

Open Source MANO

An ETSI OSM Community White Paper

OSM RELEASE ONE

A TECHNICAL OVERVIEW

October 2016

Authors:

Adrian Hoban, Chair OSM TSC, Intel

Alfonso Tierno Sepulveda, OSM Resource Orchestration MDG, Telefónica

Gerardo García de Blas, OSM TSC, Telefónica

Kiran Kashalkar, OSM User Interface MDG, Rift.io

Mark Shuttleworth, OSM TSC, Canonical

Matt Harper, OSM TSC, Rift.io

Rajesh Velandy, OSM NW Service Orchestration MDG, Rift.io

Editor:

Chris Buerger, Convener OSM Marketing Task Force, Intel

ETSI

06921 Sophia Antipolis CEDEX, France

Tel +33 4 92 94 42 00

info@etsi.org

www.etsi.org



Contents

Contents	2
Introduction & Scope	4
OSM Scope	5
Run-Time Scope	5
Design-Time Scope	6
OSM Development Themes	7
OSM Release ONE Overview	8
Interoperability	8
User Experience	8
Modelling Enhancements	8
Release ONE Highlights	8
OSM Release ONE Details	10
On-boarding experience & VNF Packaging	10
Cloud Init Support	10
Create Networks at VIM	10
Remove NSD Datacenter Network Reference	11
Error Messages and Logging	11
YAML Format Descriptor Display	11
Enhance Visual Differentiation between NS Catalog and VNF Catalog	11
Restructure Layout of Service Primitive Page	11
Package Creation Command Line Utility	12
EPA based Resource Allocation	12
(Network) Service Modelling	12
Juju-2.x	12
Network Types in RO	13
Allow IP Parameters for Networks	14
Configuration/Service Primitive Model Enhancements	14
Multi-VIM	14
New VIM connector for VMware vCloud Director	14
OpenVIM as Reference VIM with EPA Capabilities	15



Datacenter Capabilities	15
Support for VIM Accounts	15
Multi-Site	15

Introduction & Scope

ETSI OSM is an operator-led ETSI community that is delivering a production-quality open source Management and Orchestration (MANO) stack aligned with ETSI NFV Information Models and that meets the requirements of production NFV networks.

OSM Release ONE has been engineered, tested and documented to allow for rapid installation in operator labs worldwide that seek to create a scalable and interoperable open source MANO environment. It substantially enhances interoperability with other components (VNFs, VIMs, SDN controllers) and creates a plug-in framework to make platform maintenance and extensions significantly easier to provide and support.

In addition, Release ONE improves administrator and developer experience, both in terms of usability and installation procedure as well as enhances modelling of VNFs and network services. In line with the goals of the OSM open source project, the output of this modelling work will be contributed to ETSI NFV. Release ONE also provides extremely flexible virtualized network functions (VNF) configuration and advanced networking management as well as improved troubleshooting capabilities, with advanced logging.

This White Paper outlines the main architecture of OSM, the new capabilities developed and open-sourced as part of Release ONE. It also provides insight into a number of the development themes the OSM Technical Steering Committee (TSC) has been pursuing in close collaboration with the OSM End-User Advisory Group (EUAG) for this release.

More information about ETSI OSM, its community and how to download OSM Release ONE can be found here: <https://osm.etsi.org/>. The authors of this White Paper would like to extend a sincere 'Thank You' to the entire OSM community to make Release ONE possible.

OSM Scope

The OSM community has defined an expansive scope for the project covering both design-time and run-time aspects related to service delivery for telecommunications service provider environments. The express goal is that the OSM code base can be leveraged in these environments as-is in a Roll-Your-Own context, or in whole or part of a commercial product offering.

Figure 1 shows the approximate mapping of scope between the OSM components and the ETSI NFV MANO logical view (the background image has been extracted from Figure 4 in the NFV Reference Architecture Framework, ETSI GS NFV 002 V1.2.1 (2014-12))

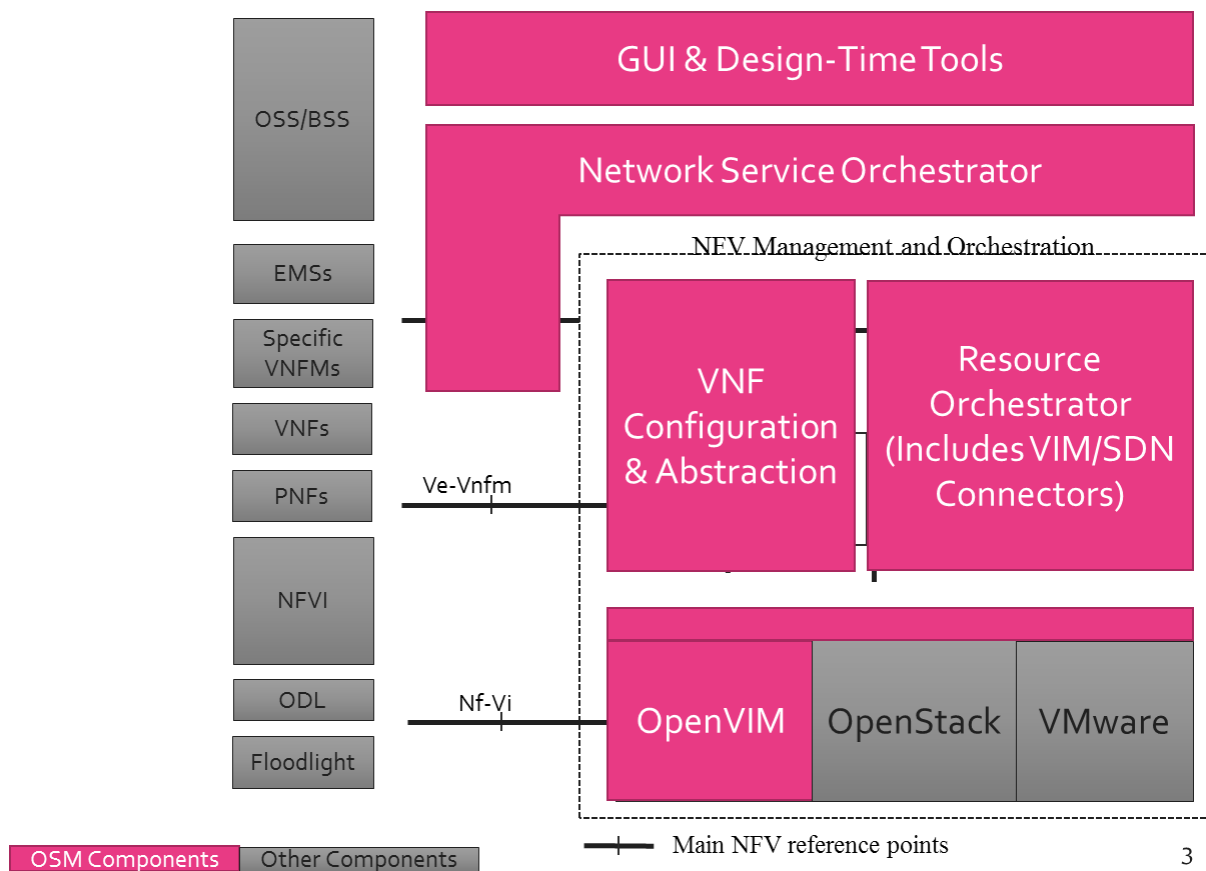


Figure 1 OSM Mapping to ETSI NFV MANO

Run-Time Scope

The run-time scope of OSM includes:

- Automating end-to-end Service Orchestration environment that enables and simplifies the operational considerations of the various lifecycle phases involved in running a complex service based on NFV.
- Superseding of ETSI NFV MANO where the salient additional area of scope includes Service Orchestration but also explicitly includes provision for SDN control.



- Delivering on a plug-in model for integrating multiple SDN controllers.
- Delivering on a plug-in model for integrating multiple VIMs
- Including one reference VIM that has been optimized for Enhanced Platform Awareness (EPA) to enable high performance VNF deployments.
- Integrating a “Generic” VNFM with support for integrating “Specific” VNFMs.
- Facilitating support for OSM to integrate Physical Network Functions into an automated Network Service deployment.
- Being suitable for both Greenfield and Brownfield deployment scenarios.
- GUI, CLI and REST interfaces to enable access to all features.

Design-Time Scope

The design-time scope of OSM includes:

- Delivering a capability for Create/Read/Update/Delete (CRUD) operations on the Network Service Definition.
- Supporting a Model-Driven environment with Data Models aligned with ETSI NFV MANO.
- Simplifying VNF Package Generation.
- Supplying a Graphical User Interface (GUI) to accelerate the network service design time phase.

OSM Development Themes

The OSM community has identified a number of development themes that are used to direct community innovation over multiple release cycles. While innovation is certainly not limited to these areas, having this focus has helped the community to deliver Release ZERO ahead of schedule with Release ONE released after just four months instead of the initial projection of six months¹.

The Release ONE development themes are:

- On-boarding experience & VNF Packaging to lower the barrier of entry for VNF vendors.
- Simplified install and upgrade process to accelerate adoption and deployment combined with an improved development environment to facilitate an expansion of the developer community.
- EPA based resource allocation to facilitate high performance VNF deployments with lower Total Cost of Ownership for the operator.
- Service Modelling to simplify, accelerate and standardize the design-time phase.
- Multi-VIM support expanding OSM so that VMware, OpenStack and OpenVIM are enabled.
- Multi-Site support enabling automated service delivery across multiple sites where a site is represented as a grouping of infrastructure managed by a VIM.

¹ N.B. It should be noted that OSM does intend to track more closely to an approximate six month release cadence for Release TWO onwards.

OSM Release ONE Overview

OSM Release ONE has made significant steps in advancing on each of the themes noted above. While over thirty distinct capabilities have been progressed, there are three salient categories of innovation that articulate the cohesion of this release. These were **interoperability**, **user experience** and **enhanced modelling**.

Interoperability

One of the guiding principles of OSM is that each component is both replaceable and pluggable. To that end, Release One has taken a substantial step forward to drive **interoperability** with other components (VNFs, VIMs, SDN controllers) and to create a plug-in framework to make it easier to maintain.

User Experience

A significant portion of the development effort within the OSM technical community has focused on responding to input from developers and operators to greatly enhance the **user experience**, both in terms of usability and installation procedure.

Modelling Enhancements

One of the greatest challenges facing the entire community of technologists developing specifications, standards and implementations for NFV relates to the topic of information and data modelling of the NFV solution space. For Release ONE, the OSM community has developed further **enhancements to modelling** of VNFs and Network Services. Based on the implementation experience gained in developing OSM Release ONE, the community remains committed to collaborating with the ETSI NFV Industry Specification Group (ISG) via formalized feedback to help harmonize and progress the VNF and Network Service Descriptors.

Release ONE Highlights

The OSM community have been delighted to welcome so many new members and participants to the project, with more than 45 companies joining before Release ONE was completed. Of the new members, OSM would in particular like to thank **VMware** for a significant contribution that has enabled OSM to natively support VMware based VIMs. This brings the number of supported VIMs in OSM to three, with OpenStack and OpenVIM already supported in Release Zero.

Release ZERO was able to leverage OpenVIM, but for Release ONE **OpenVIM** code has been accepted in the OSM GitHub repository and is now under its governance. OpenVIM provides a reference VIM for all-in one installations with full support of Enhanced Platform Awareness. That is, the user now does not need to have a pre-existing VIM installation in their premises before trying to install OSM. In case they do not have one, they can now use OpenVIM very easily.

Additionally, with Release One there is an improved Resource Orchestrator (RO) **plug-in model** to support simpler addition of other types of VIMs and SDN Controllers.



Support for deploying Network Services across **multiple sites** is a significant step forward for OSM. That is, it is now possible to deploy a Network Service that spans across multiple datacenters.

The inclusion of a **one-step install** process, based on containers and Juju modelling, represents a major accomplishment for OSM. This will make testing, customization and deployment much easier and more efficient. .

OSM Release ONE has extended VNF and Network Service models, which now allow **Day-Zero VNF configuration** (this is a configuration “on-the-fly”, before the first execution, that is extremely popular in cloud computing) and **advanced networking management**.

With a view towards easing the operational aspects of deploying advanced network services with OSM, there is now a collection of **enhanced troubleshooting capabilities**, with **advanced logging** and access to **VNF console** whenever needed.

As a further effort to enrich the user experience, a large number of **user interface improvements**, which simplify the overall operation and improve user experience, have been completed.

OSM Release ONE Details

This section presents a lower-level view focusing on some of the specifics of the advancements in the OSM Release ONE codebase.

On-boarding experience & VNF Packaging

Cloud Init Support

OSM has added to Release ONE tighter controls on the use of "cloud-init" procedures, commonly used in cloud environments. "cloud-init" is typically used for the initial configuration of virtual machines where the user can inject context information (configuration parameters and scripts) to configure the initial VM setup.

The new feature only permits the explicit injection of user-configured SSH keys into each VM from the NSD. No user-supplied cloud-init script is allowed. All other configuration may be performed via the VNF Configuration and Abstraction (VCA) component.

Create Networks at VIM

The ability to create a tenant network in a specific VIM from the OSM User Interface has been added to Release ONE. This is an important addition as the Network Service Descriptor (NSD) connection points need to be connected to an existing VIM network. As these networks can now be created from within OSM, the user no longer needs to directly access the VIM API to create the needed networks prior to NSD instantiation.

Two types of networks can now be created from UI or from SO:

1. Isolated tenant networks only visible to the user/tenant that creates it.
 - a. These networks can be used to connect one or several NSD connection points, allowing, for example, the joining of two NS instances.
2. External networks that in addition to 1) can be connected to the outside world and can be made visible to other users/tenants.

This functionality is only available to users with admin rights. The way the network is connected to the outside must be specified by the user/SO, and can be one of:

- A router to another "public" network where the User/SO provides this public network name to connect to, and the router parameters.
- A provider network (e.g. in the case of OpenStack) where the User/SO provides the needed tag that the networks uses.
- An external switch port (OpenVIM) where the User/SO provides the physical switch port name and the VLAN tag if any.
- A multitenant network to connect to where the User/SO provides the network name to connect to and the VLAN tag.

Remove NSD Datacenter Network Reference

In Release ZERO the implementation of a NS Descriptor (NSD) contained "external" networks, or connection points intended for connecting outside networks. Management networks are a typical example of these types of connections, but they could also have been data plane connections. These networks needed to exist at the VIM with the same name used for the NSD. This relationship between descriptor name and deployment is undesirable because it forced the need for defining the deployment of a particular NSD for each VIM. As a result, the user was required to be able to map what VIM network is used for a concrete NSD "external" network or connection point.

Release ONE has advanced this area as it is now possible to map an NSD "external" network name to an existing VIM network.

Error Messages and Logging

Detailed feedback when deploying and configuring Error Message from RO & VCA can be found [here](#) and [here](#) respectively.

This pair of advancements enable collecting detailed error message when deploying and configuring from RO/VCA and propagating the messages to the UI/Logs files. This is especially useful in the case of errors to simplify troubleshooting of issues when using OSM.

YAML Format Descriptor Display

The UI Composer/Catalog Manager is a web-driven application that uses JSON as the transport for all data between the browser and the backend. JSON, however, is not the most human-readable language and it can get difficult to keep track of parameters in NSDs, VNFDs and VLDs when viewing them in raw JSON format.

YAML is a human-friendly language and using it to display the descriptors in the User Interface (UI) makes it much easier for the user to understand the parameter details of the various descriptors.

The UI now displays descriptor details in YAML format when a user requests to see the raw descriptors that have either been saved to the backend or are being worked upon in the UI.

Enhance Visual Differentiation between NS Catalog and VNF Catalog

This enhancement adds visual differentiators to the UI at strategic locations so that the user can keep context between sessions while working on them.

It was implemented based on feedback for OSM Release ZERO from the OSM EUAG that it was difficult to recall the context of whether a user was working on a VNFD, VLD or NSD if the user paused working on them for a while. This was because there was no differentiation between the visual representation of a NSD, VNFD or VLD in the catalog manager.

Restructure Layout of Service Primitive Page

Based on feedback received for Release ZERO from the End User Advisory Group (EUAG) and the community, there are a set of enhancements to the service primitive page that make it less difficult to keep the context of the parameters when a service primitive is triggered between different sessions and users. The changes include:

- Display primitive jobs triggered in reverse order of trigger time (last one to be triggered appears on top).
- Display the parameters with which a service primitive was triggered.
- Display errors from VCA if provided for failed service primitive jobs.

Package Creation Command Line Utility

A new command line utility has been added to create NS/VNF packages.

The Utility allows the creation of the following:

- Create descriptor package from folders.
- Create descriptor for simple VNF/NS with the folder structure.
- VNF accepts name, flavor, number of interfaces, and interface type as inputs. This enables a convenient method of creating packages tailored for different deployment flavors.

EPA based Resource Allocation

For Release ONE, a significant advancement is the inclusion of the OpenVIM implementation as a reference VIM. This reference VIM has a considerable set of EPA-focused capabilities that are useful for demonstrating the capabilities of the Service Orchestration, Resource Orchestration and VNF Configuration capabilities in OSM to deploy truly high-performance VNFs.

(Network) Service Modelling

Juju-2.x

OSM Release ZERO used Juju 1.x for modelling VNFs.

In Juju 2.x, there are a number of capabilities that the OSM community enables with Release ONE:

- Multi-model controller
 - A single instance of Juju can manage multiple models. This means that the SO will be able to stand up multiple independent scenarios using a single Juju instance, rather than needing separate Juju controllers per scenario. This reduces the footprint of OSM and improves operations for production users.
- Workload status
 - VNF charms provide clear status and error messages to the SO to interpret and present to the user. With this capability, VNF providers will find it easier to express the status of their VNFs in a complex topology.
- LXD containers
 - The "local" provider, which is used for DevOps iteration, moves to LXD to offer much faster and more flexible local deployment operations. Therefore, VNF providers will have a more efficient test and development experience providing OSM VNFs.

- Multi-user controller
 - A single Juju controller can track multiple users to have independent access to particular models managed by that controller. To offer audit and ops access at each level of the stack (essential for a number of production deployment use cases) Juju 2.x enables this in a PCI-DSS compliant fashion.
- Resources
 - A charm can now specify a set of resources that will be delivered to the charm at runtime. Common examples would be the charm's preferred JVM or a Docker image. Juju 2.x mediates access to these resources for efficiency; the net result is an improved VNFD experience for VNF vendors.

Network Types in RO

For OSM Release ZERO, the network types in Resource Orchestrator described the actual enforcement done in the host bridges/switches and in the physical switches. This feature changed the network types in RO functions and databases to describe the functional behavior instead of the actual enforcement.

Today network types in RO are:

- `ptp`, for point-to-point networks (E-Line) enforced in an underlay switching infrastructure. VM interfaces must be either port pass-through or SR-IOV.
- `data`, for L2 networks (E-LAN) enforced in an underlay switching infrastructure. VM interfaces must be data plane interfaces (either port pass-through or SRIOV).
- `bridge`, for L2 networks (E-LAN) enforced in a Linux Bridge or OVS-based virtual switch. VM interfaces are based on virtio interfaces.

The new network types enabled in the RO are:

- "E-Line", for E-Line networks, independently on the kind of enforcement in place.
 - E-Line networks are L2 point-to-point networks which mean that whatever frames come from a network port go to the other network port, and vice versa, no matter which MAC addresses are in the frames.
- "E-LAN", for E-LAN networks, independently on the kind of enforcement in place.
 - E-LAN networks are L2 point-to-multipoint networks where the destination network port is determined by the destination MAC address.

Both kinds of networks have a parameter related to the type of enforcement (e.g. virtual networks or underlay) that can be specified by the NS builder or by the VNF provider to make the enforcement explicit. If specified, the specific network configuration will be done either via virtual networks (Linux bridges or virtual switches, depending on the technology used by VIM) or via underlay switching infrastructure. If unspecified, the enforcement will depend on the actual interfaces being connected. For the moment, with the current state of the art in VIMs, it is assumed that:

- An E-Line network cannot connect VM interfaces based on virtio, but can only connect data plane interfaces (port pass-through or SRIOV).
- An E-LAN network can only connect VM interfaces of the same type: either virtio interfaces or data plane interfaces.

As an example, if three virtio-based VM interfaces are connected in an NS or VNF to an E-LAN network, the enforcement will be done via virtual networks based on Linux bridges or OVS, depending on the specific technology used by the VIM, but they won't be connected via underlay switching. Alternatively, if two virtio-based VM interfaces and one SR-IOV are intended to be connected in an NS, this will be rejected by the RO at instantiation time. This behaviour could be changed in the future if the state of the art in VIMs evolve.

In addition, E-LAN networks can have extra parameters to be passed to the VIM during network creation such as L3 configuration parameters (e.g. DHCP enabling, DHCP range, FW enabling, FW rules to be applied, IPv4 subnet and mask, etc.). The internal data structures in RO allow room for potential addition of new parameters, but the specification of these extra parameters is out of the scope for Release ONE.

Allow IP Parameters for Networks

OSM Release ZERO deployments cannot define IP parameters for networks (ranges, DHCP settings DNS) at the VNFD and NSD. Although this can be configured after deployment, it is convenient to allow the deployment with these parameters already configured at the VIM.

The parameters that can now be configured for Release ONE include IPv4/IPv6, IP address range, DHCP enable and range, Gateway, DNS and security group. This extension only affects L3 networks and does not apply to E-LINE type networks.

The specific configuration can be defined at 3 levels:

- A VNFD can contain parameters for its internal networks.
- An NSD can override this parameters for each one of its networks.
- Finally, the end user can override this parameters at instantiation time.

Configuration/Service Primitive Model Enhancements

This change enhances the VNF level configuration primitive model to be consistent with the NS level service primitive model and also renames config-primitive under NSD as "service-primitive"

Both the VNFD and NSD are updated with these changes and there is no functional impact due to these changes other than the changes in the descriptor.

Multi-VIM

New VIM connector for VMware vCloud Director

As noted above, VMware's vCloud Director has been included in OSM Release ONE as one of the natively supported VIMs, thus significantly broadening the scope of environments on which OSM can be used to deploy automated end-to-end network services.

OpenVIM as Reference VIM with EPA Capabilities

Also as noted above, OpenVIM has been included in OSM Release ONE as one of the natively supported VIMs and as a reference VIM, particularly in the context of its advanced EPA capabilities. OpenVIM can be leveraged by OSM to demonstrate the deployment of VNFs with high performance and efficiency.

Datacenter Capabilities

EPA and other capabilities differ from one datacenter to another depending on the type of VIM and its supported capabilities. For instance, OpenVIM supports EPA capabilities (e.g. use of hugepages memory, pinning to HW threads or cores, simultaneous support of SR-IOV and passthrough interfaces, E-Line and E-LAN L2 networks, configuration of an underlay switch for inter-VM communication), but it does not support L3 parameters in E-LAN networks. On the other hand, OpenStack supports a subset of EPA capabilities, while it offers several L3 parameters in E-LAN networks based on OVS.

It is desirable that some actions towards the VIM (e.g. deployment of a VM with data plane interfaces) could be rejected at RO level based on the declared capabilities of the VIM, without the need to run the long process of interacting to the VIM, running the associated task in VIM, and monitoring the status in VIM (e.g. VM in active state) to know if there was an error at deployment time. Moreover, in some cases the actions in the VIM do not lead to errors detected at VIM, but to sub-optimal performance that can only be detected inside the VNF (e.g. an inappropriate pinning leading to inadequate performance). The rejection at RO level based on VIM capabilities would also avoid these sub-optimal situations.

This new feature includes the ability to collect the different relevant capabilities of each type of VIM. When registering a new datacenter in the RO (datacenter creation primitive), these capabilities will be loaded in the RO database as part of the VIM information.

Support for VIM Accounts

The cloud-account in the User Interface has been renamed as VIM account to be consistent with ETSI terminology. While this is purely a cosmetic change in Release ONE it is another indicator of OSMs goal to align closely with the ETSI NFV MANO specifications.

Multi-Site

The implementation in Release ZERO allowed the deployment of a Network Service instance in a single datacenter chosen by the user.

This feature enables supporting NS instances with VNFs running in multiple different datacenters. The feature assumes that the networks interconnecting several datacenters are pre-provisioned.

There are several use cases where a multi-site NS applies, e.g.

- The deployment of a full core network with several Provider Edge (PE) Routers in different datacenters.
- The deployment of VoIP network where some elements are centralized in a datacenter while others are distributed (e.g. SBC).



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

This White Paper is issued for information only. It does not constitute an official or agreed position of ETSI, nor of its Members. The views expressed are entirely those of the author(s).

ETSI declines all responsibility for any errors and any loss or damage resulting from use of the contents of this White Paper.

ETSI also declines responsibility for any infringement of any third party's Intellectual Property Rights (IPR), but will be pleased to acknowledge any IPR and correct any infringement of which it is advised.

Copyright Notification

Copying or reproduction in whole is permitted if the copy is complete and unchanged (including this copyright statement).

© European Telecommunications Standards Institute 2016. All rights reserved.

DECT™, PLUGTESTS™, UMTS™, TIPHON™, IMS™, INTEROPOLIS™, FORAPOLIS™, and the TIPHON and ETSI logos are Trade Marks of ETSI registered for the benefit of its Members.

3GPP™ and LTE™ are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

GSM™, the Global System for Mobile communication, is a registered Trade Mark of the GSM Association.