Open Source MANO

OSM White Paper

OSM Release FIVE
Technical Overview

1st edition – January 2019

ETSI

Authors

Adam Israel	Canonical
Alfonso Tierno Sepúlveda	Telefónica
Andy Reid	BT
Felipe Vicens	ATOS
Francisco Javier Ramón Salguero	Telefónica
Gerardo García de Blas	Telefónica
Gianpietro Lavado	Whitestack
Mark Shuttleworth	Canonical
Matt Harper	Rift.io
Michael Marchetti	Sandvine
Ricard Vilalta	CTTC
Silvia Almagia	ETSI
Vanessa Little	VMware

Contents

1 Introduction to Release FIVE

The rapidly expanding Open Source MANO Ecosystem has resulted in an explosive growth of innovation in the OSM platform.

Release FIVE, the sixth major OSM release, is evidence that OSM has become a reliable and viable option for NFV orchestration over production environments. With five successful Hackfests, dozens of PoCs and demonstrations, active participation to all the ETSI NFV Plugtests, and a stretch collaboration with the main 5G European research projects, the community has not only attracted more attention from academics and scientists in the field, but also from companies in the NFV industry which have already started to offer commercial solutions around OSM, including training, consulting, VNF onboarding services and MANO distributions for production environments.

Release FIVE, the most ambitious release to date, sees the addition of an exciting set of new features, following industry priorities through the operator-led OSM End-User Advisory Group (EUAG). These features include:

- Network slicing for 5G.
- Support for physical and hybrid network functions.
- Multi-domain orchestration providing dynamic multi-site inter-datacentre connectivity.
- Enhancements to the monitoring and policy framework providing closed-loop operations.
- New capabilities for VNF configuration and deployment.
- Enhancements to general user experience.

Furthermore, Release FIVE is the most stable release to date, with significant security enhancements, readiness for high-availability deployments, and continuous improvements to the CI/CD pipeline.

OSM continues to focus on "ease of use" for its existing community, and in order lower participation barriers for new members, by adding to the lean installation requirements already incorporated in the previous release. As a result, Release FIVE is distributed as a set of pre-built downloadable containers that can be easily deployed in minutes, with docker images available for anyone in Docker Hub, so that users willing to quickly deploy and experiment OSM do not have to go through the whole process of building the platform from source.

In summary, , with the significant advancements incorporated into Release FIVE, the OSM community foresees a successful adoption of the platform as a key enabler for carrier-class NFV MANO deployments around the world.

2 Architectural Progression

Release FIVE builds on the micro-service architecture adopted in Release FOUR, with the incorporation of a Time Series Database as a new common service for metrics storage, and the addition of a policy module (POL) for close loop automation.
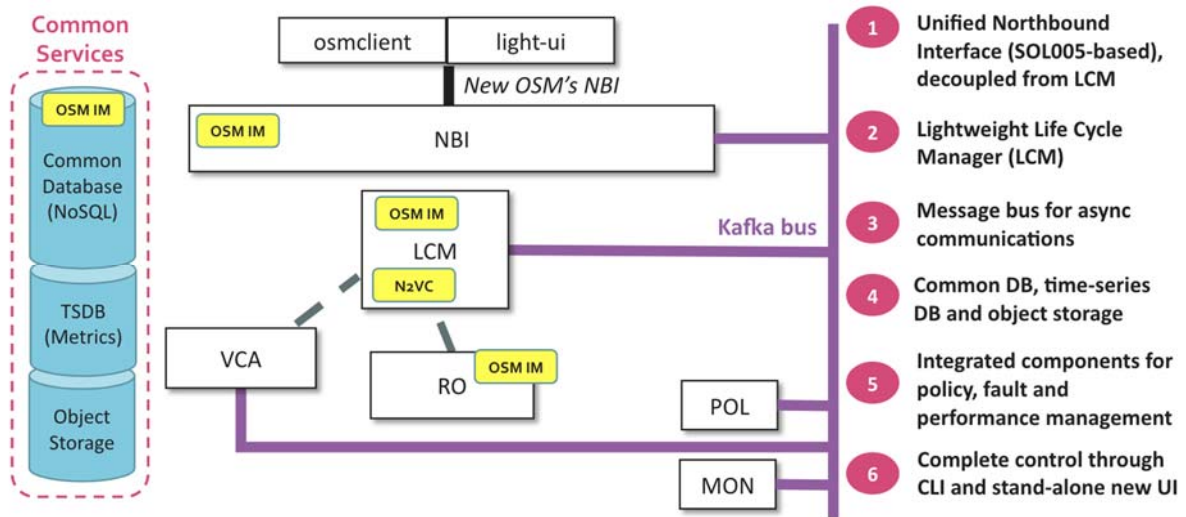

Figure 1. OSM Release FIVE Architecture

3 Release FIVE Features Overview

3.1 Network Slicing for 5G

One of the key new features within OSM Release FIVE is Network Slicing. This feature was proposed to improve the existing network infrastructure resources usage and management as described in [NGMN-MSCH16] and contributed to OSM in collaboration with the 5G Tango research project

A Network Slice Instance is defined as a set of network functions and the resources for these network functions which are arranged and configured, forming a complete logical network to meet certain network characteristics.

The implementation of Network Slicing in OSM Release FIVE follows the recommendations for management and orchestration in [3GPP-TS28.801], aligns with the data model specified in [3GPP-TS28.541] and follows the guidelines for network slicing implementation in NFV architectures described in [ETSI-NFV-EVE021].

In the OSM Information Model, a network slice has been defined as an interconnected list of network services, which provides enough flexibility for deploying unique network slices, while re-using the necessary constructs for deploying network services.
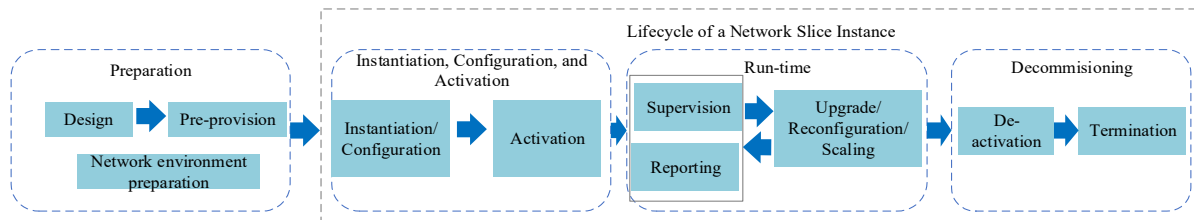

Figure 2. Life Cycle of a Network Slice Instance

OSM distinguishes between Network Slice Templates, where the definition of the network slices is provided (similar to network service descriptors), and Network Slice Instances, where the run-time record of the deployed Network Slice is provided (similar to network service records).
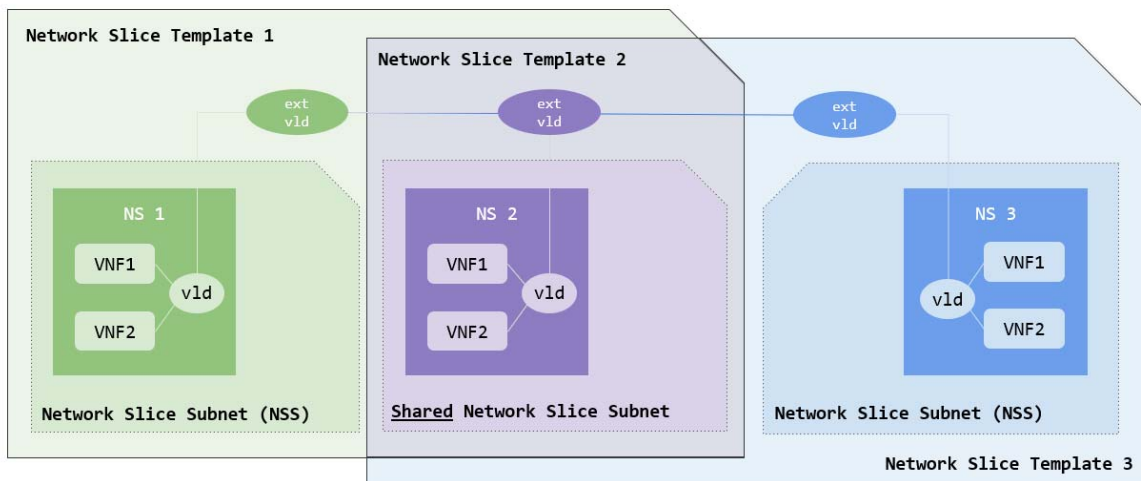

Figure 3. Network Slice Templates and Subnets

3.2 Support of PNFs, HNFs, and pools of PDUs

Virtualized Network Functions are not always the only components in an end to end network service. In this release, the Information Model has been enhanced to support Physical Deployment Units and Hybrid Network Functions.   The Northbound API has been extended to control these PNFs and HNFs, and these physical devices can now take advantage of the life cycle management and configuration capabilities of charms. In addition, similar PDUs can be grouped and managed as a pool to streamline the operational lifecycle.
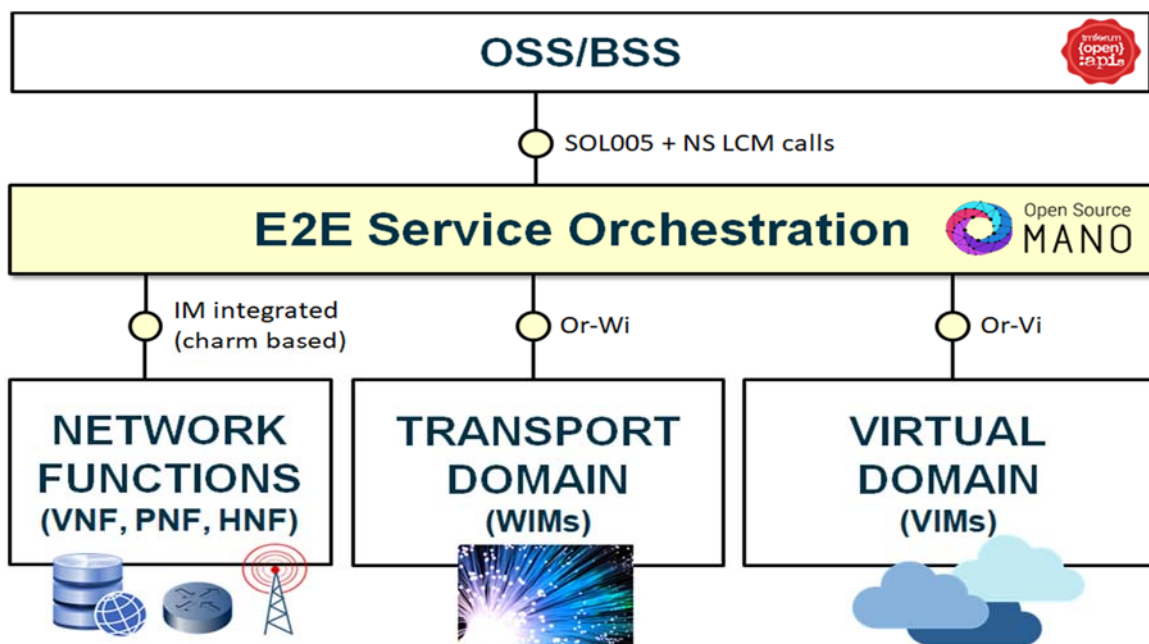

Figure 4. Cross-domain End-to-End Service Orchestration with OSM

3.3 Multi-site Inter-DC connectivity through a WIM plugin model

 Dynamic multi-site deployments are key in the evolution of orchestration of complex network services. In collaboration with the Metro-Haul research project, a WAN Infrastructure Manager (WIM) plug-in model has been created to enable the integration of different WAN control API technologies in the dynamic establishment of multi-site network services across WAN networks. This feature can also be leveraged to enable multi-site failover scenarios for disaster recovery or maintenance.
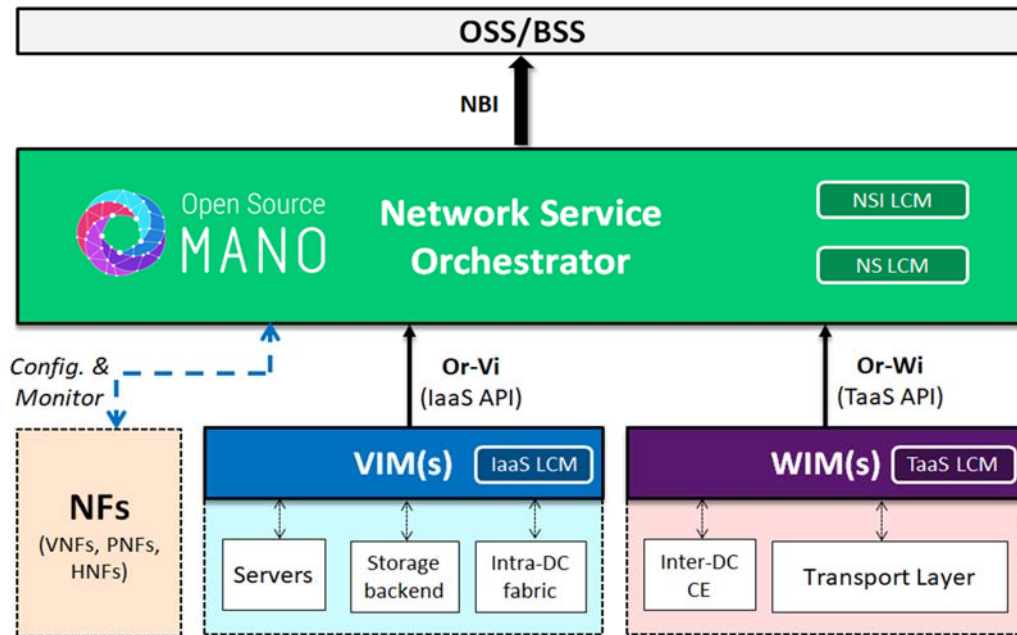
Figure 5. New WIM Plugin Model

3.4 Monitoring and Policy framework enhancements

3.4.1 Enhanced Metrics collection (MON)

Metrics collection is now supported from both the infrastructure (VIM) and directly from VNFs. This is made possible due to the new architecture of OSM MON, which now includes the following components:

- **mon-central**: Handles VIM and alarms-related messages in the Kafka bus.
- **mon-collector**: Iterates continuously over the VNFRs and collects metrics using VIM and VCA plugins, then exposes them to be consumed by the Prometheus Time Series Data Base (TSDB).
- **mon-evaluator**: Iterates continuously over alarms and evaluates them. In case of triggering, it sends a notification over the Kafka bus.
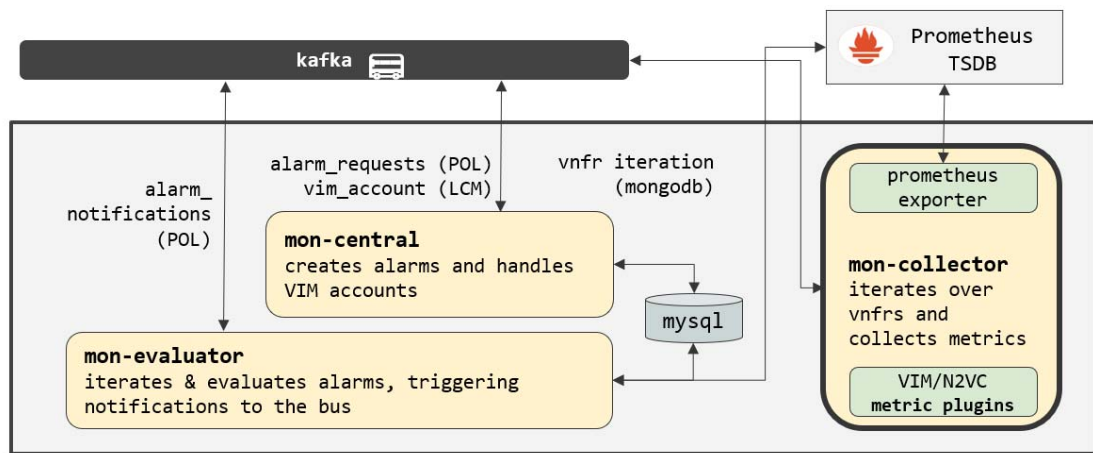
Figure 6. OSM MON internal architecture

3.4.2 VNF Alarm Management

Through the new Policy Manager (POL) and mon-evaluator components, OSM is now able to create, manage and trigger alarms based on infrastructure or VNF events and metrics. In this release, alarms are created as part of auto-scaling rules based on metric thresholds. In future releases, alarms will be associated to any kind of event that is relevant for the proper operation of Network Services.

3.4.3 Time Series Data Base for Metrics

OSM has now the ability to store and correlate metrics in a local, highly scalable and performant Time Series Data Base provided by the well-known and supported Prometheus project. This addition brings enormous possibilities on enhancing lifecycle automation based on metrics aggregation and correlation, independently of their source.

3.4.4 Service Assurance with Autoscaling VNFs

Through the new Policy Manager and in coordination with LCM orchestrator, OSM is now capable of automating the horizontal scaling decisions with a VDU granularity. This feature can be easily activated through including a scaling descriptor in the VNFD, bringing elasticity of services triggered by any kind of metric. Furthermore, OSM is also completing the set of SOL005-based calls by adding manual control to horizontal scaling as described in [ETSI-NFV-SOL005].

3.5 New capabilities for VNF configuration

3.5.1 Support for Multi-VDU VNFs

The information model for VDUs has been extended to support charms at the VDU level. This way, existing single-VDU VNFs work as previously, but more complex VNFs can specify different VNF configuration charms for each VDU. This feature will have additional enhancements in Release SIX.

3.5.2 Full Charm Support

It is now possible to have full native charm support for installation, configuration and day 2 life cycle operations as well as scale-in and scale-out operations. Support for VNFs that are installed and configured via a charm has also been included. This is an elegant complement to what was already implemented in earlier releases with proxy charms.

3.5.3 Secure Key Management

OSM Release FIVE provides automatic ssh key generation and injection into newly deployed VMs, and centralized key management.

3.6 New capabilities for VNF deployment

3.6.1 Selective support of MAC address and IP Profiles

Many legacy, non cloud-native, VNFs assume that the VIM can provide specific MAC addresses to specific interfaces of their VDUs/VMs.

To support those legacy VNFs, OSM Release FIVE allows to specify the MAC address of the VDU interfaces either in the descriptor or as an instantiation parameter.

In the case of IP profiles, a coherency check is mandatory to prevent these MAC addresses created by the VIM from creating conflicts between different instances of the same NS/VNF or between VNFs interconnected in a NS.

In order to guarantee that coherency, MAC addresses can only be set in a descriptor where they correspond to VLs or CPs that are considered ″internal″ to the corresponding VNFD or NSD. Thus:

- At VNFD level, a specific MAC address can be set only in internal interfaces/CP.
- At NSD level, same rules as above would apply, but in the context of a NS template (now, many ″external″ VNFD's CPs or VLDs might become ″internal″ for the NSD).
- If required, parameters passed at instantiation time can allow to set MAC addresses of external CP of the NS.

3.6.2 Specify Management IP addresses at deploy time
Having a management interface ready and reachable right after boot is critical for driving Day 1 and Day 2 operations of VNFs. Previously, OSM assumed that the VIM assigned an IP address to the VNF from the pool in the management network, which the VNF could retrieve via DHCP or metadata service, if available.

However, many legacy OSS assume that the VNFs should have a specific IP address assigned from a pool of addresses that they manage (or that they can learn from other systems).

For those cases, OSM can provide a mean to be told, at VNF instantiation time, which management IP address it to be used. Please note that this address cannot be part of the VNFD or the NSD or, otherwise, several instances of the VNF/NS could end up having the same IP address, leading to collisions, etc.

3.6.3 Service Function Chaining
Service Function Chaining is a key NFV technique that allows operators to steer traffic across VNFs, and an important part of network service composability since it enables to pick and choose different VNFs, interconnect them and force the traffic to pass through them in the desired order.

Service Function Chaining (SFC) orchestration has been extended in OSM Release FIVE to support OpenStack neutron-sfc plugin features. Also, the support of multiple input/output ports has been added, making OSM the orchestrator of choice for SFC automation in other projects, such as OPNFV. OSM Release FIVE supports asymmetrical chains as defined by [IETF-RFC7665].

3.7 Platform enhancements

3.7.1 RBAC for the platform
OSM Release FIVE provides extended capabilities for the Role Based Authentication and Control of the platform, which will continue to evolve in Release SIX. The NBI interface can now be authenticated against an external database. Release FIVE supports Keystone integration and future releases will enable the support of additional user databases. As of Release FIVE, this feature is classified as experimental.

3.8 User Experience Enhancements

3.8.1 Deployment
The Open Source Mano platform is now easier to deploy than ever before. With downloadable containers available in DockerHub, vagrant images and virtual machine image installs. OSM users can now download and have an operational system in a matter of minutes. Further details are provided in Section 6 DevOps update.

3.8.2 GUI Composer
Significant enhancements have been made to the Graphical User Interface to include a network service composer, making it very simple to model and modify network services.

3.8.3 Event and Log Visualization
With the integration of Prometheus and Graphana in the OSM platform, telemetry and performance views are now consistently provided across the platform, providing a uniform look and feel to the monitoring interface regardless of which monitoring engine is plugged in.

4 OSM DEVOPS, CI/CD and Remote Labs

OSM continues to improve the automated CI/CD pipeline and build and test automation in every release. The diagram below depicts the different stages and technologies involved in the pipeline.
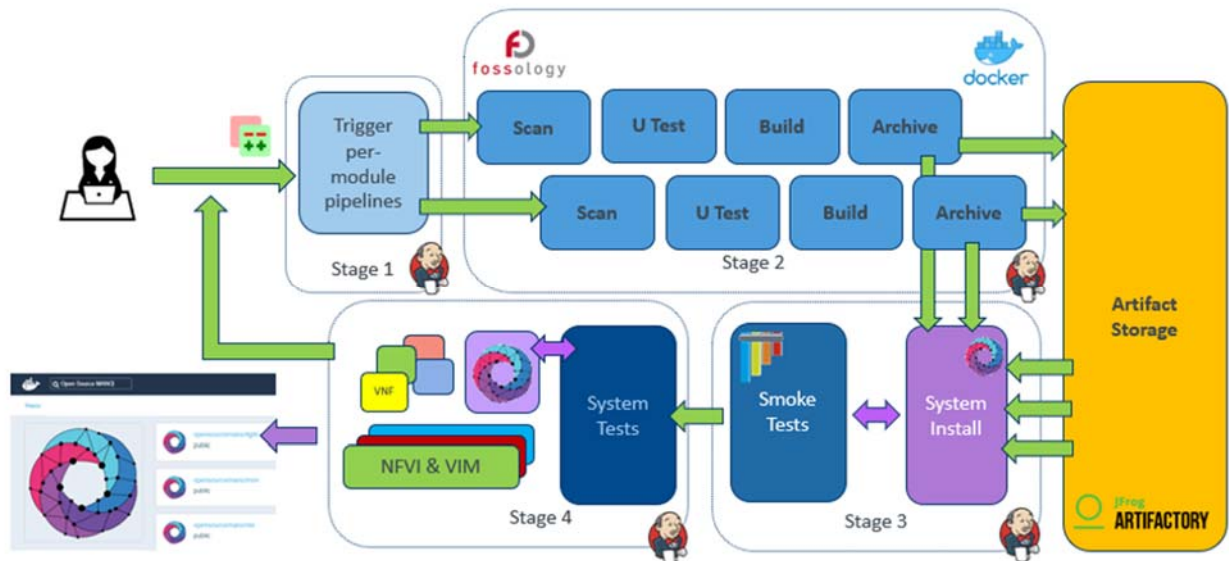

Figure 7 OSM CI/CD pipeline

From Release FIVE, at the end of the pipeline, OSM Docker images are published to Docker Hub for each successful build, and vagrant images are published to Vagrant Cloud for easy download and install of stable OSM releases.

OSM CI/CD Stage 4 leverages the Hub for Interoperability and Validation at ETSI (HIVE) to access a growing pool of OSM Remote Labs. These labs, provided by the OSM community, allow to test OSM with different combinations of HW, VIMs and SDN Controllers, and fully validate OSM end-to-end functionality and interoperability with external components.
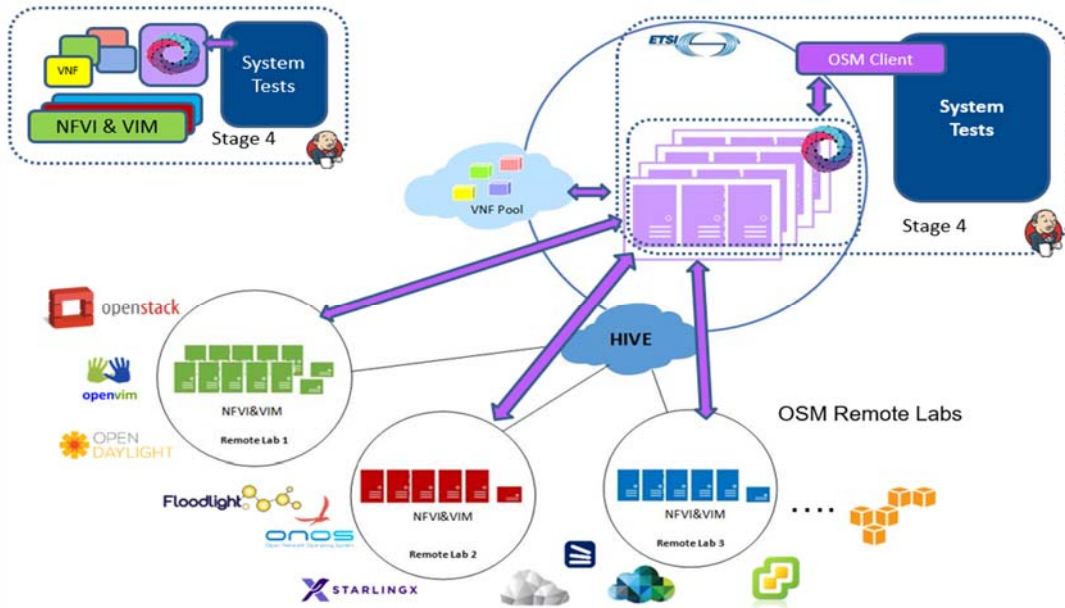
Figure 8 OSM CI/CD Stage 4 & Remote Labs

The OSM Remote Labs and HIVE are also key enablers of the OSM Hackfests, where participants have the opportunity to get hands-on experience with OSM and test experimental features on different environments provided by the community
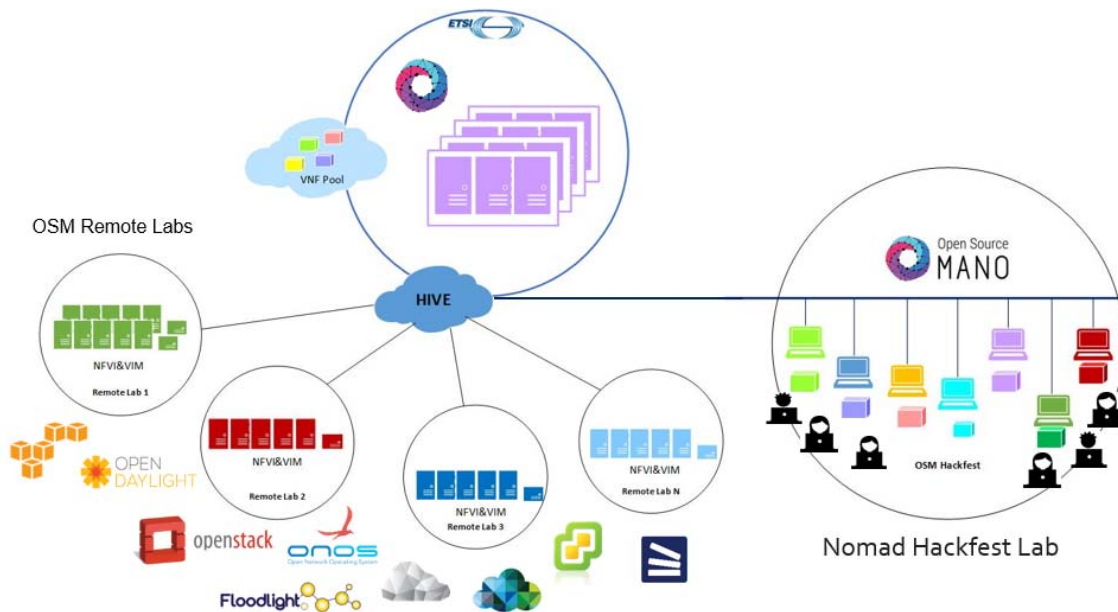
Figure 9 OSM Hackfest & Remote Labs

As a side note, OSM has been integrated as upstream project in OPNFV XCI (cross community CI) where it is actively tested with OpenStack and OpenDaylight upstream code.

5 VNF On-Boarding

The support for event and performance monitoring of VNFs means the OSM is now able to provide complete automation for all operations across the full lifecycle of a VNF as well as of NSs and network slices which can be composed of VNFs. From the perspective of OSM's End User Advisory Group (EUAG) this is a major milestone. However, the EUAG also notes that this full automation that OSM can provide can only be realized if all relevant processes covering the full lifecycle of the VNFs, from Day-0 and Day-1 processes to Day-2 operations, are properly considered during VNF on-boarding and included in the corresponding VNF Package.

Full automation is one of the main objectives of NFV as this is where most of the cost benefits lie. In order to realize these cost benefits, it is essential that human intervention is removed from all operational processes in the lifecycle. Moreover, this has further benefits beyond the simple cost saving of avoiding the human processes.

- First, full automation allows operation at a scale and with a consistency which is hard to achieve with human intervention.
- Second, full automation allows NSs to be encapsulated and fully controlled by other software systems through the OSM NBI in such a way that automation can be applied recursively (also greatly aiding operation at scale).

Both of these are important features identified by the EUAG members. As a simple illustration, having OSM automatically stand up a VNF but then require a skilled technician to configure the VNF as appropriate (including activating any licensing) would fall significantly short of the automation objective, and would be an underuse of the capabilities that OSM provides. Similarly, requiring a skilled technician to interrogate and interpret event and performance information in order reoptimize poor performance of an NS, equally would fall significantly short of the automation objectives.

OSM now provides for automation across the full lifecycle of the entities it orchestrates:

- Day 0: the initial instantiation of an entity by instantiating its components and connecting them together (this also naturally also encompasses the reverse process of decommissioning entities and freeing resources).
- Day 1: the initial configuration of the instantiated components sufficient to bring the entity into full operation.
- Day 2: includes two kind of processes:
  - o Any subsequent reconfiguration of the components upon the call of primitives, and providing the mechanism to implement those transactions in a reliable manner.
  - o handling of event and performance information and appropriate responses (again, identifying the possible triggers and responses is an example of information which needs to be included in the descriptors at the on boarding stage).

As a result, the on-boarding of a VNFs (and preparation of the corresponding package) is a phase substantially more relevant than simply specifying the components to be instantiated and how they should be interconnected. Therefore, not all on-boarding may be of equal value if some VNFs are on-boarded with insufficient detail in their package to automate the full lifecycle of the VNF. The EUAG would therefore greatly encourage the OSM community to make full benefit of all the automation features across the full lifecycle when VNFs are on-boarded including the full range necessary specification in the appropriate packages as part of VNF on boarding process.

6 References

[3GPP-TS28.801]   Telecommunication management; Study on management and orchestration of network slicing for next generation network

[3GPP-TS28.541]   Management and orchestration; 5G Network Resource Model (NRM); Stage 2 and stage 3.

[ETSI-NFV-EVE012]   Network Functions Virtualisation (NFV) Release 3; Evolution and Ecosystem; Report on Network Slicing Support with ETSI NFV Architecture Framework

[ETSI-NFV-SOL005]   Network Functions Virtualisation (NFV) Release 2; Protocols and Data Models; RESTful protocols specification for the Os-Ma-nfvo Reference Point

[IETF-RFC7665]   Service Function Chaining Architecture

[NGMN-MSCH16]   Mschner K. et all Hedmar, P. Description of network slicing concept. NGMN Alliance document, January 2016.

ETSI
06921 Sophia Antipolis CEDEX, France
Tel +33 4 92 94 42 00
info@etsi.org
www.etsi.org