the 5GinFIRE MANO platform

The SO provides the point of contact for external entities (e.g., the 5GinFIRE portal) to interact with the OSM system. It supports the lifecycle management of network services, coordinating the creation and deletion of network services composed of multiple VxFs. For this purpose, the SO interfaces with the RO and the VCA modules of the OSM architecture. Additionally, the SO provides other essential enabling functionalities, such as the management of NS/VNF descriptors and packages. In OSM Release TWO, the SO module includes a graphical user interface[2], which provides an intuitive mechanism to ease the on-boarding of NS/VNF packages and the lifecycle management of NS instances.

The RO module coordinates the allocation and configuration of computing, storage and network resources under the control of the different VIMs and SDN controllers, in order to support the execution and interconnection of VxFs. OSM Release TWO supports multiple types of VIMs through a plugin model, including OpenVIM, OpenStack, VMWare vCloud

[1] At the time of writing, ETSI Open Source MANO (ETSI OSM) has announced the availability of OSM Release THREE (see http://www.etsi.org/news-events/news/1239-2017-11-news-etsi-open-source-mano-announces-release-three). Considering the early state of this release and the deadlines established within 5GinFIRE, the project will keep the consolidated agreement of using OSM Release TWO for the first version of the 5GinFIRE MANO platform, considering the migration to Release THREE during the second year of the project lifetime.

[2] From Release THREE, this graphical user interface is separated from the SO module.

Director and Amazon Web Services Elastic Compute Cloud. Additionally, this plugin model enables the RO to directly manage a number of SDN controllers, including OpenDaylight, Floodlight and ONOS.

The VCA module is aligned with the VNF Manager defined by the ETSI NFV reference architectural framework, supporting day-1 configuration of VNFs. With this purpose, the VCA has an interface to Juju, which allows configuring VNFs through the execution of Juju charms[3] that can be specified within VNF packages.

The orchestration service provided by the OSM software stack is capable of managing and orchestrating the deployment of network services across the set of datacentres set up by the 5GinFIRE testbed providers (TID&UC3M, ITAv, UNIVBRIS&BIO). Each of these datacentres includes a specific NFV infrastructure (NFVI), made available for experimentation purposes in the scope of the project and being under the control of a VIM, which enables the isolation among experiments at its corresponding site. The VIM solution that has been independently chosen by 5GinFIRE testbed providers, for the first stable version of the 5GinFIRE MANO platform, is OpenStack Ocata [7].

2.3 Approach and methodology

The approach followed to design and deploy the initial 5GinFIRE MANO platform has been iterative, where components were tested, deployed when they were proven to be functional and integrated. The work started with the installation on an operational small-scale MANO stack, which was replicated at each of the sites of WP4 partners. This allowed the work to progress independently at each partner location.

Every two weeks, there is a work package progress meeting where every partner can discuss the status of each deployment, problems found and possible solutions. In these progress meetings, information related to other work packages is discussed. All the information generated during these meetings has been made available to the other work packages, to allow direct collaboration within the project. Besides these periodic meetings, online communication has always been available to resolve issues and exchange new information that would help others. This approach allowed us to test different configurations at each site, validating the operation of the MANO platform with these configurations. Using an agile methodology, the following steps have been identified and carried out to design and deploy the 5GinFIRE MANO platform:

- Test the MANO software stack in isolation.
- Test the MANO software stack with the VIM and NFVI components.
- Test the MANO software stack with multiple VIMs and NFVIs.
- Interconnect different VIMs and NFVIs with the common orchestrator.

These steps are visible in the work plan defined Figure 3, where, in an iterative way, we started from a small-scale deployment, and evolved into a bigger and more complex one (in this figure, green colour denotes completed activities; yellow colour refers to activities in progress; blue colour means planned activities).

[3] Creating your own VNF charm (Release TWO), OSM Wiki (last access: December 2017): https://osm.etsi.org/wikipub/index.php/Creating_your_own_VNF_charm_(Release_TWO).

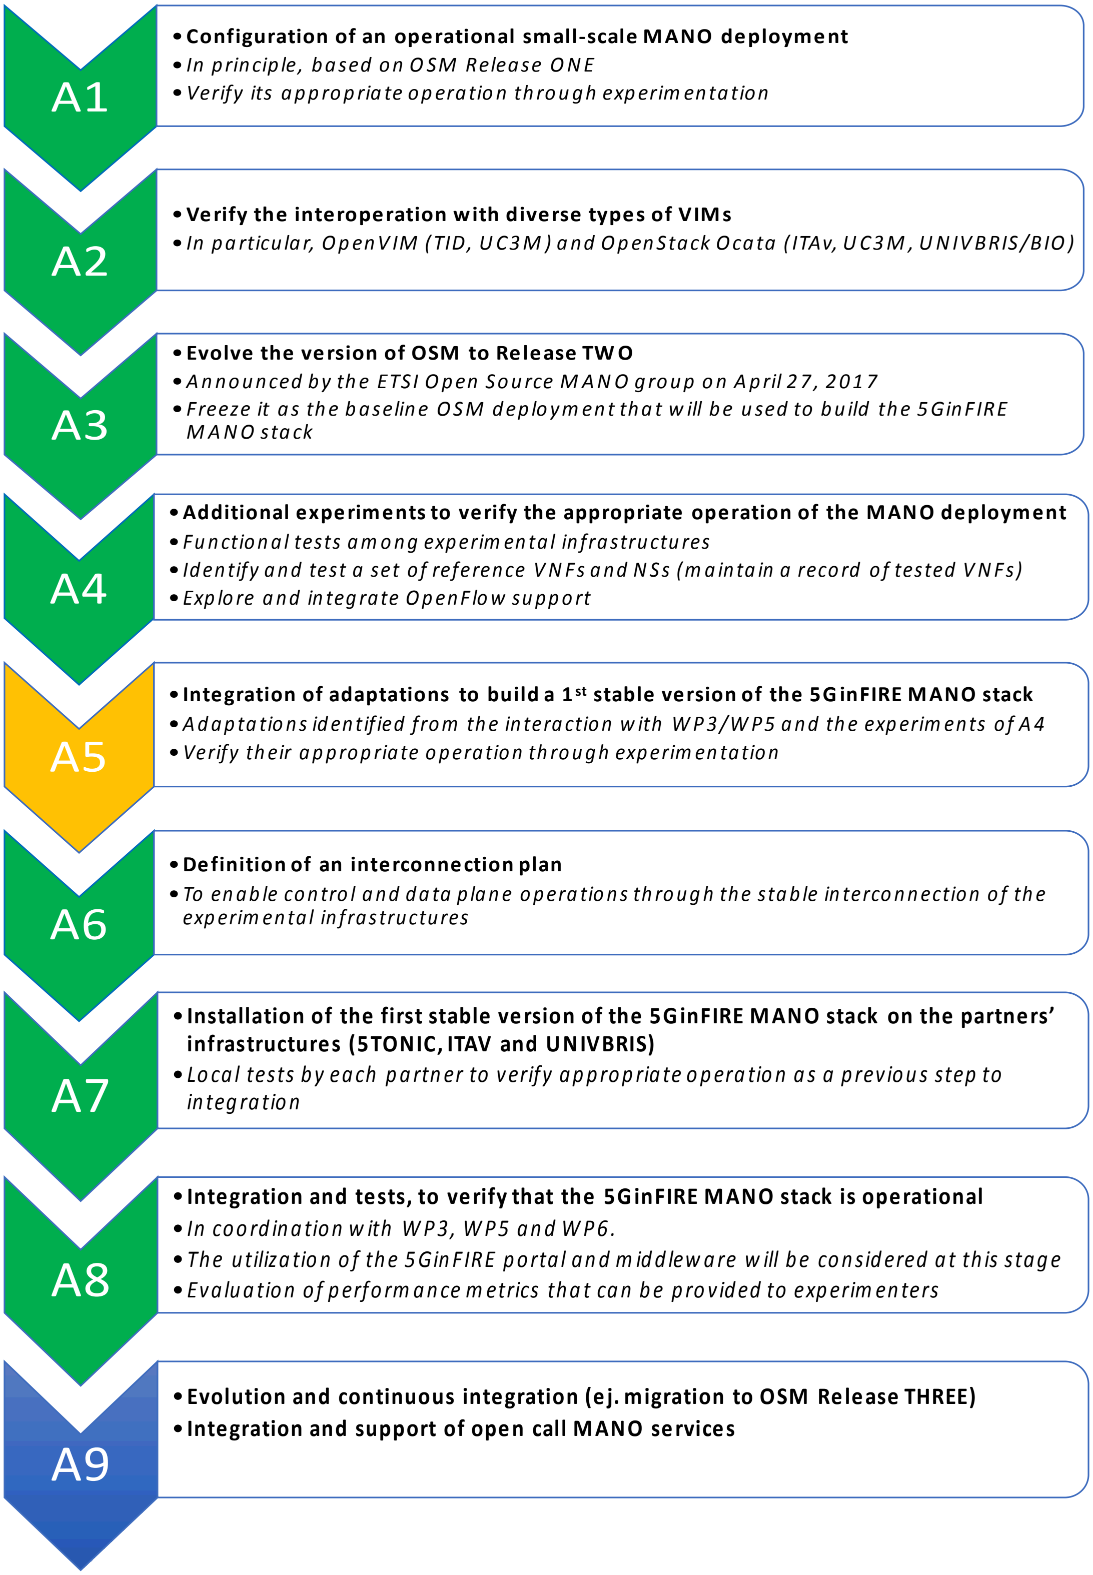

A1

- **Configuration of an operational small-scale MANO deployment**
- *In principle, based on OSM Release ONE*
- *Verify its appropriate operation through experimentation*

A2

- **Verify the interoperation with diverse types of VIMs**
- *In particular, OpenVIM (TID, UC3M) and OpenStack Ocata (ITAv, UC3M, UNIVBRIS/BIO)*

A3

- **Evolve the version of OSM to Release TWO**
- *Announced by the ETSI Open Source MANO group on April 27, 2017*
- *Freeze it as the baseline OSM deployment that will be used to build the 5GinFIRE MANO stack*

A4

- **Additional experiments to verify the appropriate operation of the MANO deployment**
- *Functional tests among experimental infrastructures*
- *Identify and test a set of reference VNFs and NSs (maintain a record of tested VNFs)*
- *Explore and integrate OpenFlow support*

A5

- **Integration of adaptations to build a 1st stable version of the 5GinFIRE MANO stack**
- *Adaptations identified from the interaction with WP3/WP5 and the experiments of A4*
- *Verify their appropriate operation through experimentation*

A6

- **Definition of an interconnection plan**
- *To enable control and data plane operations through the stable interconnection of the experimental infrastructures*

A7

- **Installation of the first stable version of the 5GinFIRE MANO stack on the partners' infrastructures (5TONIC, ITAV and UNIVBRIS)**
- *Local tests by each partner to verify appropriate operation as a previous step to integration*

A8

- **Integration and tests, to verify that the 5GinFIRE MANO stack is operational**
- *In coordination with WP3, WP5 and WP6.*
- *The utilization of the 5GinFIRE portal and middleware will be considered at this stage*
- *Evaluation of performance metrics that can be provided to experimenters*

A9

- **Evolution and continuous integration (ej. migration to OSM Release THREE)**
- **Integration and support of open call MANO services**

Figure 3: Overview of the WP4 workplan

3 Deployment of the 5GINFIRE MANO platform

3.1 Description of the NFV experimental infrastructures

At the core of the 5GINFIRE MANO deployment are the participating experimental facilities, each one with its independent resource management element (VIM). It is important to describe their nature and structure to better understand how the general 5GINFIRE MANO framework had to be deployed.

3.1.1 5TONIC site

The global 5G Telefonica Open Network Innovation Centre (5TONIC) [4] has been established in Madrid (Spain) as a leading European hub for knowledge sharing and industry collaboration in the area of 5G technologies. The laboratory provides an open research and innovation ecosystem for industry and academia that will promote joint project development, joint entrepreneurial ventures, discussion fora, and a site for events and conferences, all in an international environment of the highest impact. 5TONIC will also serve to evaluate and demonstrate the capabilities and interoperation of pre-commercial 5G equipment, services and applications. Currently, the 5TONIC laboratory has nine members: Telefonica, Institute IMDEA Networks, Ericsson, Intel, Commscope, Universidad Carlos III de Madrid, Cohere Technologies, Artesyn Embedded Technologies, and InterDigital.

The 5TONIC laboratory, as a multipurpose environment, counts with multiple racks, which may be flexibly interconnected according to any experimentation requirements, along with a common infrastructure to aid experimentation, trials and demonstrations with 5G products and services. In particular, secure external access may be provided via VPN gateways, allowing different solutions to support management, control and data operations from remote network locations, depending on specific requirements. Due to confidentiality reasons, we cannot disclose all the software and hardware available in the laboratory, which also includes experimental prototypes from the industry and academic members. Consequently, in the following we keep our description concrete, describing the main infrastructure and equipment that will be offered for experimentation to 5GinFIRE. A graphical representation of this infrastructure is schematized in Figure 4.

With respect to the orchestrator software stack, the first stable version of the 5GINFIRE MANO platform will be based on OSM Release TWO, and will run in a virtual machine using a server computer with 16 cores, 128 GB RAM, 2 TB NLSAS hard drive and a network card with 4 GbE ports and DPDK support. This server computer will also host a VIM instance based on OpenStack Ocata [7]. An additional and independent instance of OpenStack Ocata will be deployed in a separate server computer with six cores, 32GB of memory, 2TB NLSAS and a network card with four GbE ports and DPDK support. This server computer will also host a virtual machine providing a relay functionality (a jump machine) to support the access of experimenters to VxFs (see section 3.4). In both VIM instances, the OpenStack networking service was installed to support layer-3 services, and the ML2 plug-in of OpenStack was configured to use Linux bridges.

The utilization of two separate VIMs will allow allocating experiments to two separate NFV infrastructures. On the one hand, the laboratory will include a dedicated NFVI that will be allocated to experimentation activities within 5GINFIRE. This infrastructure (5GINFIRE NFVI in Figure 4) will consist of a set of three server computers, each with six cores, 32GB of memory, 2TB NLSAS and a network card with four GbE ports and DPDK support. These servers will be interconnected by a GbE data-plane switch. On the other hand, the 5TONIC laboratory will also offer a second NFVI for experimentation, based on two high-profile servers, each equipped with eight cores in a NUMA architecture, 128GB RDIMM RAM, 4TB SAS and eight 10Gbps Ethernet optical transceivers with SR-IOV capabilities. These servers are currently interconnected in the data plane by a 24-port 10Gbps Ethernet switch. The latter NFVI (IMDEA NFVI in Figure 4) forms part of the infrastructure of the IMDEA Networks Institute at 5TONIC, and may be used to support high resource demanding experiments. This specific infrastructure will be available for experimentation in 5GINFIRE under specific terms and conditions as described at the project website [1]. Finally, the experimentation infrastructure offered to 5GinFIRE includes a number of server computers (not shown in the figure) to support complementary functionalities, such as aiding experimentation (e.g. hosting client applications), deploying a network management system, performing access-control functionalities, storing data (e.g. disk backups, VM images or code), etc.

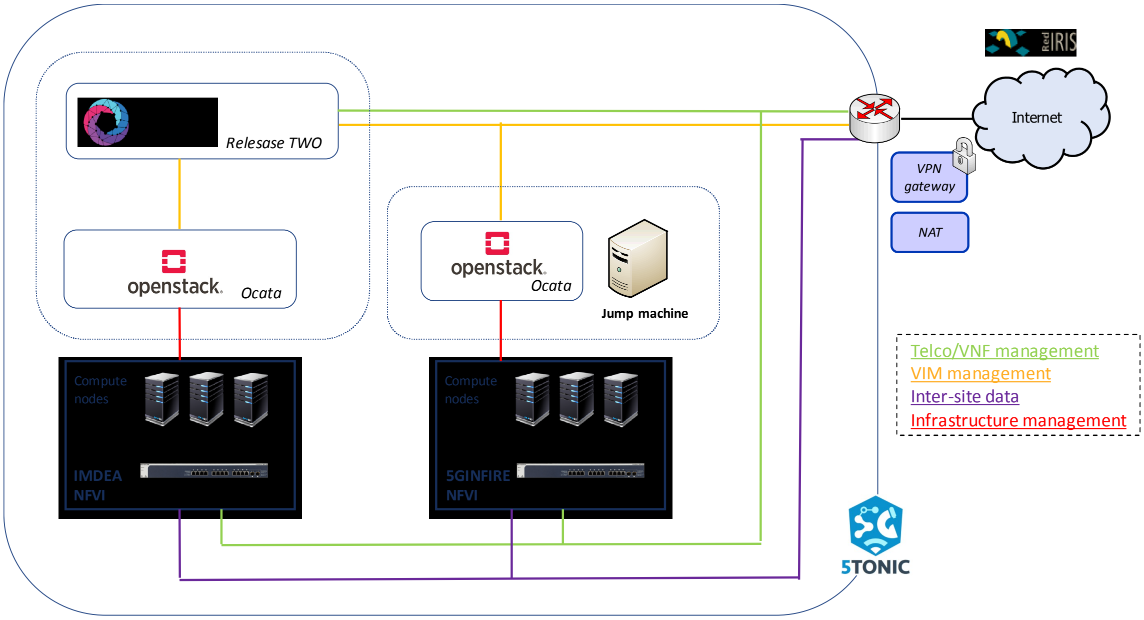

Figure 4: Overview of the 5TONIC equipment available in 5GinFIRE

Last but not least, a relevant feature of the deployed MANO platform is that it would be feasible to extend the experimentation infrastructure offered by 5GinFIRE with the public cloud support of OSM Release TWO.

3.1.2 ITAv site

The ITAv site provides an infrastructure based on Openstack Ocata. This infrastructure is connected to the 5GINFIRE OSM-based orchestrator (see Figure 5). Two servers compose the NFVI in ITAv:

1. Orphic
   - Cores: 24
   - Memory: 192 GB
   - Network: 4 x 1Gbps interfaces (supports passthrough, DPDK and SR-IOV)

- Storage: 2 x 1TB SAS3 drives

2. Aeolus
   - Cores: 16
   - Memory: 256 GB
   - Network: 4 x 1Gbps interfaces (supports passthrough)
   - Storage: 2 x 1TB SAS2

ITAv's Openstack deployment has one controller node where all the services are installed and then there the components mentioned above, where only the compute service is installed.

ITAv's Openstack deployment has three VLAN networks, which are: Control, Data and Management. The Control VLAN is for the control plane packets, while the Data VLAN is used for the data plane packets. The Management VLAN is used to access the nodes in the Openstack deployment, and it is used by the Openstack services to communicate with each other, thus not being visible in the other VLANs.

Networks inside the Openstack deployment are handled using Neutron and Linux Bridges. Projects internal networks are created using VXLAN, and the external ones use VLAN networks.

This infrastructure is connected to a 24-port 1Gbps switch. The switch supports the control and data plane, which are separated using VLANs. Each of these VLAN networks can be extended in order to support and interconnect more devices (using layer 2 or layer 3 technologies).

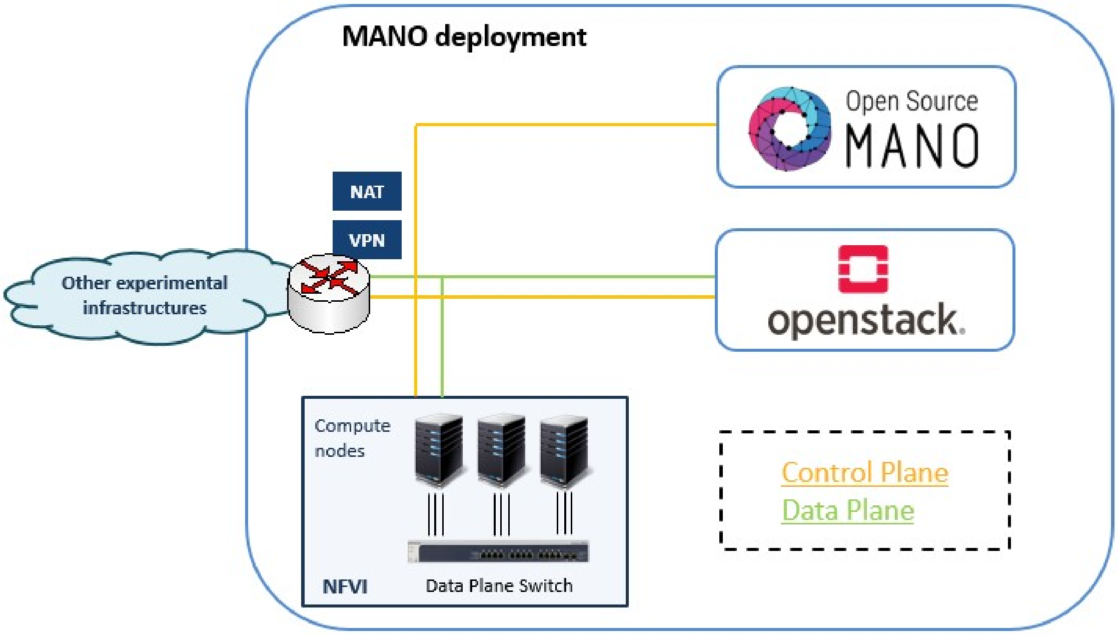

Figure 5: Deployment configuration at the ITAv site

At the ITAv site, there is also an OSM Release TWO deployed, intended to be used for local deployments and for testing purposes.

3.1.3 UNIVBRIS & BIO site

The overall 5GinFIRE UNIVBRIS-BIO site is based on the Network Function Virtualization Infrastructure (NFVI) foundation specified by ETSI NFV reference architecture [6]. The UNIVBRIS-BIO site is composed by three main building blocks: compute, storage and network. All the virtualized resources being provided by the UNIVBRIS-BIO site are available through the OSM-based common orchestrator located at the 5TONIC site. In particular, the UNIVBRIS-BIO site's purpose is to provide excellent testing platform for heterogeneous experimentations and at the same time guarantee computational resources or/and slicing for hosting, deploying, instantiating and supporting VxF's life cycle enabling to conduct rigorous, transparent and replicable testing of NFV ecosystem.

Figure 6 shows the main building blocks of UNIVBRIS-BIO site.

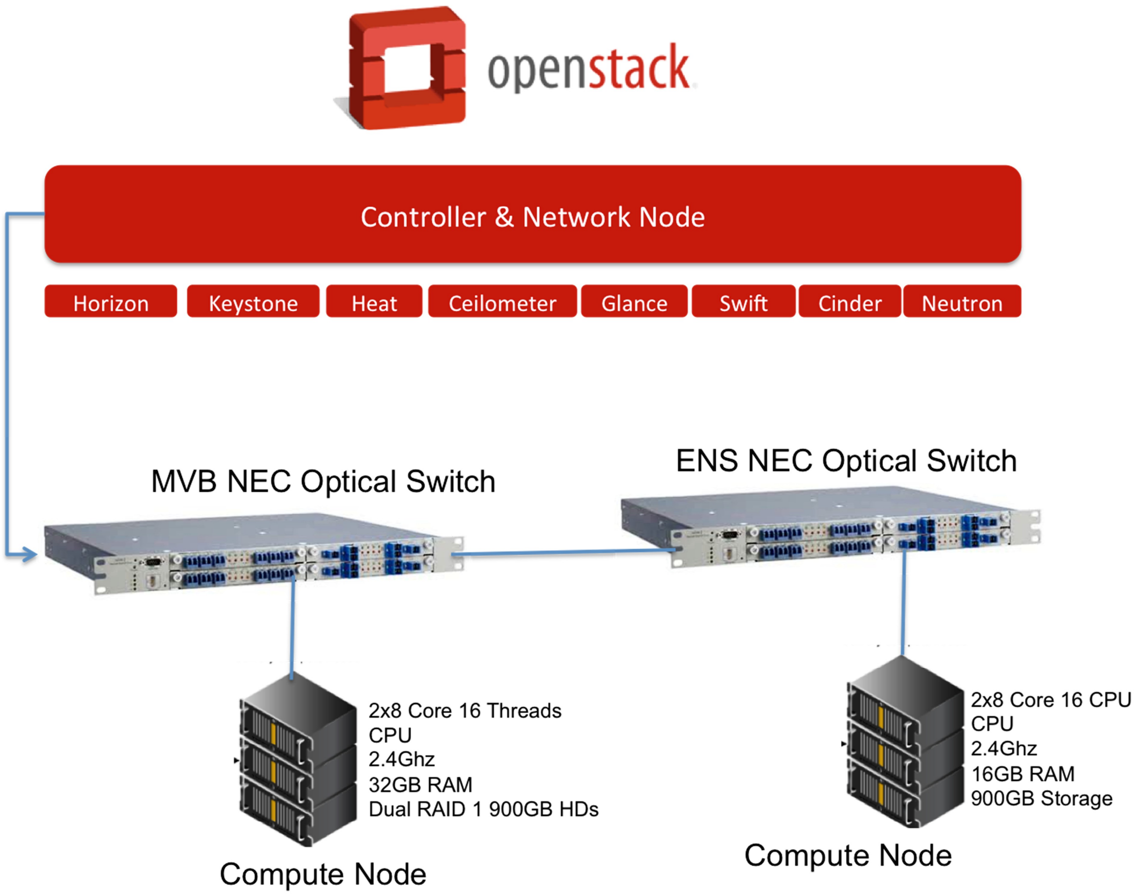

Figure 6: UNIVBRIS-BIO Site Architecture

The cloud environment consists of an OpenStack Ocata instance that operates on top of CentOS 7 operation system. Self-service networks capability is provided by the site where new networks can be created by clients as needed, and connected to their virtual routers. A VXLAN tunnel connects the controller and compute nodes and networking nodes. For network slicing there is a trunk around the Bristol city containing VLANs which can be accessed via switch access ports and through individual SSIDs broadcasting specific VLANs according to experimenter needs. Each VLAN gateway resides on a virtual router for each

OpenStack project. Depending on the experimenters needs the VLAN can be routed locally or separately with a public IP.

The site includes also a public API proxy node hosting OpenStack dashboard and DNS. The OpenStack public APIs are proxied to the API node.

While many switches claim support for Openflow, most of them only provide partial support and vital functionality/features are missing. BIO recommends using PICA8 switches as these devices fully support Openflow. BIO is upgrading their core switches from to PICA8s due to missing Openflow functionality in the current ones.

3.2 Inter-site communications

This section describes the main mechanisms that have been agreed and implemented by 5GinFIRE testbed providers to support inter-site communications. These mechanisms will also be used to support the integration of additional infrastructures into the 5GinFIRE MANO platform, such as those that are expected to incorporate through the Open Call process of the project. From the point of view of the MANO platform, inter-site communications encompass the following types of data exchanges:

1) Communications between the OSM stack, deployed at 5TONIC, and the VIMs and SDN controllers operated by testbed providers. This type of communications allows the OSM stack to coordinate the allocation and configuration of computing, storage and network resources at the diverse datacentres.
2) Communications between the OSM stack and the VxFs deployed at each datacentre, to support day-1 configuration of VxFs via Juju charms.
3) Inter-site communications between VxFs, to enable a VxF at one datacentre to exchange data with VxFs deployed at other datacentres.

The first two types of communications (1 and 2) involve the exchange of control information, and will be referred to as **inter-site control-plane communications** in the remainder of this document; the third type of communications supports the exchange of data among VxFs located at different sites, and will be denominated **inter-site data-plane communications**.

3.2.1 Establishment of inter-site control and data-plane communications

To support the exchange of control and data-plane information among 5GinFIRE sites, the approach taken by 5GinFIRE has consisted in the utilization of an overlay network architecture based on Virtual Private Networks (VPNs). This approach is schematized in Figure 7, where the three types of communications are represented (type 1 in yellow; type 2 in green; type 3 in purple}.

The motivation behind the selected approach is threefold. On the one hand, the criticality of the resources that participate in the communications, which are typically shared among the multiple projects and users of each experimental infrastructure, requires a solution that guarantees an appropriate access control to these resources. As an example, the communication of information with the OSM stack should be limited to the 5GinFIRE portal, the VIMs of testbed providers and the deployed VxFs. On the other hand, given that sites are physically separated at different network locations, inter-site communications may traverse multiple untrusted network domains. Consequently, the adopted solution must support the necessary mechanisms to guarantee the confidentiality, integrity and authenticity of these communications. Last but not least, considering the heterogeneous nature of the network technologies, topologies and services that exist at each testbed provider premises, the approach to follow must allow these testbed providers to independently configure the most appropriate mechanisms to support the delivery of control and data traffic across their network segments (i.e. testbed providers must retain the control over their network infrastructures). These challenging requirements can be satisfied with the overlay network approach considered by the project.

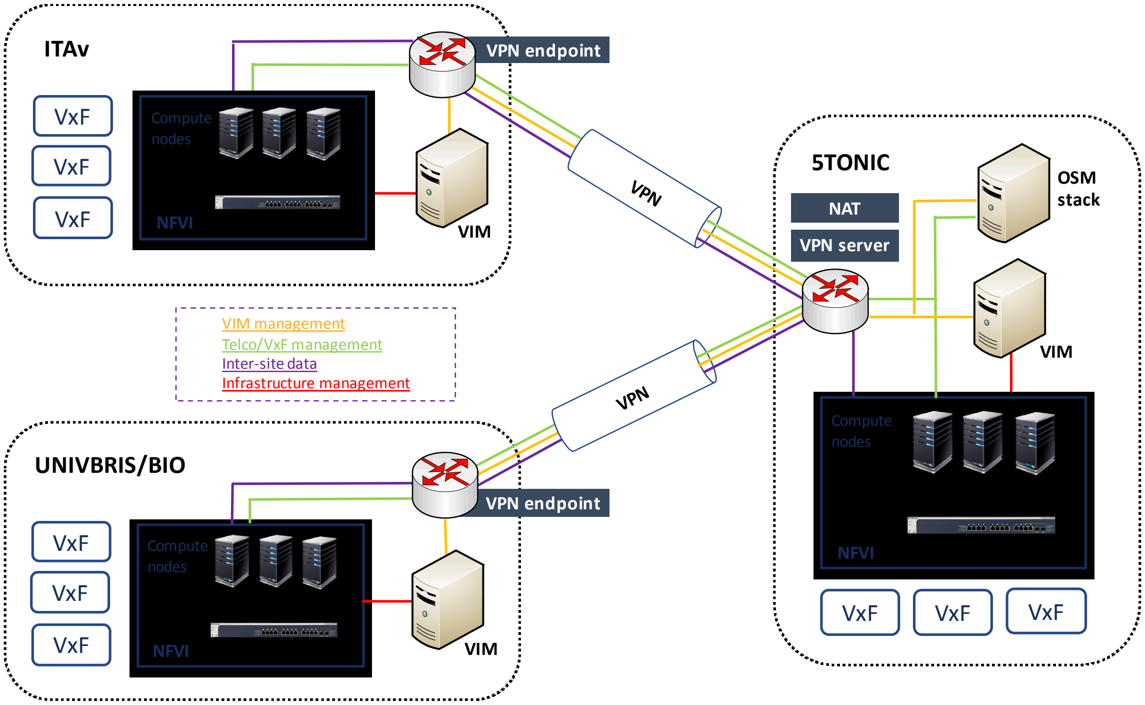

Figure 7: Approach for inter-site connectivity

In this approach, inter-site control-plane communications are enabled by a VPN service that, as agreed by 5GinFIRE partners, has been deployed at the 5TONIC site, the 5GinFIRE testbed provider that hosts the OSM stack. The VPN server of 5TONIC offers authorized partners (i.e. 5GinFIRE testbed providers) a secure access with certificate-based authentication to 5TONIC premises, enabling the exchange of control information between the OSM stack and the VIMs/SDN controllers at the remote sites (represented with a yellow line in Figure 7). Additionally, it enables the communications needed by the OSM stack to carry out the configuration of the VxFs after their deployment (shown in green colour in the figure). This way, inter-site control plane communications follow a hub-and-spoke distribution model, where information is distributed using a star topology centred at 5TONIC, where the OSM stack is installed.

The utilization of a VPN service provides a suitable solution to support the flexible incorporation of sites to the 5GinFIRE MANO platform, allowing the establishment of multiple protected access links between the VPN server, at the 5TONIC site, and each 5GinFIRE testbed provider, which may deploy several VPN endpoints. This enables the project to accommodate particular requirements that can be presented by interconnecting entities, for example using different VPN endpoints to support control and data communications. This is shown in Figure 8, which includes an example where two sites with different requirements are connected to the VPN service hosted by 5TONIC. In this example, *site 2* shares a VPN connection for inter-site control and data-plane communications, while *site 1* uses two independent VPN endpoints for control and data information.

In addition, the adopted approach enables 5GinFIRE testbed providers to retain the control over their networks. That is, a provider may reasonably determine the number of needed VPN endpoints (subject to approval within 5GinFIRE), and their location inside the provider site. Moreover, providers can set up the most appropriate mechanisms, according to their internal network configuration, for the exchange of content from the VPN endpoints at their sites towards the corresponding traffic destination (VIM equipment, SDN controllers and VxFs), and vice versa. The 5GinFIRE testbed providers can also configure any VIM networks that are suitable to enable the inter-site exchange of control and data information by VxFs deployed at their infrastructures. The VPN server and the VPN endpoints behave as layer-3 routers, forwarding incoming traffic (i.e. traffic that arrives from the VPN) into their corresponding site, and outgoing traffic (i.e. traffic originating from the site) through the appropriate VPN connection. This way, inter-site control traffic originated at the OSM stack is distributed through the internal network of 5TONIC, being encapsulated by the VPN server and routed towards the corresponding VPN endpoint at its destination site. The VPN endpoint then forwards the traffic, which is transmitted through the internal network of the testbed provider towards its corresponding destination equipment or VxF (in the latter case, traffic will be delivered to the VxF through a VIM network). Traffic in the reverse direction would be delivered following an analogous procedure.

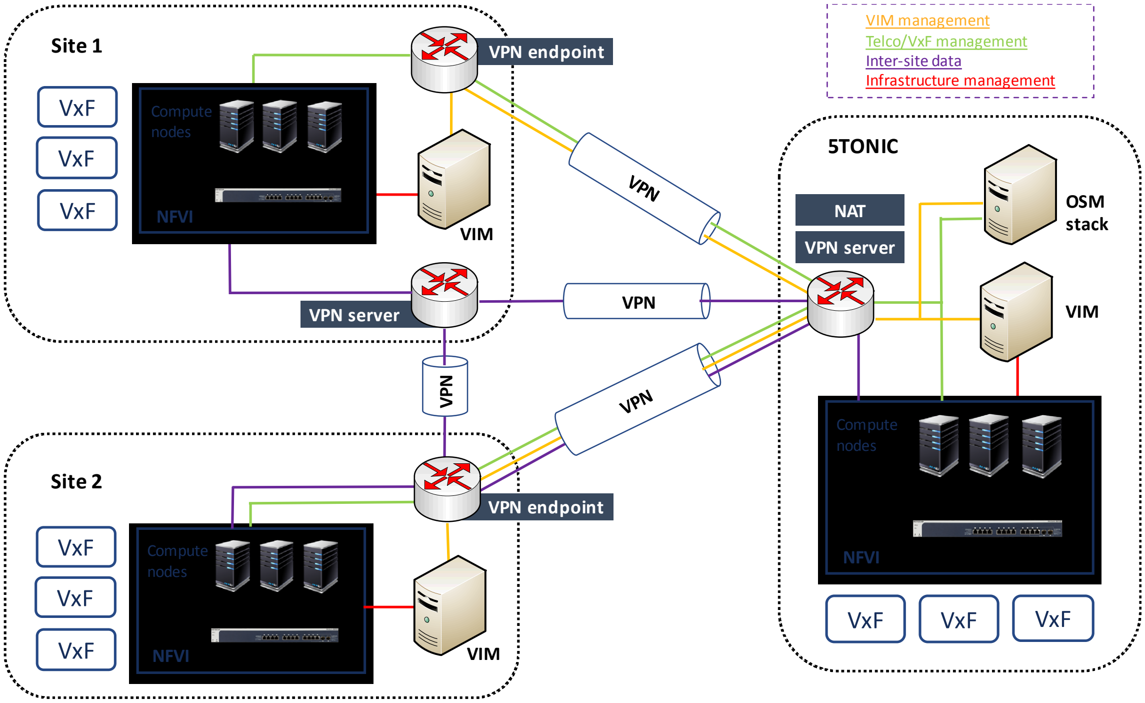

Figure 8: Examples of inter-site connectivity

With respect to the exchange of data among VxFs deployed at different sites, a first approach could be re-utilizing the aforementioned hub-and-spoke distribution scheme. In this case, the VPN server would act as a traffic relay among datacentres. However, direct site-to-site communications could be of interest in some experimentation scenarios, to avoid costly suboptimal network paths and bottlenecks close to the 5TONIC network location. This may require the installation of a VPN service in the infrastructures of other 5GinFIRE testbed providers. This is shown in Figure 8, where site 1 deploys a VPN server, enabling the establishment of direct site-to-site communications with site 2, to support inter-site data exchange among VxFs. Inter-site data-plane communications will require the creation of specific VIM networks at each site, with network connectivity towards the appropriate VPN endpoint of the site. These networks will then be used to attach any VxFs needing this type of communications.

3.2.2 Addressing plan

Enabling effective network communications among multiple sites requires a careful design of the IP address space to be used by the 5GinFIRE testbed providers. In particular, the address space utilized for control and data-plane communications should not collide with the address space already in use by any of the sites for other purposes.

To satisfy this criterion, the following agreements have been taken by 5GinFIRE regarding the IP address space:

1) 5G testbed providers will use the private address space 10.154.0.0/16 for control and data plane communications, i.e. an address range not in use at the sites that provide the first stable version of the 5GinFIRE MANO platform.
2) To simplify routing configurations inside 5TONIC, this specific location will use the private address space 10.4.0.0/16 to support control and data plane communications.

The 5GinFIRE network operations center will be in charge of the allocation of IP address ranges to entities within the address space 10.154.0.0/16. The current allocation of IP addresses to sites is shown in Table 1. In addition, the VPN service at 5TONIC has been configured to use subnetworks 10.154.255.0/24 and 10.154.254.0/24, being the latter the address range that will be allocated to VPN endpoints.

Table 1: Current allocation of IP addresses to sites

Site	IP address range
5TONIC	10.4.0.0/16
ITAv	10.154.0.0/20
UNIVBRIS&BIO	10.154.16.0/20

3.2.3 Requirements to support the interconnection of external sites

Beyond any particular requirements that may be identified when facing the interconnection of additional external sites, which will be treated on a case-by-case basis, in the following, and according to the information indicated in this section, we indicate a non-exhaustive list of requirements that must be fulfilled by external entities to connect to the 5GinFIRE MANO platform:

1) Installation and configuration of the VPN endpoints that are necessary to integrate the entity's infrastructure into the network overlay architecture of the 5GinFIRE MANO platform.
2) Utilization of a VIM solution compliant with the 5GiFIRE MANO platform (see section 2.2).
3) Configuration of the appropriate VIM networks that enable the exchange of control and data information originated and terminated at VxFs deployed at the entity's datacentre.
4) Set up of the appropriate mechanisms to support the delivery of control and data information across the local network segments of the external entity (i.e. from the VPN endpoints towards the VIM/SDN controller and VxFs, and vice versa).
5) Utilization of an appropriate IP address space, not conflicting with the address space assigned to the 5GinFIRE testbed providers so far. Interconnecting entities must use a range of IP addresses within the network prefix 10.154.0.0/16. This range will be determined by the 5GinFIRE network operations centre, according to existing allocations and the entity's needs.

3.3 Support of SDN

5GinFIRE is also exploring the utilization of SDN technologies to support the configuration of data-plane communications within the scope of a site. The integration of SDN-based control into an experimental infrastructure can be made following two different approaches:

1) Integration of the SDN controller with the VIM.
2) Direct integration with OSM, using the SDN assist capability of this orchestration platform.

A first tentative to integrate with SDN controller was performed in the context of OpenStack Mitaka, which was the VIM version chosen by UNIVBRIS&BIO in the beginning of the project. However due to Mitaka packstacks automated configuration this approach was unsuccessful, since packstack configured Neutron to use the linux bridge agent instead of the Openvswitch agent. In addition, the packstack node management became inaccessible once we had configured the Openvswitch agent and connected Openvswitch with the OpenDaylight SDN controller. Then it was decided to move to OpenStack Ocata. However, so far it has not been possible to perform the integration with Openstack Ocata.

The second option that was explored consisted of integrating a SDN controller with the OSM stack, using the SDN assist feature provided by this platform[4]. In this case, the goal is to integrate a SDN controller via the Resource Orchestrator (RO) of OSM, without directly connect to the VIM. The main components required for this integration include:

1) An external SDN Controller (e.g., OpenDayLight).
2) A dataplane switch with OpenFlow.
3) The mapping between the switch ports and the compute node interfaces.
4) Admin credentials if the VIM is OpenStack.

[4] Configure VIM SDN, OSM Wiki (last access: December 2017): https://osm.etsi.org/wikipub/index.php/Configure_VIM_SDN

BIO experimented with associating an Openflow enabled NEC PF5459 Switch with OSM MANO. An Openflow instance was setup on the switch and linked with an Opendaylight SDN controller. The DPID of the Openflow instance on the switch was provided to MANO and the SDN controller associated with the datacentre (VIM) It was discovered that the NEC PF5459s only partially support Openflow and key features are missing.

Another experiment was also performed, an Opendaylight SDN controller was associated with the Openvswitch of a single node OpenStack environment. This caused BIO to lose access to the VIMs management. At the time of writing, activities regarding the integration of SDN with OSM are still being conducted.

3.4 Provision of access to experimenters

Apart from the specific internal objectives 5GinFIRE partners have for the deployed NFV experimental infrastructure, a key aspect for the project is the support of external experiments via the Open Calls mechanism. External experimenters will be able to execute their experiments remotely on top of the 5GinFIRE infrastructure, which imposes a new connectivity requirement: experimenters must be able to access their resources under experimentation, for monitoring, reporting, etc., independently of where (i.e. on which site) their VxFs are actually deployed.

There exist several technical solutions to achieve such management connection, so the project has focused the analysis on three key parameters:

1) Availability, or easiness to replicate the solution across the different sites involved.
2) Operation simplicity.
3) Isolation degree among the different deployed elements.

Initially, the project considered the logical networks that were already functional and reaching deployed VxFs. At the sight of Figure 7, there were three of them, but only two spanned across multiple sites: the data-plane and the VxF management interconnections. Both are mandatory for any new site incorporating to 5GinFIRE, so reusing them for external management would ensure the universality of the solution. Between the two, the VxF management network was the first candidate, to avoid data-plane congestions and because it is less prone to configuration errors.

Simplicity is granted in this solution, as well, if centralized access to this management network is considered. In particular, the same certificate-based VPN mechanism provided by 5TONIC to 5GinFIRE testbed providers for inter-site control-plane communications will be replicated for external experimenters, so no additional complexity will be added.

Finally, from the protection point of view, having a single network shared by all external experimenters means that some control mechanisms need to be implemented to account for elements not being physically/logically isolated. In that sense:

- VxFs from different experimenters will not be isolated in the management plane, so it will be a requirement for experimenters not to configure trivial credentials (e.g. admin/admin).
- Certain access rules will be required to isolate 5GinFIRE control elements (e.g. OSM), blocking undesired flows and permitting only those which are strictly required.

These security policies will be implemented at a “jump machine”, controlled by 5GinFIRE, and which will be the only one permitted by the VPN access rules, as depicted in Figure 9 below. Any new experimenter (blue line) will be connected to the VPN server, and will be granted access to the “jump machine” only. ACLs will then be implemented in the green interface of this node, controlling incoming connections to the VxF management network, and ensuring security for the orchestration modules.

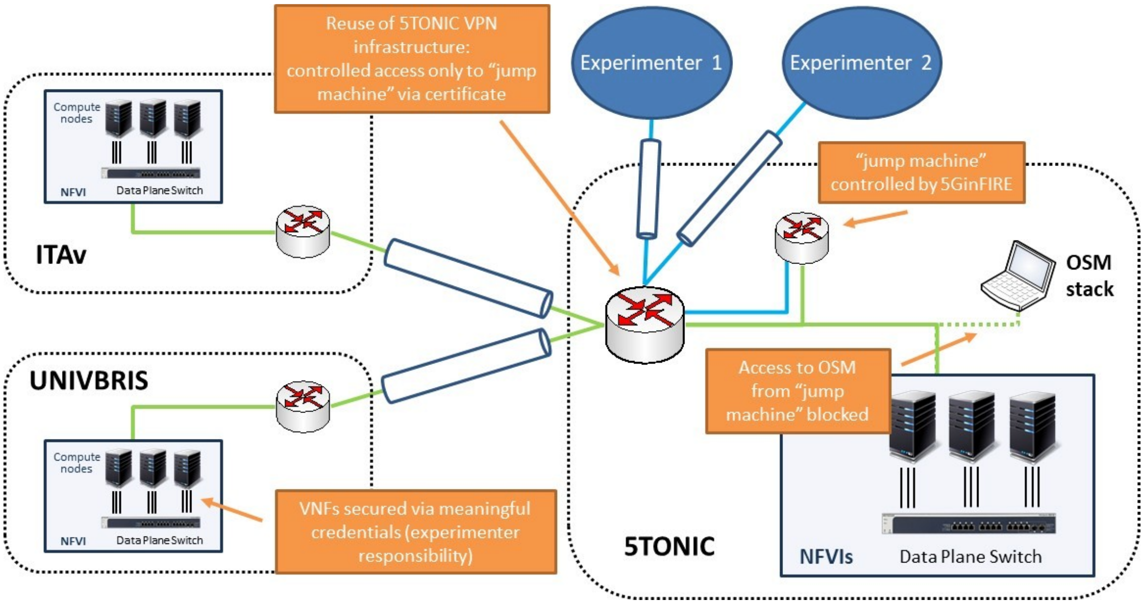
Reuse of 5TONIC VPN infrastructure: controlled access only to “jump machine” via certificate
Experimenter 1
Experimenter 2
Compute nodes
NFVI
Data Plane Switch
ITAv
“jump machine” controlled by 5GinFIRE
OSM stack
UNIVBRIS
Access to OSM from “jump machine” blocked
Compute nodes
NFVI
Data Plane Switch
VNFs secured via meaningful credentials (experimenter responsibility)
5TONIC
NFVIs
Data Plane Switch

Figure 9: External access for experimenters

3.5 Mechanisms for evolution and continuous integration

One of the most salient characteristics of current software development practices is the support for continuous integration techniques, allowing a seamless evolution of the operational environments as their software bases evolves. 5GinFIRE testbed providers will take care of maintaining this characteristic for the 5GinFIRE MANO platform and, given the extremely active nature of the foreseen MANO software evolution, proactively incorporate any breakthrough features that are incorporated in the platform software base and that require a platform update beyond continuous integration mechanisms. These updates will be planned and executed at each site without impacting the operational availability of the MANO framework as a whole. In the following, we describe the mechanisms for evolution and continuous integration that will be independently set up by 5GinFIRE testbed providers.

UC3M will maintain an independent small-scale MANO deployment, to support testing of new releases of OSM and OpenStack, which provide the basis for the MANO framework deployed at 5TONIC. This small-scale deployment will also support experimentation with other OSM components and plugins, e.g., to test new types of VIMs, SDN controllers and public cloud support, as well as the validation and debugging of new components and extensions to the 5GinFIRE MANO platform. With this purpose, the small-scale MANO platform at UC3M will have a functional OSM stack and two VIM instances based on OpenStack. The OSM stack and the two VIMs will each run on a mini-ITX computer (Intel Core i7 2.3 GHz, 8 GB RAM, 128 GB SSD, 4 GbE ports with DPDK capabilities).

One of the VIM instances will manage a NFVI consisting of 4 compute nodes, each consisting of a mini-ITX computer with the same hardware characteristics as the aforementioned equipment. To support future research and experimentation activities of UC3M, related to virtualization over single board computers, the second VIM instance will support the interaction with the OSM stack to enable the deployment of lightweight VNFs using containers over an NFVI conformed by three Raspberry Pi 3 Model B (this corresponds to a work in progress at UC3M, see section 5.3). The small-scale MANO deployment of UC3M is schematized in Figure 10.

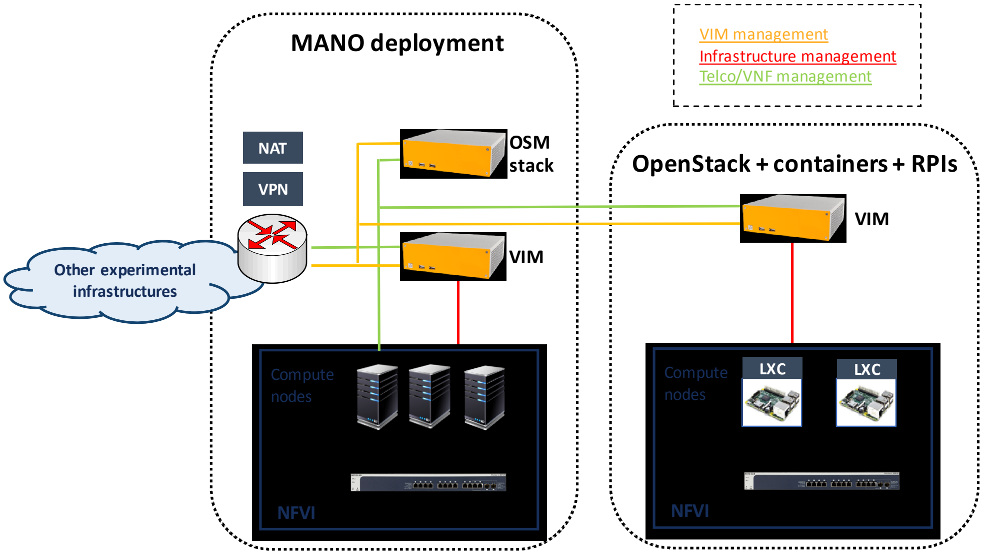

Figure 10: Small-scale deployment at UC3M

ITAv will maintain an independent small-scale MANO deployment, to support testing of new releases of OSM and OpenStack, which provide the basis for the MANO framework deployed at ITAv. This small-scale deployment will also support experimentation with other OSM components and plugins. It will also help validate new components developed and new hardware/software integration. In this small-scale deployment, it will be possible to test the project's AAA component under development (see section 5.2) without causing testbed downtime. It will also allow to test a third-party platform/solution to monitor VNFs and their configuration process. This software will help us in the future to determine the reasons why a VNF failed or why the configuration was not successful.

UNIVBRIS/BIO will maintain a NFVI, which is SDN enabled providing to the experimenters the proper platform for VxF experimentations. The number of compute nodes will be increased according to the experiment demands. Mobile Edge Compute (MEC) nodes will be available over the BIO infrastructures to support the Smart City Safety scenario mainly to include communication with Raspberry PIs Model B. The MEC nodes will enable to move the VxFs from the Cloud to the Edge

Future evolution of the UNIVBRIS/BIO platform will focus on the integration of OpenStack with ODL, as well as with OSM MANO. For this end, all the procedures are going to be performed in a separated infrastructure, called BIO Staging, before being deployed at production environment. BIO Staging is a controlled environment inside of the laboratory at UNIVBRIS that enables different kind of configurations and testing to guarantee the proper operation of the extension being experimented. Once everything is operated adequately the next step is to perform the migration to the production environment.

4 Functional validation

In this section, the different functional tests and experiments that have been applied to verify the availability of the 5GINFIRE MANO platform are introduced and briefly discussed.

4.1 Tested and validated NSs and VNFs

The importance of documenting the processes for testing individual VNF on-boarding and activations, and NS deployment, was evident from the very moment the first tests with the orchestration element were run. To that end, we created documents that describe the VNFs, NSs and tests. It also stores the relations between VNFs and NSs, which tests were performed and results yielded.

The VNFs and NSs tested are the reference ones provided by OSM (more details about the reference VNFs and NSs, besides what is provided in this section, can be found in [9][5]), to validate them for further platform testing. All the tests were done using the small-scale MANO deployment at ITAv.

The tests designed to be performed on VNFs are shown in Table 2.

Table 2: VNFs tests

Test ID	Name	Version	Description	Date
1	Onboarding VNF	1.0	Onboarding VNF into the OSM	01/09/2017
2	Instantiating VNF	1.0	Instantiating VNF	01/09/2017
3	Configuring VNF	1.0	Configuring VNF	01/09/2017

The tests designed to be performed on the NSs are shown in Table 3.

Table 3: NSs tests

Test ID	Name	Version	Description	Date
1	Onboarding NS	1.0	Onboarding NS into the OSM	01/09/2017
2	Instantiating NS	1.0	Instantiating NS	01/09/2017

[5] Besides this, we used the Ping Pong network service, a well-known NS for testing and debugging purposes. It is available at (last access: December 2017): https://osm-download.etsi.org/ftp/old-examples/ping_pong_ns/. Additionally, the reference NS cirros_2vnf_ns is available at (last access: December 2017): https://osm-download.etsi.org/ftp/old-examples/cirros_2vnf_ns/

Test ID	Name	Version	Description	Date
3	Configuring NS	1.0	Configuring NS - Configure VNFFG	01/09/2017

The VNFs tested are indicated in Table 4.

Table 4: VNFs tested

VNF ID	Name	Short Name	Vendor	Version	Image	Date
1	cirros_vnfd	cirros_vnfd	OSM	1.0	cirros034	01/09/2017
2	Ref_VNF_11	Ref_VNF_11	ETSI	1.0	cirros034	01/09/2017
3	Ref_Vnf_21	Ref_Vnf_21	ETSI	1.0	cirros034	01/09/2017
4	ping_vnf	ping_vnf	RIFT.io	1.1	Fedora-x86_64-20-20131211.1-sda-ping.qcow2	01/09/2017
5	pong_vnf	pong_vnf	RIFT.io	1.1	Fedora-x86_64-20-20131211.1-sda-pong.qcow2	01/09/2017

The NSs tested are described in Table 5.

Table 5: NSs tested

NS ID	Name	Short Name	Vendor	Version	Date
1	cirros_2vnf_nsd	cirros_2vnf_nsd	OSM	1.0	01/09/2017
2	Ref_NS_1	N/D	ETSI	N/D	01/09/2017
3	ping_pong_ns	ping_pong_ns	RIFT.io	1.1	01/09/2017

The composition of the NSs is indicated in Table 6.

Table 6: NSs composition

NS ID	NS Name	VNF ID	VNF Name
1	cirros_2vnf_ns	1	cirros_vnfd

NS ID	NS Name	VNF ID	VNF Name
	d		
2	Ref_NS_1	2	Ref_VNF_11
2	Ref_NS_1	3	Ref_Vnf_21
3	ping_pong_ns	4	ping_vnf
3	ping_pong_ns	5	pong_vnf

The VNF test results are illustrated in Table 7.

Table 7: VNFs testing results

VNF Test Run ID	VNF Test ID	VNF ID	Date	Result	Comments
1	1	1	01/09/2017	Success	N/D
2	2	1	01/09/2017	Success	N/D
3	1	2	01/09/2017	Success	N/D
4	2	2	01/09/2017	Success	N/D
5	1	3	01/09/2017	Success	N/D
6	2	3	01/09/2017	Success	N/D
7	1	4	01/09/2017	Success	N/D
8	2	4	01/09/2017	Success	N/D
9	3	4	01/09/2017	Success	Sometimes VCA component freezes in blocked state
10	1	5	01/09/2017	Success	N/D
11	2	5	01/09/2017	Success	N/D
12	3	5	01/09/2017	Success	Sometimes VCA component freezes in blocked state

The NS test results are indicated in Table 8.

Table 8: NSs testing results

NS Test Run ID	NS Test ID	NS ID	Date	Result	Comments
1	1	1	01/09/2017	Success	N/D
2	2	1	01/09/2017	Success	N/D
3	1	2	01/09/2017	Success	N/D
4	2	2	01/09/2017	Success	N/D
5	1	3	01/09/2017	Success	N/D
6	2	3	01/09/2017	Success	N/D

Some VNFs do not require VNF configuration, so that test was not run for them. The VNFs that do not require VNF configuration are the following: cirros_vnf, ref_vnf_11 and ref_vnf_21.

The NSs tested did not required NS configuration, so the test was not run.

The next VNFs and NSs to be tested will explore service chaining (VNFFG), PCI-Express pass-through, and SR-IOV (technologies that were not tested in previous VNFs and NSs).

4.2 Support of inter-site communications

This subsection overviews the functional tests and experiments that have been carried out to validate the overlay network architecture solution adopted by the project, to support inter-site control and data plane communications.

To do these tests, an experimental infrastructure has been set up at UC3M, consisting of: (a) an access gateway, (b) a VIM based on OpenStack Ocata and (c) an NFVI with a single compute node. Each of these components has been provided as a mini-ITX computer platform (Intel Core i7 2.3 GHz, 8 GB RAM, 128 GB SSD, 4 GbE ports with DPDK capabilities). The access gateway was used to deploy a VPN endpoint of the overlay network architecture of 5GinFIRE (see Section 3.2), supporting this way the exchange of inter-site control and data-plane information with the experimental infrastructure deployed at 5TONIC.

The functional tests described in Table 9 have been performed. In a first test (test 1), the OSM stack at 5TONIC was used to deploy a reference NS at the external site provided by UC3M. The NS was composed of two interconnected VNFs (see Figure 11), which did not require additional configuration through the Juju interface of the OSM stack[6]. Upon the successful deployment of the NS, both VNFs were completely functional and capable of exchanging data using the virtual data link that was configured between them. This test has

[6] More details about this reference NS can be found in (last access: December 2017): https://osm-download.etsi.org/ftp/old-examples/cirros_ns/

served to validate the capacity of the overlay network architecture of 5GinFIRE to support control plane communications between the OSM stack and an external VIM, to coordinate the allocation and configuration of computing, storage and network resources at an external site.

Table 9: functional tests to validate the overlay network architecture of 5GinFIRE

Test ID	Description	Result
1	Deployment of a reference NS in the external site configured at UC3M, not requiring day-1 configuration of VNFs	Success
2	Deployment of a reference NS in the external site configured at UC3M, requiring day-1 configuration of VNFs	Success
3	Deployment of a reference NS using the multi-site capabilities of the 5GinFIRE MANO stack	Success

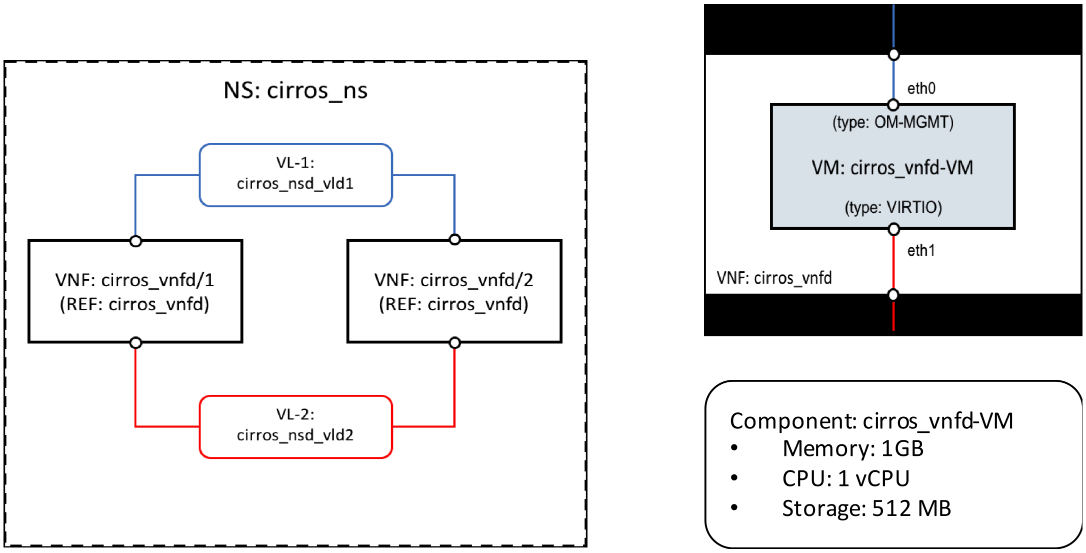

Figure 11: Reference NS #1 used to validate inter-site communications

In a second test (test 2), the OSM stack of 5GinFIRE was utilized to instantiate a different reference NS at the UC3M experimental infrastructure. In this case, the NS consisted of two interconnected VNFs[7] that required day-1 configuration operations via Juju. The NS is represented in Figure 12. A VIM network was pre-created at the UC3M NFVI, to support the

[7] As already commented, the Ping Pong network service is available at (last access: December 2017): https://osm-download.etsi.org/ftp/old-examples/ping_pong_ns/

IP connectivity of the VNFs with the VPN endpoint of UC3M (this network was also created for test 1, though it was not required the configuration of the VNFs). After deploying the virtual machines of the NS, both VNFs were successfully instantiated and configured through their corresponding Juju charms, and the NS was completely functional, being both VNFs capable of exchanging data through a virtual data link. This experiment has served to validate the capacity of the overlay network architecture of 5GinFIRE to support inter-site control communications related to the lifecycle management of VNFs.

Finally, in a third test (test 3), the OSM stack hosted at 5TONIC was used to deploy the same network service as in the previous test, although leveraging the multi-site support of the 5GinFIRE MANO platform. In this case, the OSM stack was instructed to deploy one of the VNFs at 5TONIC, while the other was instantiated at UC3M. Besides the VIM networks that are necessary to support control plane communications, this experiment also required the creation of an additional VIM network at 5TONIC, to provide the VNF with IP connectivity to the VPN server of the site; analogously, an additional VIM network was pre-created at UC3M, providing the VNF with IP connectivity to the VPN endpoint. After the successful deployment of the NS, both VNFs were able to exchange data through the VPN-based overlay network of 5GnFIRE, validating the suitability of the adopted solution to support the exchange of inter-site data.

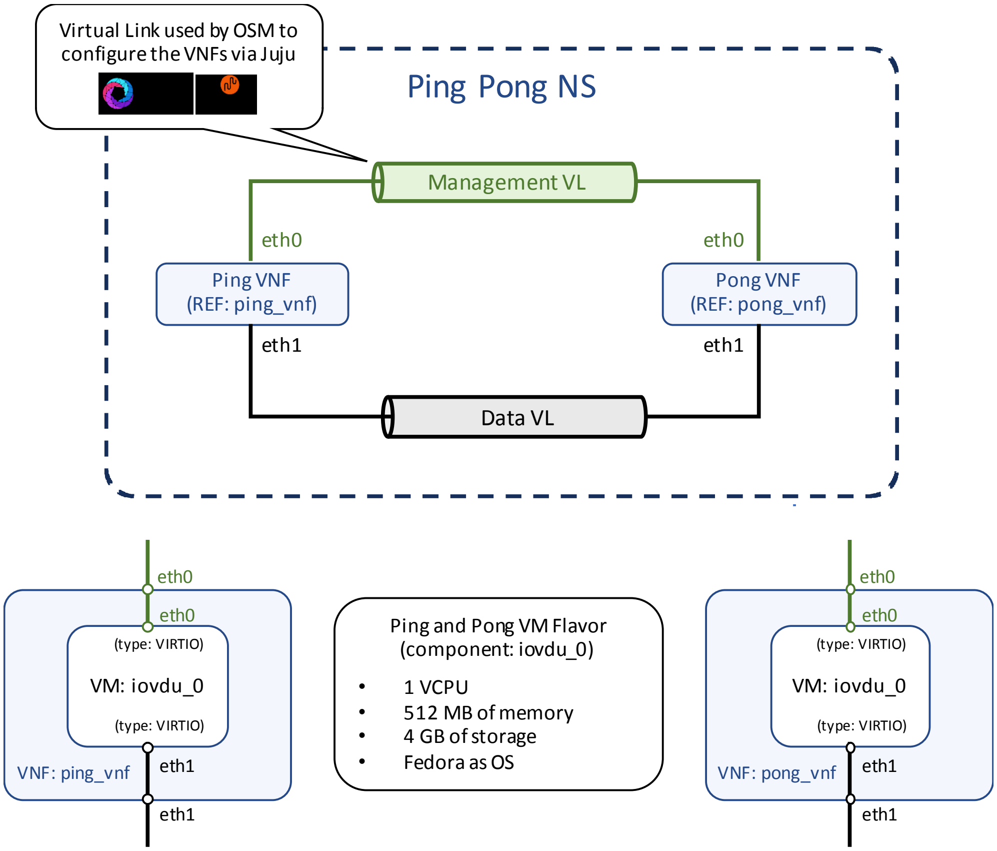

Figure 12: Reference NS #2 used to validate inter-site communications

5 Enhancements to the orchestration platform

In order to properly address all the requirements for the orchestration platform of the 5GINFIRE environment, the team has started to work on several enhancements to the current platform elements. This section describes the already initiated work on these enhancements, which have a different degree of maturity. All of them will be contributed (actually, the one on Ansible has already been) to the relevant upstream communities.

5.1 Integration of Ansible into the Juju framework of OSM

As we already commented, the VCA module is the OSM component in charge of the configuration of VNFs, generally aligned with the VNF manager entity of the ETSI NFV reference architectural framework. It presents an interface to Juju, allowing the configuration of VNFs through the execution of a limited form of Juju charms, called VNF Configuration charms or proxy charms[8].

The configuration of VNFs via proxy charms is summarized in Figure 13 where, after the creation of a virtual machine by the RO, as instructed by the SO (step 1 in Figure 13), the SO deploys the proxy charm using a LXD container (step 2 in Figure 13). The configuration of the VNF is done with the invocation of Juju actions by the SO, which cause the proxy charm to configure the VNF (step 3 in Figure 13).

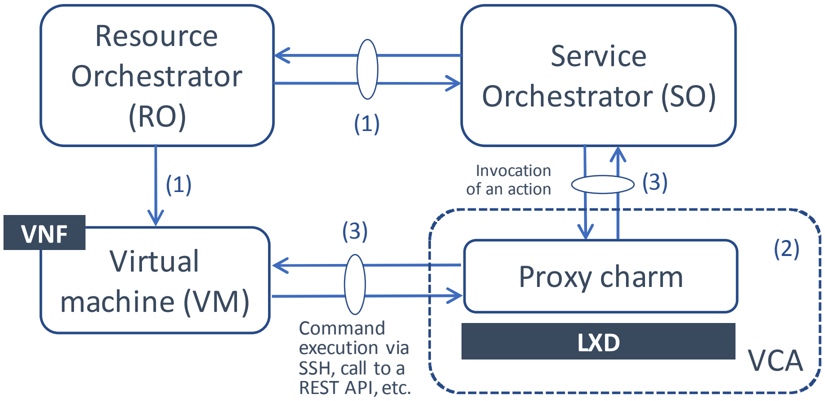
Resource Orchestrator (RO)
Service Orchestrator (SO)
(1)
(1)
Invocation of an action
(3)
VNF
Virtual machine (VM)
(3)
Proxy charm
(2)
LXD
VCA
Command execution via SSH, call to a REST API, etc.

Figure 13: VNF configuration using Juju

Aiming at increasing the range of configuration options available to VNF developers in OSM, as well as facilitating the portability of existing VNF developments, 5GinFIRE has explored the utilization of other well-know and wide-used mechanisms to support the configuration of VNFs, identifying Ansible [10] as a technology of particular interest.

Under the aforementioned considerations, we have developed a base charm layer, *ansible-charm*, that allows the configuration of a VNF using an Ansible playbook. The base charm layer includes the Juju base layers vnfproxy[9] and ansible-base[10], and provides a template

[8] Creating your own VNF charm (Release TWO), OSM Wiki (last access: December 2017): https://osm.etsi.org/wikipub/index.php/Creating_your_own_VNF_charm_(Release_TWO).

[9] vnfproxy, Juju charm layer (last access: December 2017): https://github.com/AdamIsrael/vnfproxy

[10] Ansible Base Layer for Charms (last access: December 2017): https://github.com/chuckbutler/ansible-base

ready for customization that allows creating a proxy charm that supports the execution of an Ansible playbook using the Juju framework of OSM. The ansible playbook to be run is provided as a file under the directory structure of the base charm laYer *ansible-charm*. This directory structure is shown in Figure 14.

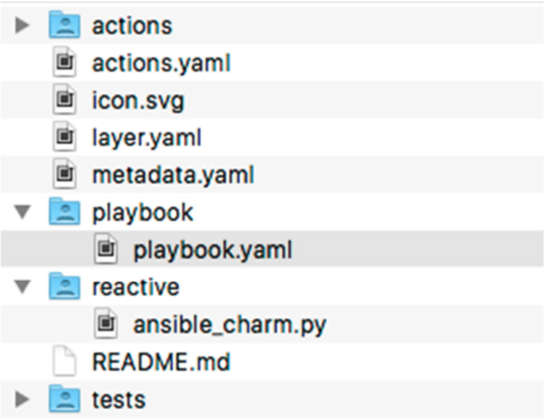

```
▶ actions
  actions.yaml
  icon.svg
  layer.yaml
  metadata.yaml
▼ playbook
    playbook.yaml
▼ reactive
    ansible_charm.py
  README.md
▶ tests
```

Figure 14: File structure of the base charm layer *ansible-charm*

In the following, we summarize the step-by-step instructions that can be followed to use the base charm layer:

1) Include the playbook to be run under the playbook folder of the base charm layer, naming it as *playbook.yaml* (this file is highlighted in Figure 14). Alternatively, the VNF developer can open the file that already exists in this directory and paste the desired playbook in this file.
2) The base charm layer already implements a Juju action, *ansible-playbook*, within the *reactive* folder. By default, this action runs the playbook *playbook/playbook.yaml*. Additional actions can be defined, if needed by the VNF developer.
3) Build the charm, via the *charm build* command.
4) Update the VNF descriptor (VNFD) to use the charm. In particular: (a) specify the name of the Juju charm in the VNF configuration; (b) include the action *ansible-playbook* with no arguments as a service primitive and as an initial configuration primitive.
5) Include the compiled charm in the VNF package.

The development of the base charm layer *ansible-charm* has been contributed to the OSM community[11]. Our future work includes the maintenance and consolidation of this contribution across the new releases of OSM.

5.2 Security framework: applicability of Keystone

5.2.1 Introduction

To enhance the 5GinFIRE MANO platform, we want to create a component for OSM that is responsible for authentication, authorization and accounting. This component will help to

[11] Example VNF Charms (last access: Dec. 2017) https://osm.etsi.org/wikipub/index.php/Example_VNF_Charms

manage the MANO stack, by providing an easier way to manage users and roles and account for every action in the system. While integrating this AAA component into OSM, we also want to add a new feature, specifically to allow project isolation.

The system uses Openstack Keystone, so it can handle user, role and project management, thus allowing the reuse of a widely used component that has been thoroughly tested. Keystone is a project within OpenStack, providing support for the management of users, roles and projects. It also adds the benefit of integrating with other user/organization management systems, like LDAP, Kerberos or Microsoft's Active Directory. Using Keystone also allows addressing authentication and authorization matters, since it is a source of user and project information. For accounting, this project will use logging. This approach can evolve to a SQL database if need be.

5.2.2 Requirements

The main drivers for the creation of this component are the aim to centralize user access control to OSM, and account for every action made by the users. The requirements are specified in the following list:

- Centralize user access control
- Centralize user access roles
- Provide user access control to 5GinFIRE portal
- Provide user access control to OSM
- Reuse the Keystone component from OpenStack
- User management
- Role management
- Project management
- Project isolation
- Action accounting
- Non-repudiation of actions

The centralisation of user access control is necessary to provide a single identity provider that makes user management simpler and easier, removing the burden of managing multiple accounts for the same entity.

The centralisation of user access roles increases security, allowing both the 5GinFIRE portal and OSM to have access to the same information regarding an entity, enabling better decisions, avoiding duplication of information about an entity and disallowing conflicting roles.

We propose the reuse of Keystone, a component already existing in the architecture of OpenStack. Keystone allows the use of other technologies widely used to store user information, such as LDAP, Kerberos or Microsoft's Active Directory.

User and role management is another requirement to make possible the creation, retrieval, update and deletion of that information.

Given its original inception to solve current operational problems and grow once a stable version was available, OSM lacks the notion of users and projects, and this is important for entity isolation in a multi-user and multi-project environment like 5GINFIR. Isolation is needed when multiple users are working in the system, and one's work must not interfere with the next one's. This isolation should allow users to work on the same system without affecting one another and can be achieved using the notion of projects and project management at the level of OSM.

Action accounting is important to understand which user did what. In the future, it might be also beneficial for charging purposes.

Non-repudiation of actions is an additional important security aspect that will allow verifying the whole process and its integrity, verifying that no malicious action was taken against the orchestration framework.

5.2.3 Authentication

Authentication is the act of verifying the user's identity. In our case, the user must provide a username and a password. If the credentials match any of the stored credentials, the user will be granted an authorization token, that identifies her/him in the system.

The authentication provided by Keystone will be used here, since it already supports a myriad of backends, like LDAP, Kerberos or Microsoft's Active Directory.

The AAA layer also provides an external API to be used by the 5GinFIRE portal, as mentioned in deliverable D3.1, allowing for centralized access control between the two systems.

5.2.4 Authorization

Authorization is the act of verifying if the user has permissions to do the requested operation. The AAA component bases itself in Keystone that uses an RBAC model, where a user has roles in projects, and according to those roles, he has certain permissions.

5.2.5 Roles

The roles to be used are presented in Table 10.

Table 10: Roles to be used

No	Role	Role description
1	Owner	Is attributed to each user that creates an experience in that project. An owner can manage its project
2	Administrator	Is attributed to each administrator so they can manage all the projects

These roles help isolate projects by creating clear boundaries between them, and between the users responsible for them. Each user can only access projects where (s)he has a role.

The Administrator role is a special role reserved for the users that manage the MANO framework and the infrastructure. These users are responsible for managing the projects and associating VIMs to them, meaning that a VIM is not available to all projects by default. It is relevant also for isolation, where each project can have its isolated VIM, or the project can share the VIM with other projects.

5.2.6 Accounting

Accounting is the act of recording all the interactions of the users and components of the system. It is important to account for all the actions made by the users because this information is extremely important in a multitude of situations, like for example:

- Audit the system
- Debugging the system
- Charging actions

For this information to be trustworthy, there is a need for a non-repudiation mechanism to guarantee that the requests and responses can have their origin verified.

5.2.7 Processes

To illustrate how the concepts presented before can integrate with OSM processes, we present the following two subsections, where authentication, authorization and accounting examples can be found.

5.2.7.1 User authentication

In the user authentication process, there are two steps: authentication and accounting (see Figure 15).

Authentication is the process by which users are granted access based on the credentials they provide. This process will be recorded by Accounting, allowing for security audits. Accounting will record which user was used, the result of the operation and the time it occurred.

From this process, there are two possible outcomes: success or error. In case of success, the user is logged in, and OSM provides her/him with an authentication token. In case of error, the user is prompted to try again, because the credentials are wrong or an error in the system occurred.

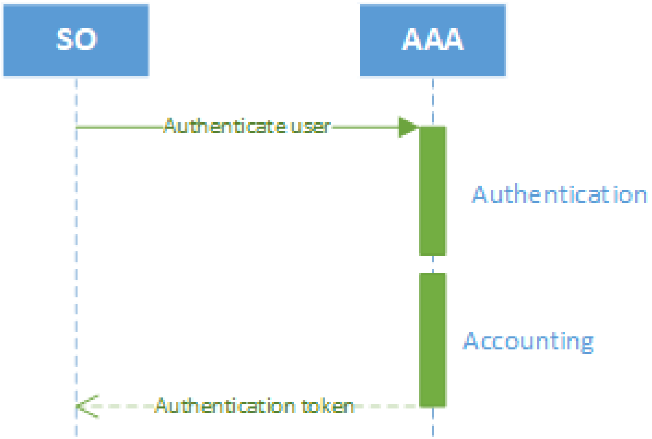

Figure 15: User authentication process flow

5.2.7.2 Virtual machine creation

In the virtual machine creation process, there are two steps: authorization and accounting (see Figure 16).

Authorization verifies if the user has permissions to do the requested operation, in this case, create a virtual machine.

Accounting records that the user tried to do an operation and how the operation impacted the whole system.

In this process, we can check that the operation is verified by both the component that requested it and the component that fulfils it. These steps increase security by verifying that the operation is valid and that the user issuing it has the right permissions to do so.

After the operation is complete, the state of the system must be accounted for and recorded by the AAA system; this allows to check for errors, record resource usage and user interactions with the infrastructure.

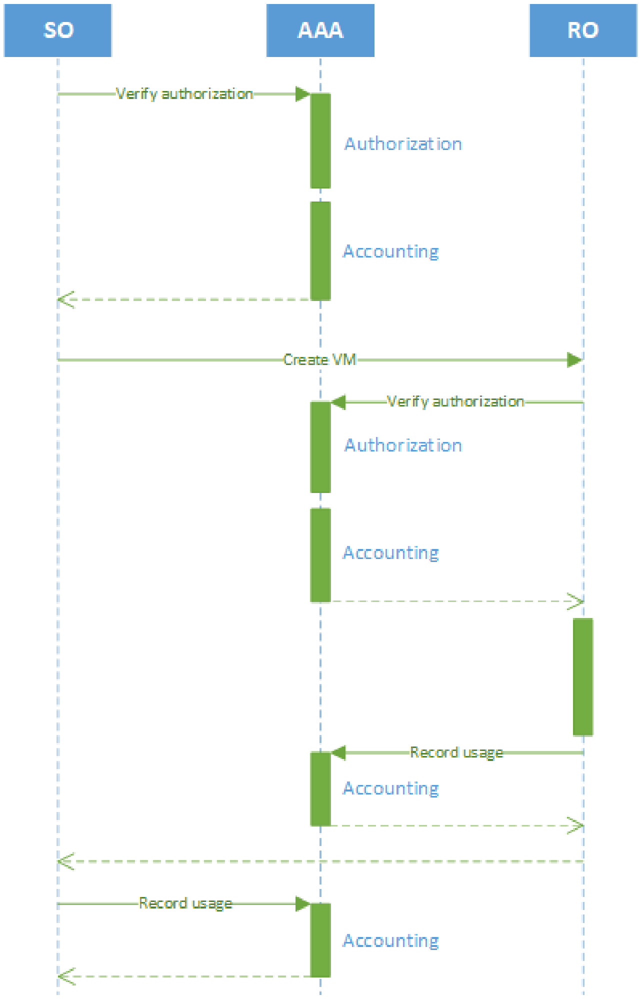

Figure 16: Virtual machine creation process flow (simplified)

5.2.8 Implementation

The AAA component is implemented in Python, using the Django REST Framework and Openstack Keystone. Python was the selected language to agree with the codebase of OSM, allowing for easier integration and maintenance of the project.

The AAA component supplies a REST API so that other components can communicate with it. It is through this API that the other OSM components will be able to interact with the AAA component.

Each component in OSM is deployed inside a Linux container, and for consistency, the AAA component is also deployed inside one (including Keystone).

5.2.9 Integration into OSM

As mentioned previously, we are making all the efforts so that the AAA component can be compliant with OSM design patterns, so when the component is ready, we expect that it can be integrated as a component into OSM.

5.2.10 Project status

Currently, the project is in the development phase, where Keystone is being integrated into the AAA component. There is also an effort into integrating Keystone with the 5GinFIRE portal, as mentioned in deliverable D3.1. The next phase will be the integration of the component above into OSM.

5.3 Virtualization based on OpenStack and containers for SBCs

One enhancement that is being considered to the current 5GinFIRE orchestration platform is the use of Single Board Computers (SBCs) as compute nodes within an operational cloud-computing platform. This enhancement aims at providing a limited-capacity and inexpensive infrastructure capable of instantiating lightweight VNFs, that is, VNFs which computational cost is not considerably high in comparison with the capacity provided by current infrastructure resources.

Towards providing the abovementioned enhancement, the Raspberry Pi (RPi) model 3[12] has been selected as SBC to act as the compute node in the proposed infrastructure, and OpenStack (release Ocata [7]) as the open-source cloud-computing platform. This composition presents a significant challenge since the RPi does not support the hardware acceleration needed for the instantiation of virtual machines, which encourages the use of containers for virtualization. In this respect, LXC containers[13] provided by Linux operating system has been the selected technology to build and check the performance of our designed system.

With this, we developed an early prototype of a virtualization system at UC3M, allowing the execution of lightweight VNFs over RPIs with the utilization of OpenStack and Linux

[12] Raspberry Pi – Teach, Learn, and Make with Raspberry Pi (last access: December 2017): https://www.raspberrypi.org

[13] Linux Containers (last access: December 2017): https://linuxcontainers.org

containers. The prototype is now functional at UC3M, and we are considering evolving this system into different directions. In particular:

- Analyse the performance of the virtualisation system.
- Explore the integration of LXD containers, as an enhancement with respect to the utilization of LXC containers.
- Expand the OpenStack networking configuration applied from layer 2 to layer 3, to provide routing services within the platform.
- Integrate the Open Source MANO framework within the system, to support the orchestration and deployment of network services composed by lightweight VNFs.
- Apply this evolved platform to a specific use case with micro drones as compute nodes, making possible the instantiation of lightweight and mobile VNFs in this emergent type of devices.

5.4 Utilization of public cloud infrastructures

From the version in use in the 5GINFIRE MANO framework (Release TWO) OSM has the ability to interact with VIMs corresponding to the cloud management systems of public cloud infrastructure..

There are some concerns regarding this usage, such as:

- Limitations of the orchestration capabilities. For example, in what relates to VNF placement decisions, EPA support, or SDN-based direct control of connectivity
- Data confidentiality and privacy requirements
- Data plane performance issues

But it is clear that the possibility of offloading some functions to a public cloud would not only add a new dimension to 5GINFIRE scalability, but also allow users to play with further degrees of freedom when planning their experiments, at least for some particular ranges of functions, or some specific kind of experiments. Furthermore, 5GINFIRE would allow for experimentation with hybrid NFV/cloud services, contributing to the blurring of the line separating IT and telecommunications.

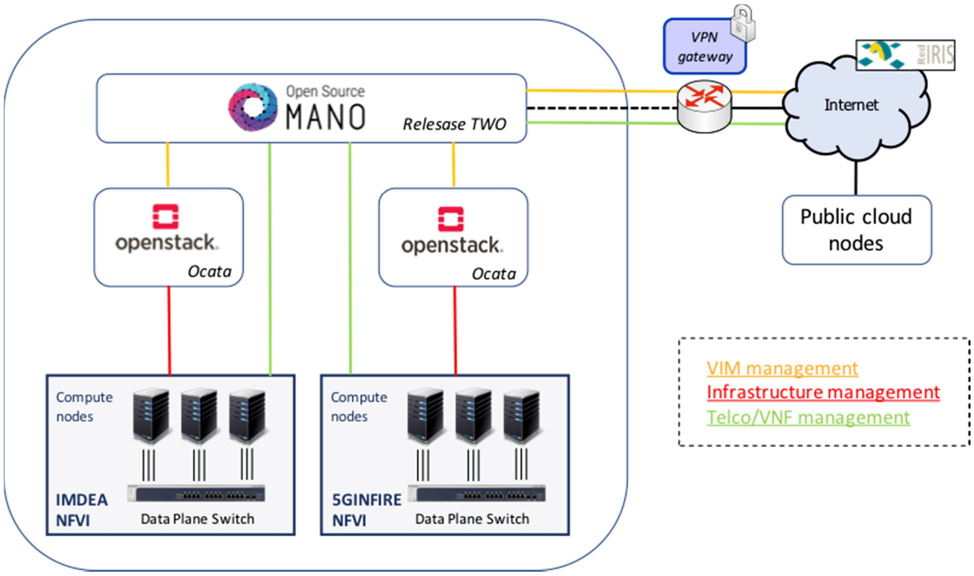

Figure 17: Integration of public cloud

Figure 17 above illustrates the foreseen scenario of public cloud integration in the 5GINFIRE infrastructure, over an Internet connection, in addition to the N participating testbeds providing a dedicated infrastructure with a tighter integration with the common orchestrator. There are issues to be addressed and verified regarding the different connectivity aspects discussed in previous sections, as well as isolation guarantees that need to be assessed, but given the potential advantages, the project team will start some initial tests once the first infrastructure deployment is stable.

5.5 UFU site

At the Federal University of Uberlândia (UFU), located in Minas Gerais state, in Brazil, the 5GINFIRE infrastructure is under deployment. Figure 18 presents an overview of UFU's site.

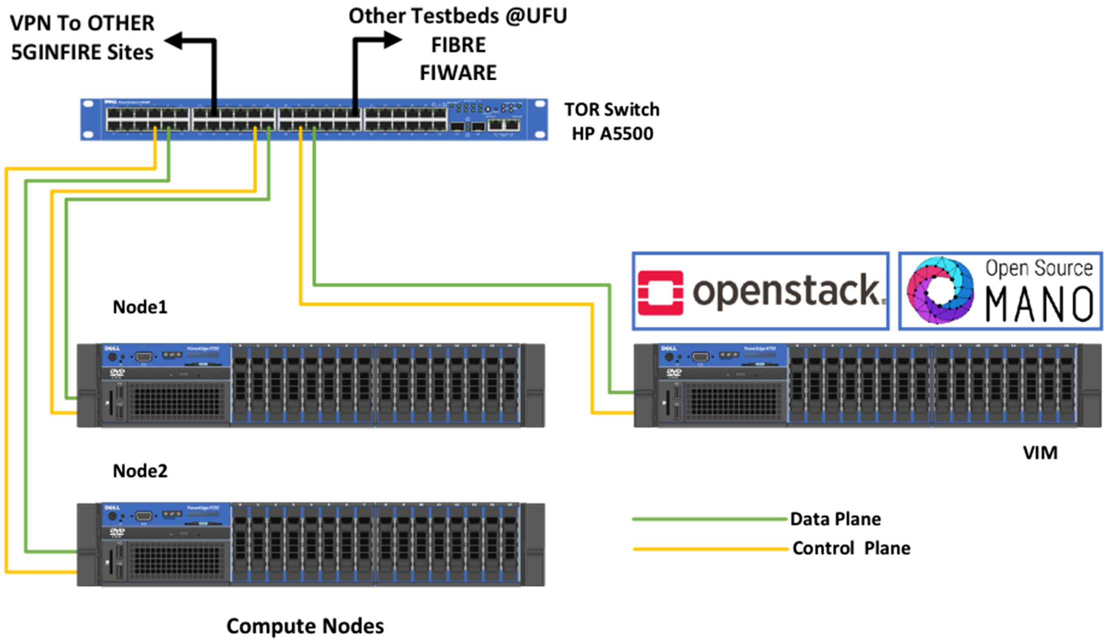

Figure 18: 5GINFIRE Infrastructure at UFU

The NFVI at UFU uses two compute nodes described below:

- Compute Node1: Dual processor Intel Xeon E5-2650 v2 with 32 cores; 96 GByes of RAM; 4 TBytes of storage; 4 x 1 Gbps network network interface cards;
- Compute Node2: Intel Xeon E5440 with 8 cores; 32 GBytes of RAM; 292 GBytes of storage; 4 x 1 Gbps network interface cards.

The VIM server, which hosts Openstack Ocata, has an Intel Xeon E5405 with eight cores, 24 Gbytes of RAM, 730 Gbytes of storage and 4 x 1 Gbps network interface cards. This server will also host an OSM release TWO that will be used locally for testing purposes. A Top of Rack (TOR) switch interconnects all the servers using a data plane network and a control plane network physically isolated. UFU's site will interconnect with 5TONIC and the others 5GINFIRE sites using a VPN.

As result of previous work, currently, at UFU there are two other testbed deployments of earlier projects: FIBRE [11] and FIWARE [12]. The unique infrastructure at UFU will allow to explore the integration scenarios between a FIRE based testbed (FIBRE) with 5GINFIRE.

5.6 Multi-Domain Orchestration

Network Function Virtualization will help to bring life the 5G vision [13] [14], and a 5G based service can reply on different VNFs on its deployment.

Considering that the operator is the owner of the infrastructure, to provide an end-to-end service, is necessary that the MANO interacts with different VIMs of each data center controlled by the operator. These data centers may be on edge and or in the core. This scenario has just one domain administrator, and it is called a single domain scenario.

A possible scenario for the 5G networks considers the existence of a multi-domain infrastructure where the edge cloud and the core cloud are can be different administrative domains, each one using its MANO components, as presented in Figure 21. Each operator has its MANO deployed, such as OSM. OSM orchestrates the resources across all the operator's data centers (represented by DC1 to DCn in Figure 19). This scenario is the one that is being focused on the 5GINFIRE facility in the first moment.

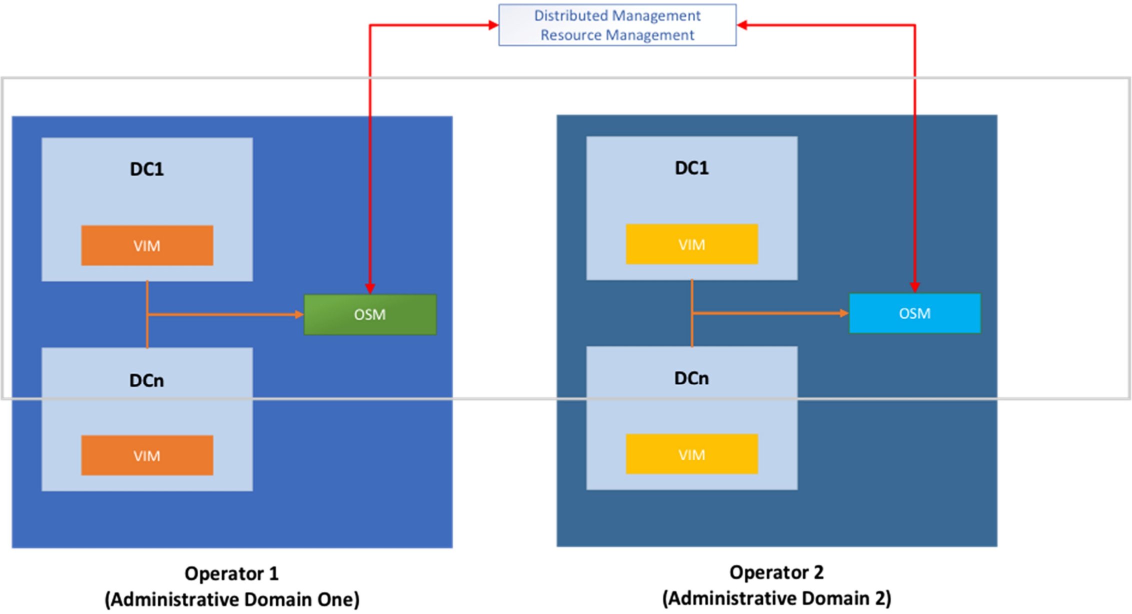

Figure 19: Multi-domain orchestration scenario

The multi-domain orchestration use case can be present in a resource sharing scenario, where more than one operator can pool together resources and share their infrastructure focusing CAPEX/OPEX reductions by creating a 5G slice with different infrastructure owners [15]. Another possible scenario that requires a multi-domain orchestration it would be an experimental deployment of an instance of a 5G vertical where the resources are in a FIRE based testbed and in the 5GNFIRE facility.

In this case, a multi-domain orchestrator component will be necessary to orchestrate the resources in each domain. This scenario poses several challenges [16] to support the distributed management and the resource management.

OSM release TWO does not contemplate the inter-orchestrator communication, whether inter-domain or not. A possible enhancement foreseen to 5GINFIRE orchestration platform is to support the Multi-Domain Orchestration (MdO). This be can accomplish by an OSM module that plays the role of a MdO. To do this enhancement it is relevant to explore results from a 5G related project, called 5GEX [17], that is considering this MdO aspect to implement network service and resource orchestration across multiple administrative domains. 5GINFIRE Open Calls can also address this enhancement.

6 Conclusion

One of the initial challenges for 5GINFIRE was to create a multi-domain stable MANO platform, suitable to orchestrate a widely diverse set of network-hosted functions, beyond VNFs in the strict sense, and to support an evolution path both for the MANO platform itself, and for the functions and services to be managed.

The WP4 5GINFIRE team has addressed this challenge and produced such a MANO platform, considering the aspects related to a multi-domain and multi-user environment. The MANO platform has been validated, and its ability to deploy and execute cross-site services demonstrated. A set of valuable documentation on these tests, usable as future reference by 5GINFIRE experimenters, has been produced. Subsequent integration tests will be carried out in WP5, which will also verify the appropriate operation of the 5GinFIRE MANO platform for the use cases considered by the project, coordinated with the other management elements. In addition, the mechanisms related to monitoring and maintenance of the diverse components of the MANO platform and the NFVIs will be covered by WP6.

At the same time, the team has established the path for the MANO platform evolution, open to all the possible ways for achieving it:

- Direct contribution by project partners.
- Results from the open calls.
- Updates by the upstream projects.

For the first two cases, contribution to the upstream projects (OSM, OpenStack, OpenDaylight…) will be attempted whenever possible and, in fact, has already been made in one case (see section 5.1). To guarantee a seamless evolution, the team is working in establishing a continuous integration environment.

References

[1] http://5GINFIRE-5g.eu/ © 5GINFIRE consortium 2017

[2] Open Source MANO. ETSI-hosted project (last access: Dec 2017): https://osm.etsi.org

[3] Riwal Kerherve et al., “Evolving FIRE into a 5G-Oriented Experimental Playground for Vertical Industries”, 5GinFIRE Deliverable D2.1, May 2017.

[4] 5TONIC: an open research and innovation laboratory focusing on 5G technologies (last access: Dec 2017): https://www.5tonic.org

[5] Adrian Hoban et al., “OSM Release TWO, A Technical Overview”, ETSI OSM Community White Paper, April 2017.

[6] ETSI GS NFV 002 V1.2.1, “Network Functions Virtualisation (NFV); Architectural Framework”, version 1.2.1, Dec. 2014.

[7] OpenStack Ocata (last access: Dec 2017): https://releases.openstack.org/ocata/

[8] The 5GinFIRE orchestration service and the NFV infrastructure at the 5TONIC laboratory (last access: Dec 2017): https://5ginfire.eu/5tonic/

[9] Reference VNF and NS Descriptors (Release TWO), OSM Wiki, (last access: Dec 2017): https://osm.etsi.org/wikipub/index.php/Reference_VNF_and_NS_Descriptors_(Release_TWO)

[10] Ansible, Automation for Everyone, Red Hat (last access: Dec 2017): https://www.ansible.com

[11] FIBRE, “Future Internet Brazilian Environment for Experimentation (FIBRE),” 2017. [Online]. Available: http://fibre.org.br/. [Accessed: 19-Dec-2017].

[12] FIWARE, “FIWARE,” 2017. [Online]. Available: https://www.fiware.org/. [Accessed: 19-Dec-2017].

[13] S. Abdelwahab, B. Hamdaoui, M. Guizani, and T. Znati, “Network function virtualization in 5G,” IEEE Commun. Mag., vol. 54, no. 4, pp. 84–91, Apr. 2016.

[14] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. D. Turck, and R. Boutaba, “Network Function Virtualization: State-of-the-Art and Research Challenges,” IEEE Commun. Surv. Tutor., vol. 18, no. 1, pp. 236–262, First quarter 2016.

[15] K. Samdanis, X. Costa-Perez, and V. Sciancalepore, “From network sharing to multi-tenancy: The 5G network slice broker,” IEEE Commun. Mag., vol. 54, no. 7, pp. 32–39, Jul. 2016.

[16] R. Guerzoni et al., “Analysis of end-to-end multi-domain management and orchestration frameworks for software defined infrastructures: an architectural survey,” Trans. Emerg. Telecommun. Technol., vol. 28, no. 4, p. n/a-n/a, Apr. 2017.

[17] 5GEX, “5G Exchange Project (5GEx),” 2017. [Online]. Available: https://www.5gex.eu. [Accessed: 19-Dec-2017].