ETSI GR NFV-IFA 035 V5.1.1 (2023-10)

GROUP REPORT

Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on network connectivity integration and operationalization for NFV

Disclaimer

The present document has been produced and approved by the Network Functions Virtualisation (NFV) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference

DGR/NFV-IFA035

Keywords

API, management, MANO, NFV

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from:
https://www.etsi.org/standards-search

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format at www.etsi.org/deliver.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx

If you find errors in the present document, please send your comment to one of the following services:
https://portal.etsi.org/People/CommiteeSupportStaff.aspx

If you find a security vulnerability in the present document, please report it through our Coordinated Vulnerability Disclosure Program:
https://www.etsi.org/standards/coordinated-vulnerability-disclosure

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.
No recommendation as to products and services or vendors is made or should be implied.
No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.
In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.
The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2023.
All rights reserved.

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™** and **LTE™** are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**® and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This Group Report (GR) has been produced by ETSI Industry Specification Group (ISG) Network Functions Virtualisation (NFV).

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document studies inter-site and intra-site connectivity technologies, and how these can be used to realize the intra-site and inter-site connectivity services and fulfil networking requirements regarding VNF connectivity within the scope of NFV framework.

The present document also analyses the intra-site interactions on the [NF-Vi]/N reference point, network provisioning across network domain boundaries, OAM aspects for intra and inter-site connectivity, overlay and underlay network provisioning, and management and operations for connectivity services with recent traffic engineering technologies.

Finally, based on the analysis the present document provides a set of recommendations for enabling intra- and inter-site connectivity services within the NFV framework.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long-term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ETSI GR NFV 003 (V1.5.1): "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[i.2] ETSI GS NFV-INF 005 (V1.1.1): "Network Functions Virtualisation (NFV); Infrastructure; Network Domain".

[i.3] ETSI GS NFV-EVE 005: "Network Functions Virtualisation (NFV); Ecosystem; Report on SDN Usage in NFV Architectural Framework".

[i.4] ETSI GR NFV-IFA 022 (V3.1.1): "Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Report on Management and Connectivity for Multi-Site Services".

[i.5] ETSI GS NFV-EVE 003: "Report on NFVI Node Physical Architecture Guidelines".

[i.6] IETF RFC 7938: "Use of BGP for Routing in Large-Scale Data Centers.

[i.7] ETSI GS NFV-IFA 005 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Or-Vi reference point - Interface and Information Model Specification".

[i.8] IETF RFC 6241: "Network Configuration Protocol (NETCONF)".

[i.9] IETF RFC 7047: "The Open vSwitch Database Management Protocol".

[i.10] ONF TS-025: "OpenFlow Switch Specification Version 1.5.1 (Protocol version 0x06)".

[i.11] IETF RFC 4271: "A Border Gateway Protocol 4 (BGP-4)".

[i.12] IETF RFC 5440: "Path Computation Element (PCE) Communication Protocol (PCEP)".

[i.13] OpenDayLight

[i.14] Tungsten Fabric.

[i.15] IEEE™ 802.1Q-2018: "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks", 10.1109/IEEESTD.2018.8403927.

[i.16] IETF RFC 8014: "An Architecture for Data-Center Network Virtualization over Layer 3 (NVO3)".

[i.17] IETF RFC 7348: "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks".

[i.18] IETF RFC 7637: "NVGRE: Network Virtualization Using Generic Routing Encapsulation".

[i.19] IETF RFC 4023: "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)".

[i.20] IETF RFC 8299: "YANG Data Model for L3VPN Service Delivery".

[i.21] IETF RFC 9182: "A YANG Network Data Model for Layer 3 VPNs".

[i.22] ETSI GR NFV-SOL 017: "Network Functions Virtualisation (NFV) Release 3: Protocols and Data Models Report on protocol and data model solutions for Multi-site Connectivity Services".

[i.23] IETF RFC 4364: "BGP/MPLS IP Virtual Private Networks (VPNs)".

[i.24] IETF RFC 8466: "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery".

[i.25] IETF RFC 9291: "A YANG Network Data Model for Layer 2 VPNs".

[i.26] Technical Specification MEF 6.2: "EVC Ethernet Services Definitions Phase 3", August 2014.

[i.27] ETSI GS NFV-IFA 014: "Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Network Service Templates Specification".

[i.28] IETF RFC 5921: "A Framework for MPLS in Transport Networks".

[i.29] Recommendation ITU-T G.709/Y.1331: "Interfaces for the optical transport network".

[i.30] IETF RFC 7364: "Problem Statement: Overlays for Network Virtualization".

[i.31] IETF RFC 7365: "Framework for Data Center (DC) Network Virtualization".

[i.32] Technical Specification MEF 6.3: "Subscriber Ethernet Service Definitions", November 2019.

[i.33] ETSI GS NFV-SOL 005: "Network Functions Virtualisation (NFV) Release 3; Protocols and Data Models; RESTful protocols specification for the Os-Ma-nfvo Reference Point".

[i.34] ETSI GS NFV-IFA 032: "Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; Interface and Information Model Specification for Multi-Site Connectivity Services".

[i.35] IETF RFC 2890: "Key and Sequence Number Extensions to GRE".

[i.36] IETF RFC 4664: "Framework for Layer 2 Virtual Private Networks (L2VPNs)".

[i.37] IETF RFC 2764: "A Framework for IP Based Virtual Private Networks".

[i.38] IETF RFC 4761: "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling".

[i.39] IETF RFC 4762: "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling".

[i.40] IETF RFC 4448: "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling".

[i.41] IETF RFC 6071: "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap".

[i.42] IETF RFC 7296: "Internet Key Exchange Protocol Version 2 (IKEv2)". .

[i.43] IETF RFC 4302: "IP Authentication Header".

[i.44] IETF RFC 4303: "IP Encapsulating Security Payload (ESP)".

[i.45] IETF RFC 7209: "Requirements for Ethernet VPN (EVPN)".

[i.46] IETF RFC 7432: "BGP MPLS-Based Ethernet VPN".

[i.47] IETF RFC 4760: "Multiprotocol Extensions for BGP-4".

[i.48] IETF RFC 8277: "Using BGP to Bind MPLS Labels to Address Prefixes".

[i.49] IETF RFC 8430: "RIB Information Model".

[i.50] IETF RFC 1655: "Application of the Border Gateway Protocol in the Internet".

[i.51] IETF RFC 4456: "Route Reflection: An Alternative to Full Mesh Internal BGP (iBGP)".

[i.52] IETF RFC 5065: "Autonomous System Confederation for BGP".

[i.53] IETF RFC 7911: "Advertisement of Multiple Paths in BGP".

[i.54] IETF RFC 5492: "Capabilities Advertisement with BGP-4".

[i.55] IETF RFC 8365: "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)". .

[i.56] IETF RFC 8402: "Segment Routing Architecture".

[i.57] IETF RFC 9256: "Segment Routing Policy Architecture".

[i.58] IETF RFC 8754: "IPv6 Segment Routing Header (SRH)".

[i.59] IETF RFC 8986: "Segment Routing over IPv6 (SRv6) Network Programming".

[i.60] ETSI GS NFV 006 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Architectural Framework Specification".

[i.61] ETSI GR NFV-IFA 038 (V4.1.1): "Network Functions Virtualisation (NFV) Release 4; Architectural Framework; Report on network connectivity for container-based VNF".

[i.62] ETSI GR NFV-IFA 043 (V0.0.4): "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on enhanced container networking".

[i.63] ETSI OSG TeraFlowSDN.

[i.64] ETSI GR NFV-EVE 022 (V5.1.1): "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on VNF configuration".

[i.65] ONF TAPI (v2.4.0): "ONF Transport API SDK 2.4.0", December 2022.

[i.66] DMTF Redfish DSP2046 (2021.4): "Redfish Resource and Schema Guide", published 2021-12-02.

[i.67] DMTF Redfish DSP0266 (1.15): "Redfish Specification", published 2021-12-02.

[i.68] ETSI GS NFV-IFA 008 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Ve-Vnfm reference point - Interface and Information Model Specification".

[i.69] ETSI GS NFV-IFA 006 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; Vi-Vnfm reference point - Interface and Information Model Specification".

[i.70] ETSI GS NFV-IFA 011 (V4.4.1): "Network Functions Virtualisation (NFV) Release 4; Management and Orchestration; VNF Descriptor and Packaging Specification".

[i.71]	ETSI GR NFV-IFA 046 (V5.1.1): "Network Functions Virtualisation (NFV) Release 5; Architectural Framework; Report on NFV support for virtualisation of RAN".
[i.72]	ETSI GS NFV-IFA 053 (V0.1.0): "Network Functions Virtualisation (NFV) Release 5; Management and Orchestration; Requirements and interface specification for Physical Infrastructure Management".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in ETSI GR NFV 003 [i.1], ETSI GS NFV-IFA 032 [i.34] and the following apply:

inter-site connectivity service: connectivity service enabling network connectivity among two or multiple NFVI-PoPs

intra-site connectivity service: connectivity service enabling network connectivity within the same NFVI-PoP

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in ETSI GR NFV 003 [i.1] and the following apply:

AC	Attachment Circuit
AH	Authentication Header
BGP	Border Gateway Protocol
CE	Customer Edge
DEI	Drop Eligible Indicator
eBGP	exterior BGP
EPL	Ethernet Private Line
EP-LAN	Ethernet Private LAN
EP-Tree	Ethernet Private Tree
EVC	Ethernet Virtual Connection
ES	Ethernet Segment
ESI	Ethernet Segment Identifier
ESP	Encapsulation Security Payload
EVPN	Ethernet Virtual Private Network
GRE	Generic Routing Encapsulation
ICV	Integrity Check Value
iBGP	interior BGP
IGP	Interior Gateway Protocol
LDP	Label Distribution Protocol
LER	Label Edge Router
LSP	Label Switching Router
MP-BGP	Multi-Protocol BGP
NLRI	Network Layer Reachability Information
NVA	Network Virtualization Authority
NVE	Network Virtualization Edge
NVGRE	Network Virtualization using Generic Routing Encapsulation
NVO3	Network Virtualization over Layer 3
NSP	Native Service Processing
OAM&P	Operations, Administration, Management and Provisioning
PCP	Priority Code Point
P	P-node (Provider node)
PE	Provider Edge

PSN	Packet Switched Network
RD	Route Distinguisher
RIB	Routing Information Base
RR	Route Reflectors
RT	Route Target
SA	Security Association
SBI	Service-Based Interface
SP	Service Provider
SR	Segment Routing
SPI	Security Parameter Index
TPID	Tag Protocol Identifier
UNI	User Network Interface
VSID	Virtual Subnet Identifier
VPLS	Virtual Private LAN service
VPRN	Virtual Private Routed Network
VRF	Virtual Routing and Forwarding
VTEP	Virtual Tunnel Endpoint
VXLAN	Virtual eXtensible Local Area Network

4 Overview of network connectivity technologies in NFV

4.1 Introduction

4.1.1 General overview

Several connectivity technologies, specified by the IETF or other standards organizations, can be utilized in the NFV context to support network connectivity among multiple virtualized entities. Whereas in this clause a general overview of connectivity technologies is given, the continuation of the present document analyses more in detail the use of such technologies with regards to the relevant NFV reference points and interfaces. A general overview is also given on routing and traffic engineering technologies that can be used as facilitators to establish the connectivity services (clause 4.4), as well as on the use of Software Defined Networking (SDN) in the NFV architectural framework (clause 4.5).

In the context of the present document, connectivity services are mainly divided into two categories: intra-site connectivity services and inter-site connectivity services.

The connectivity technologies that are introduced and analysed in the present document can be used for one of the two categories or for both, as the two categories are not meant to be mutually exclusive.

4.1.2 VPN

A Virtual Private Network (VPN) is a virtual network used to connect two or multiple networks over a common network infrastructure (such as public Internet). A VPN configuration can expand a private network across geographically disparate locations or can connect a group of devices (such as VMs) using their local private networks operating on top of a shared network fabric inside the same site, e.g. NFVI-PoP.

There are different mechanisms and technologies available to offer a VPN connection within one single site or over multiple sites. The connectivity technologies described and analysed in the present document aim at enabling a VPN connection between VNFs either within the same NFVI-PoP or across multiple NFVI-PoPs. Each technology described has specific requirements and characteristics to be considered, and therefore might be particularly suitable for specific scenarios and configuration setup.

4.1.3 NFVI-PoP infrastructure architecture and virtual networks

In ETSI-NFV, a virtual link [i.1] is defined as a set of connection points along with the connectivity relationship between them and any associated target performance metrics (e.g. bandwidth, latency, QoS), and is supported by a virtual network inside the NFVI. The Infrastructure Network domain was studied in ETSI GS NFV-INF 005 [i.2], however, information regarding the following details have been limited in relevant ETSI NFV specifications:

- how the NFV network infrastructure is composed,
- how the virtualised network is instantiated within the network infrastructure,
- how the virtualised network is interconnected with physical networks.

The NFVI Node physical architecture was studied in ETSI GS NFV-EVE 003 [i.5], including computing, storage and network resources that provide processing, storage and connectivity to VNFs through the virtualisation layer (e.g. hypervisor). The connectivity is typically supported by configurations in a multi-tiered network topology in a single NFVI-PoP where the network nodes are categorized as follows:

- a network node, which hosts compute and storage nodes, and connects to other network nodes (e.g. Top of Rack (ToR) switch, access switch, leaf switch);
- a network node, which interconnects other network nodes (e.g. aggregation switch, spine switch);
- a network node, which connects the NFVI-PoP to transport network (e.g. gateway router).

The network infrastructure that characterizes a NFVI-PoP is similar to a generic multi-tiered data centre network architecture described in [i.6]. A typical multi-tiered network topology often met in practise is the "Leaf and Spine network", where the top tier (i.e. Tier 1) node of this multi-tiered network infrastructure is a data centre gateway node that provides connectivity to external networks (e.g. Internet) over a Wide Area Network (WAN). The tier 2 node forms the spine switches and tier 3 node(s) form(s) the leaf switch(es).

4.2 Summary of Intra-site Connectivity aspects in NFV

4.2.1 Overview of Virtualised Network Resources in NFV

ETSI GS NFV-IFA 005 [i.7] specifies the *Allocate Virtualised Network Resource* operation for the *Virtualized Network Resources Management* interface produced by the VIM on the Or-Vi reference point, in order to instantiate intra-site connectivity inside an NFVI-PoP. Specifically, the information about the virtual network resource to be created is specified by the *typeNetworkData* input parameter, which specifies. amongst others. attributes such as *networkType, segmentationId layer3Attributes* and *isShared*.

These attributes characterize the intra-site connectivity, for instance:

- The *networkType* attribute specifies the type of network that maps to the virtualised network, such as "vlan", "vxlan", "gre".
- The virtualised networking technologies have an inherent feature that isolates virtualised network resources from each other. The *segmentationId*attribute enables the identification of the respective isolated virtualised network resources, by using a value of vlan, vxlan and gre identifier, etc., depending on the type of virtualised network. If *segmentationId*attribute is not present, then it characterizes a flat network. For example this is a type of virtualised network resource in which the network traffic is handled without vlan tags.
- The *layer3Attributes* attribute enables setting up a network providing defined layer 3 connectivity. This attribute specifies further sub-attributes, like *ipVersion*, *gatewayIp*, *isDhcpEnabled*, further specifying whether the virtualized network is IPv4 or IPv6, the IP address of the gateway, and whether a DHCP service is enabled or not.
- The *isShared* attribute defines whether the virtualised network is shared among network service users or not.

EXAMPLE: In the case of OpenStack® as described in Annex A, there are two types of virtualised network resource instantiations:

- **Tenant (or self-service) network:** A virtualised network managed in an OpenStack project and isolated from virtualised networks in other OpenStack projects.
- **Provider network:** A virtualised network managed by an administrator, and which can be exposed to other networks.

END OF EXAMPLE

NOTE: The OpenStack® Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. ETSI is not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

4.2.2 Network isolation technologies

Virtualised network resources can be isolated with VLAN-based solutions and/or encapsulation technologies.

Virtual Local Area Network (VLAN) is a well-known technology specified as IEEE 802.1Q [i.15]. The VLAN networks share the same physical network infrastructure, but they are isolated from each other in layer 2. The individual VLAN networks can be identified with the VLAN identifier. The valid VLAN identifiers range between 0 and 4 095.

Another technique for enabling network isolation is by using encapsulation technologies. L2-frames/L3-packets belonging to a virtualised network are forwarded by encapsulating them with an outer packet header complying with the underlay network infrastructure protocol. As described in IETF RFC 8014 [i.16], the frames/packets are identified with an encapsulation identifier that identifies the virtualised network to which they belong. As examples, Virtual eXtensible Local Area Network (VXLAN) [i.17] has a 24-bit VXLAN Network Identifier (VNI), Network Virtualization using Generic Routing Encapsulation (NVGRE) [i.18] has a 24-bit Tenant Network ID (TNI) and MPLS-over-GRE [i.19] provides a 20-bit label field, which are used for identifying the virtualized network of the respective L2-frame/L3-packet. These encapsulation technologies enable using a broader range of identifiers for virtualised network resources than the VLAN technology.

4.2.3 Routers/ Gateways for interconnecting virtualised network resources

In the *Allocate Virtualised Network Resource* operation defined in ETSI GS NFV-IFA 005 [i.7], a subnet gateway can be specified in the *NetworkSubnetData* information element. The subnet gateway terminates a network and relays traffic to other networks. In the context of the NFVI-PoP environment, the following types of gateways can be considered:

- A gateway to connect an internal virtualised network resource with another internal virtualised network resource.
- A gateway to connect an internal virtualised network resource with an external virtualised network resource, inside the NFVI-PoP.
- A gateway to connect an external virtualised network resource with networks external to the NFVI-PoP.

EXAMPLE: In an OpenStack environment, for example, the gateways between a pair of internal virtualised network resources, or between a pair of internal and external virtualised network resources are instantiated as a router, which is an OpenStack networking service, on an NFVI compute node. The router to connect an internal virtualised network resource with an external virtualised network resource has a network interface to the NFV physical network infrastructure. On the other hand, the gateway between an external virtualised network resource and a network external to the NFVI-PoP is instantiated at a networking node hosting the external virtualised network resources (see provider networks in Annex A). In the NFVI-PoP environment, this networking node is known as an NFVI-PoP gateway node connecting to other networks such as WAN. The gateway is specified when an external virtualised network resource is allocated in the *Allocate Virtualised Network Resource* operation.

4.3 Summary of Inter-site Connectivity aspects in NFV

4.3.1 Connectivity service for L3VPN Service

As described in ETSI GS NFV-IFA 032 [i.34] and ETSI GR NFV-SOL 017 [i.22], L3VPN technology can be used to support a Multi-Site Connectivity Service (MSCS). The connectivity service provides capabilities to exchange routing information and support encapsulation of data plane packets for traffic isolation over the backbone network (e.g. Wide Area Network). Besides, the connectivity service provides route isolation on a shared network infrastructure, even if different VPNs use an overlapped address space. As shown in Figure 4.3.1-1, L3VPN service is facing towards the service user, and connectivity service is represented by a L3VPN network model.

The YANG model for L3VPN service delivery is described in [i.20]. The L3VPN service model is described by *VPN service information* and *site information* on where the customer wants to establish the L3VPN service. In more detail:

- *VPN service information* includes parameters such as VPN service type identification, VPN service topology (e.g. any-to-any, hub-spoke), and so on.

- *site information* includes parameters to interface targeted sites with the L3VPN service. For example, *site-network-access* is characterized by bearer parameters (e.g. Ethernet, DSL, Wireless), connection parameters (ipv4 or ipv6), routing protocols (e.g. bgp, ospf) and so on.

The instantiation of the L3VPN service is described in [i.21]. Provider Edge (PE) devices at the edge of the backbone network are network devices connecting to Customer Edge (CE) devices. In NFV context, a CE device is considered as an NFVI-PoP gateway, which terminates external virtualised network resources defined in the context of intra-site connectivity. In the context of L3VPN network model in [i.21], the *VPN Network Access* represents the network interfaces to host the requested attachment circuits in the L3VPN service model. The *VPN Network Access* is associated to a given *VPN Node*, which is mapped to a Virtual Routing and Forwarding (VRF) instance. In this way, the requested *site-network-access* between CEs and PEs is terminated by the *VPN Node* in the PE device. PE devices may host multiple *VPN Nodes* connecting to different NFVI-PoP gateways. Within the backbone network, traffic engineered networks (e.g. MPLS, Segment Routing, Optical Transport Network) can be used to support the L3VPN service, and to offer capabilities such as path computations for a sequence of router hops, guaranteed bandwidth, path protection, etc.

In BGP/MPLS IP L3VPN, as specified in [i.23], for example, Peered PE devices exchange VPN labels as Network Layer Reachability Information (NLRI) used in MP-BGP. The VPN label is put in data plane packets and is used to determine a destination VPN Node connected to the targeted NFVI-PoP Gateway.

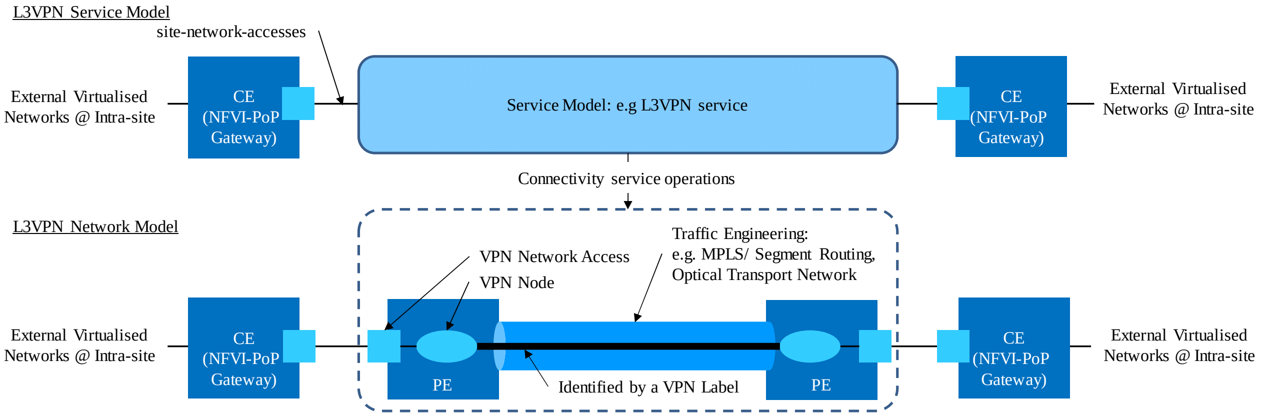
L3VPN Service Model
site-network-accesses
External Virtualised Networks @ Intra-site
CE (NFVI-PoP Gateway)
Service Model: e.g L3VPN service
CE (NFVI-PoP Gateway)
External Virtualised Networks @ Intra-site
Connectivity service operations
L3VPN Network Model
VPN Network Access
VPN Node
Traffic Engineering: e.g. MPLS/ Segment Routing, Optical Transport Network
External Virtualised Networks @ Intra-site
CE (NFVI-PoP Gateway)
PE
Identified by a VPN Label
PE
CE (NFVI-PoP Gateway)
External Virtualised Networks @ Intra-site

Figure 4.3.1-1: L3VPN service and Connectivity Service

4.3.2 Connectivity service for L2VPN Service

A Layer 2 VPN (L2VPN) service, described in ETSI GS NFV-IFA 032 [i.34] and ETSI GR NFV-SOL 017 [i.22], offers another MSCS for interconnecting multiple sites to exchange L2 traffic over a shared network infrastructure (e.g. WAN). Multiple L2VPN service types are introduced in the L2VPN service model in [i.24].

As shown in Figure 4.3.2-1, the YANG L2VPN service delivery model in [i.24] provides the customer view of the service and focuses on communication between the customers and network operators. In an NFV framework, a customer is an NFV-MANO service user who expects connectivity to interconnect resources in multiple NFVI-PoPs. At this abstraction level, the L2VPN service model describes the VPN service information and site information where the customer wants to establish the L2VPN service. In more detail:

- *vpn-service* contains a list of VPN service types (e.g. Pseudowire service, Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), EVPN service), the VPN service topology (e.g. any-to-any, hub-spoke), and so on.
- *site* includes a list of properties that are used to interface targeted sites with the L2VPN service. For example, *site-network-access* is characterized by bearer parameters (e.g. Ethernet, DSL, Wireless), connection parameters for the Layer 2 protocol parameters of the attachment (e.g. encapsulation-type, untagged/ tagged-interface, vlan/ circuit-id) and so on.

The YANG Network Data Model for Layer 2 VPNs is described in [i.25]. The PE devices at the edge of the backbone network are network devices connecting to CE devices , which in the NFV context are NFVI-PoP gateways. In the context of L2VPN network model [i.25], *vpn-network-accesses* represents the network interfaces to host the requested bearers in the L2VPN service model. *vpn-network-accesses* is associated to a given *vpn-node*, which is an instance for the L2VPN service. In this way, the requested *site-network-access* between CEs and PEs is terminated by *vpn-node* in the PE device. PE devices may host multiple *vpn-nodes* connecting to different NFVI-PoP gateways. Within the backbone network, traffic engineered networks (e.g. RSVP-TE, Segment Routing) or overlay networks (e.g. GRE, VXLAN) can be used to support the L2VPN service. Data plane packets are encapsulated with an identifier (e.g. VPN label, GRE/ VXLAN ID) to isolate routes from different NFVI-PoP gateways.

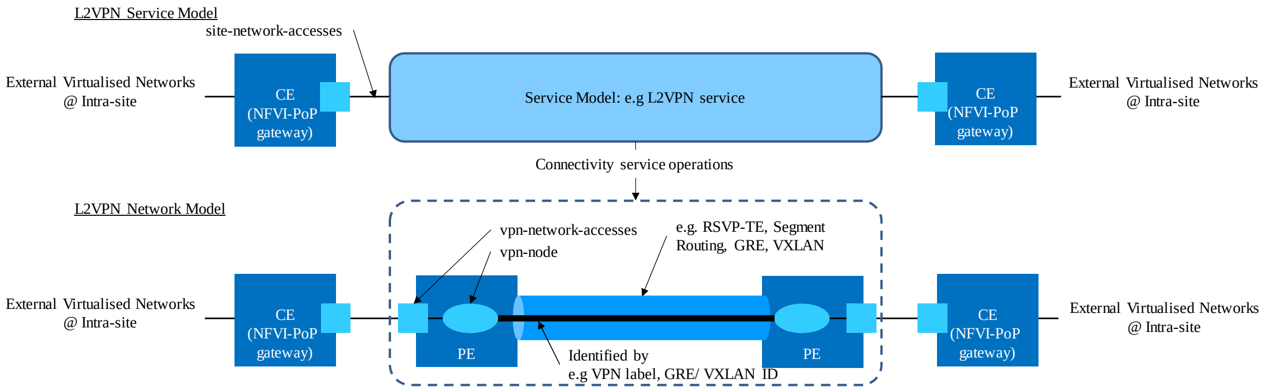
L2VPN Service Model
site-network-accesses
External Virtualised Networks @ Intra-site
CE (NFVI-PoP gateway)
Service Model: e.g L2VPN service
CE (NFVI-PoP gateway)
External Virtualised Networks @ Intra-site
Connectivity service operations
L2VPN Network Model
vpn-network-accesses
vpn-node
e.g. RSVP-TE, Segment Routing, GRE, VXLAN
External Virtualised Networks @ Intra-site
CE (NFVI-PoP gateway)
PE
Identified by e.g VPN label, GRE/ VXLAN ID
PE
CE (NFVI-PoP gateway)
External Virtualised Networks @ Intra-site

Figure 4.3.2-1: L2VPN service and Connectivity Service

4.3.3 Overview of Carrier Ethernet Service

4.3.3.1 Overview

ETSI GS NFV-IFA 014 [i.27] lists Carrier Ethernet Services specified in [i.26] as examples of connectivity referenced by a Virtual Link Descriptor. In the MEF context, User Network Interfaces (UNI) are defined between customer (e.g. VLAN) and service provider networks (e.g. WAN), and an Ethernet Virtual Connection (EVC) is defined as an association of two or more UNIs that enables the exchange of Ethernet frames. Ethernet Line (E-Line) service is based upon a Point-to-Point EVC, Ethernet LAN (E-LAN) service is based upon a Multipoint-to-Multipoint EVC, and Ethernet Tree (E-Tree) service is based upon a Rooted-Multipoint EVC.

Various transport network technologies are considered to support Carrier Ethernet Services. As discussed in clause 4.3.2, VPWS and VPLS are examples of MPLS-based transport technologies. MPLS Transport Profile (MPLS-TP) [i.28] is also introduced as another option.

In optical transport technologies, Ethernet over OTN [i.29], Ethernet over WDM and Ethernet over SONET/SDH are considered for supporting Carrier Ethernet Services. Ethernet over Direct Fibre is also an alternative.

In IEEE 802.1-based transport technologies, Provider Bridged (PB) Networks, Provider Backbone Bridged (PBB) Networks and Provider Backbone Bridged with Traffic Engineering (PBB-TE) Networks are listed as alternative transport network technologies.

4.3.3.2 Ethernet Private Line (EPL) service

The E-Line service provides a private data connection over a public network. It is enabled by connecting two Ethernet ports, thus creating a point-to-point EVC over a public network [i.29] and [i.32]. This service emulates a classical dedicated network in which a fixed bandwidth is reserved and can be extended accordingly, based on a Bandwidth on Demand (BoD) mechanism.

As shown in Figure 4.3.3.2-1, different EVC connections can be configured between pairs of PEs each, while each CE connects to the correspondent PE through an UNI. To each UNI, only one EVC can be configured. Therefore, in the classical EPL configuration, a PE managing more EVCs needs to configure different UNIs for each EVC.

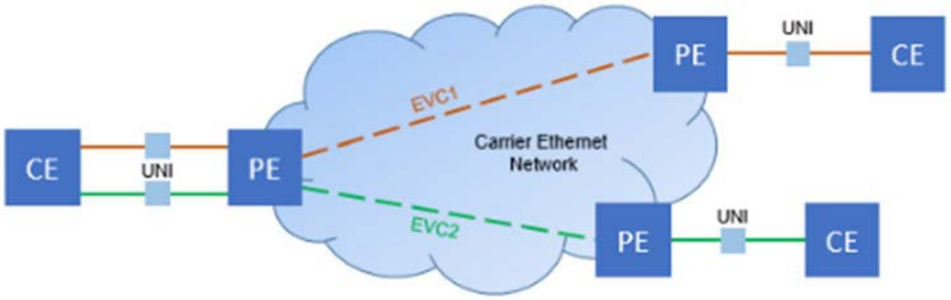

Figure 4.3.3.2-1: Point-to-point EVCs using EPL services

A particular version of EPL is the Ethernet Virtual Private Line (EVPL), in which the EVC is identified using a VLAN tag. Different EVCs can therefore be mapped to a unique UNI through the mapping of a VLAN ID (i.e. CE-VLAN ID) to each EVC, configured by the CE. This scenario is illustrated in Figure 4.3.3.2-2.

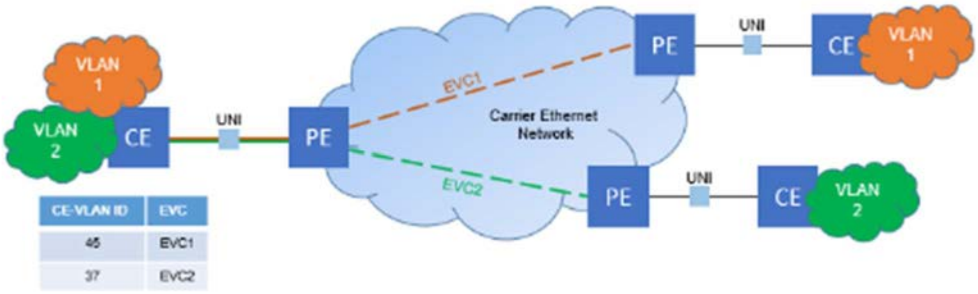

CE-VLAN ID	EVC
45	EVC1
37	EVC2

Figure 4.3.3.2-2: Point-to-point EVCs using EVPL services

4.3.3.3 Ethernet Private LAN (EP-LAN) service

The E-LAN service instantiates a multipoint-to-multipoint EVC between UNIs, with defined bandwidth requirements and a Class of Service (CoS) for traffic differentiation [i.29] and [i.32]. Figure 4.3.3.3-1 shows the multipoint-to-multipoint virtual connection, which is established in the Carrier Ethernet Network, allowing to connect multiple PEs.

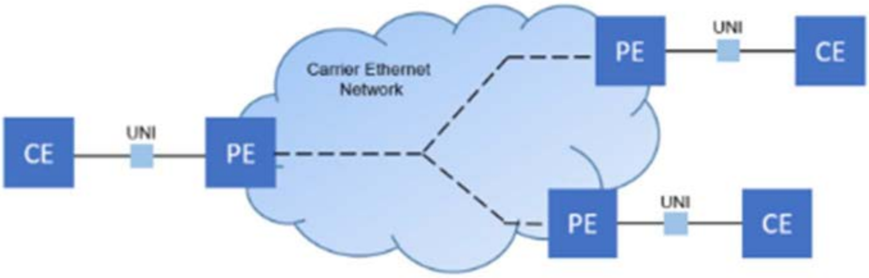

Figure 4.3.3.3-1: Multipoint-to-multipoint EVC using EP-LAN services

A particular version of EP-LAN in which VLAN is used to tag traffic and to map UNIs to a specific EVC is the Ethernet Virtual Private LAN (EVP-LAN) service.

4.3.3.4 Ethernet Private Tree (EP-Tree) service

The E-Tree service instantiates a rooted-multipoint EVC connection. An example of this virtual connection is depicted in Figure 4.3.3.4-1, where one root PE and multiple leaf PEs are represented. From the point of view of the root PE, the virtual connection established over the Carrier Ethernet Network is a multipoint connection, while from the point of view of the leaf PEs, it is as a point-to-point virtual connection.

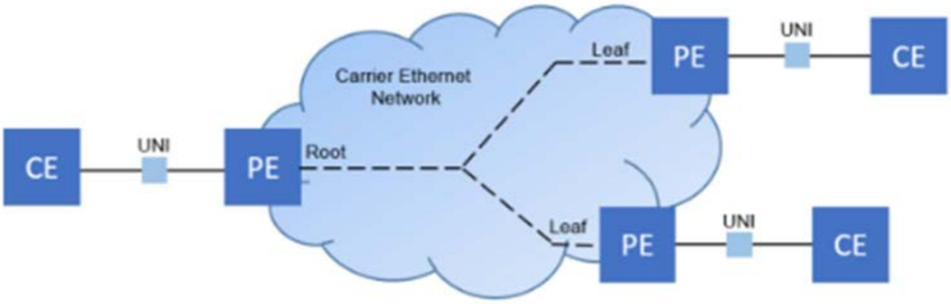

Figure 4.3.3.4-1: Rooted-multipoint EVC using EP-Tree services

4.4 Routing and Traffic Engineering technologies

4.4.1 Border Gateway Protocol for connectivity technologies

The Border Gateway Protocol (BGP) is an IETF-standardized protocol used to exchange routing and reachability information within and between Autonomous Systems (ASs). It makes decision based on the path and on networking policies and/or rules defined by the network administrator. Latest BGP version 4 is specified in IETF RFC 4271 [i.11], while a multiprotocol extension of BGP, also known as Multiprotocol BGP (MP-BGP) is specified in IETF RFC 4760 [i.47] allows to carry routing information for multiple network protocols beyond IPv4 (e.g. IPv6, IPX).

BGP runs between two or more BGP peers. When BGP is used between peers in a same AS, it is referred to as interior BGP (iBGP) and the peers are called iBGP peers. Exterior BGP (eBGP) is used between peers belonging to different ASs, and the peers are referred to as eBGP peers. As specified in IETF RFC 7938 [i.6], eBGP can be used as well in a large-scale data center design.

In the context of connectivity technologies, BGP can be used to exchange routing and reachability information in both intra-site and inter-site domains for different connectivity technologies. For example, in BGP can be used by the Service Provider to exchange the routes of a particular VPN among the different PE routers attached to the VPN [i.23], eventually allowing for distribution of MPLS labels for the route [i.48]. BGP can also be used in VXLAN-based solutions for discovery of the NVEs belonging to a specific VXLAN. IETF RFC 4761 [i.38] also proposes the use of BGP in VPLS for auto-discovery and signalling purposes, whereas MP-BGP [i.47] is used in EVPN for exchanging reachability information between PE routers.

4.4.2 Overview of Multiprotocol Label Switching Virtual Private Network

Multiprotocol Label Switching (MPLS) is a connection-oriented forwarding protocol that can be used to create VPNs for inter-site multipoint-to-multipoint connectivity through a provider's network [i.23].

Figure 4.4.2-1 shows the format of the MPLS packet and the structure of the MPLS label. As depicted, an MPLS label is encapsulated between the L2 header and the L3 header and it can be a multi-layer label referred to as a MPLS label stack. Theoretically, there is no limit to MPLS label nesting, i.e. number of inner labels. This feature makes MPLS a suitable technology for enabling VPN services that use multi-layer encapsulation of public network labels and private network labels. An MPLS label is 4 bytes, with a 20-bit label value, a 3-bit Traffic Class field for TE and QoS purposes, and an S-flag indicating if the label is at the bottom of the stack i.e. the last inner label. It also has an 8-bit Time-to-Live (TTL) field which is similar to the TTL field of the IP packet.

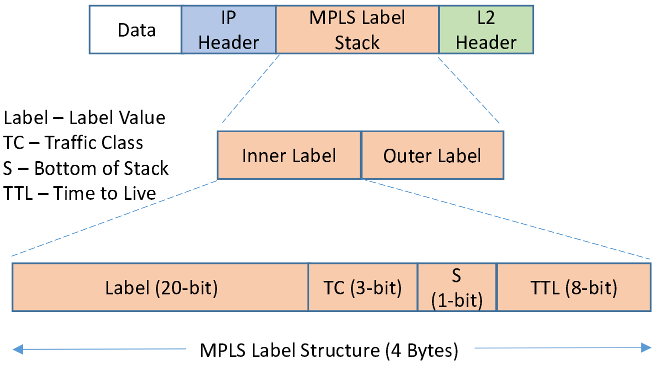

Figure 4.4.2-1: MPLS Packet Format

The MPLS labels, which have a local significance, can be used to establish a tunnel over an IP backbone network, thus supporting both L2VPN and L3VPN services. Figure 4.4.2-2 depicts the main entities involved when using MPLS VPN.

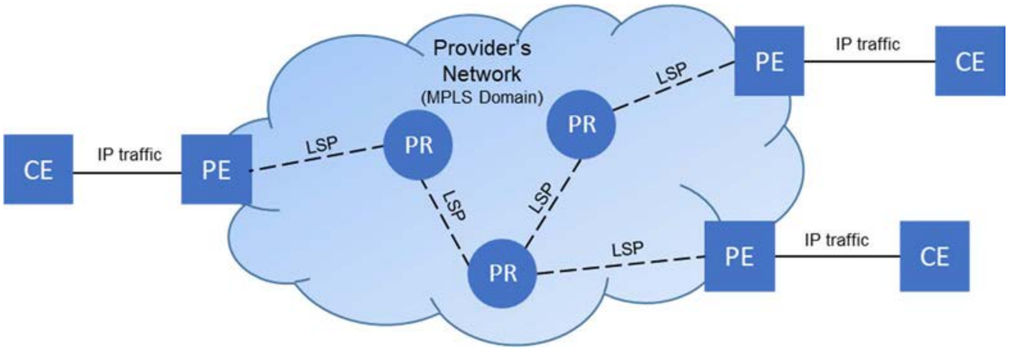

Figure 4.4.2-2: Multiprotocol Label Switching Virtual Private Network

The traffic is initially forwarded from the CE to the PE using standard IP. The PEs establish a Label Switching Path (LSP) and perform a route lookup when receiving traffic from the CE, in order to find the LSP next hop and forward traffic through the path. The ingress PE pushes an MPLS label into an IP packet and encapsulates that packet into a MPLS packet. The egress PE removes the MPLS label and forwards the original IP packet to the receiving CE. PEs involved are therefore expected to support both VPN and MPLS functionalities, while the Provider's Router (PR) inside the provider's network can only support LSP functionalities to route MPLS-labelled packets through label swapping. In the context of MPLS technology, a PE router is also called Label Edge Router (LER) and the PR is called Label Switching Router (LSR). The LSRs and the LERs use a Label Distribution Protocol (LDP) for the bi-directional exchange of label mapping information with their respective peers, also called LDP peers to build and maintain LSP databases that are used to forward traffic through MPLS networks.

In the case of a VPN Routing and Forwarding, a VRF instance is associated with each VPN. This instance includes one or more routing table, a default forwarding table, a set of policies to be applied in the VPN and routing protocols [i.23]. This VRF instance is populated with rules from other CEs and PEs in the same VPN. A locally unique number, called Route Distinguisher (RD), is also used to identify all the routes of a particular VPN.

MPLS VPN can either be used for L2VPN or L3VPN. In the first case, it is also called VPLS and routing occurs at the CEs, which are responsible for selecting a specific circuit on which traffic should be sent [i.36]. In the latter case, a MPLS L3VPN service is also called Virtual Private Routed Network (VPRN) and routing is done at the service provider's routers [i.37].

4.4.3 Segment Routing

4.4.3.1 Introduction

Segment Routing (SR) technology reassembles properties of the source routing paradigm and comprises an architecture for steering traffic through a sequence of segments along the network path. The SR architecture is specified in IETF RFC 8402 [i.56]. Each segment is identified by a Segment IDentifier (SID), while a segment list in the form of SIDs is inserted into each packet and is used by the network devices to perform certain processing and/or forwarding actions. A SID identifies a segment which can represent a node (node-SID), a local prefix (prefix-SID), a node adjacency like for example an interface (adj-SID), etc. see IETF RFC 8402 [i.56]. The segments are allocated and signalled to each node by an IGP protocol like OSPF or IS-IS (or BGP for inter-Autonomous System SIDs allocation) with the appropriate extensions. A segment list can be inserted/manipulated either in the source node or by an intermediate network node. A node individually decides to steer packets based on an SR policy (IETF RFC 9256 [i.57]). A SR policy can be built using any type of SIDs associated for example with topological or service instructions. Segment instruction examples are *forward packet through a specific interface* or *forward packet according to shortest path* etc. Since forwarding information is encoded at each packet, segment routing forwarding operations are performed over a nearly stateless data plane, with a minimum forwarding table size in the forwarding nodes. Segment routing has been realized on top of MPLS and IPv6 data planes and is able to support the necessary underlay connectivity to build L3VPN services on top.

4.4.3.2 Segment Routing over MPLS

In the case of a MPLS data plane, SIDs are encoded as MPLS labels or an index into an MPLS label space. An ordered list of segments is encoded as a stack of labels, with the top-most label on the stack being the segment that is processed. SR relies on existing MPLS forwarding operations (i.e. label Push, Swap, Pop) which are mapped to SR segment list operations (i.e. Push, continue, next). With SR there is no need for Label Distribution Protocol (LDP) to support the related MPLS operations.

4.4.3.3 Segment Routing over IPv6 (SRv6)

In SRv6, a Segment Routing Header (SRH) specified in IETF RFC 8754 [i.58] is used to insert a SR policy as an ordered list of SIDs expressed as IPv6 addresses (see Figure 4.4.3.3-1 for the structure of the SRH). The current segment to be processed by a node is indicated in the Segments List field, which is decremented by a SR-enabled router after processing the corresponding SID.

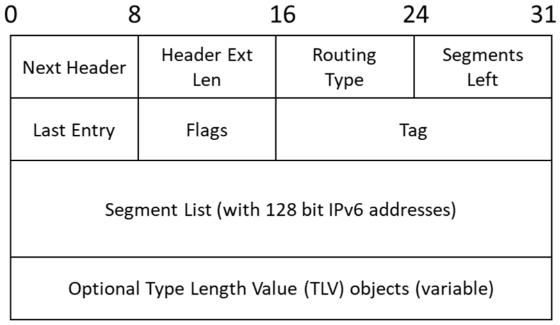
0
8
16
24
31
Next Header
Header Ext Len
Routing Type
Segments Left
Last Entry
Flags
Tag
Segment List (with 128 bit IPv6 addresses)
Optional Type Length Value (TLV) objects (variable)

Figure 4.4.3.3-1: SRH according to IETF RFC 8754 [i.58]

In SRv6, the network is divided into multiple segments, where each segment is associated with a Location and a Function to be carried out by any associated forwarding device. In more detail SRv6 SIDs in the segments list consist of the LOC:FUNCT:ARG bits, where LOC bits are used for the location/forwarding information (the address of a particular SRv6 node), followed by the FUNCT bits for the local behavior bound to the SID, followed by function arguments (ARG bits). A function could be a simple forwarding operation or even a complex operation to be performed such as network telemetry. SRv6 operations are detailed in IETF RFC 8986 [i.59].

4.5 SDN in NFV

4.5.1 Overview

A detailed report on SDN usage in NFV is provided in ETSI GS NFV-EVE 005 [i.3]. The different planes (i.e. application, management, control, data-plane) were described, like also different positioning options regarding SDN resources and SDN controllers in the NFV framework were provided. An analysis was also provided regarding different design patterns like SDN controller hierarchy and SDN control across multiple-VIMs. See also Annex A of the present document about SDN support in Openstack.

In NFV, SDN control and management mechanisms are used to perform the intended virtualized network control and management functionalities, through interaction with the VIM for the case of intra-site connectivity and WIM for the case of inter-site connectivity.

Although SDN control can be applied also for the physical network, in the context of NFV, the focus stays on the virtual network control and management aspects. Furthermore, the interface between VIM and SDN controllers is not specified in any of the referenced documentation and different implementation options can be considered.

The support of containers in the revised NFV architecture according to ETSI GS NFV 006 [i.60] imposes additional challenges on how the network segment is managed from NFV-MANO point of view. Container networking aspects have been considered in ETSI GR NFV-IFA 038 [i.61] and additional challenges like automation into container network management and network policies for container networking are investigated in ETSI GR NFV-IFA 043 [i.62].

The principles of interactions between SDN controllers and NFV-MANO are the same for VM-based solutions and for container-based environments.

Referenced documentation lacks guidelines and/or specification regarding the following aspects:

- The way an SDN controller is interacting with VIM/CISM/CCM to support intra NFVI-PoP connectivity.
- The way an SDN controller is interacting with WIM to support inter-site connectivity and MSCS management.
- The way an SDN controller is interacting with the network infrastructure. The way the network nodes and network protocol in effect (e.g. VLAN, MPLS, BGP, L2VPN) are configured and managed depends on the SDN architecture and plugins available by the SDN controller (e.g. Netconf, SNMP, PCEP).
- The way the SDN controller supports concepts like network slicing and multi-tenancy and the interfaces exposed in the northbound.

Overall NFV-MANO related operations can be realized with or without SDN control. Nevertheless, SDN control can be used to simplify network control and management from operator point of view. Recent activities in SDN support for the telco cloud are around new projects like the TeraFlowSDN by ETSI OSG [i.63] and the Tungsten Fabric by Linux Foundation [i.14]. TeraFlowSDN acts both as a SDN orchestrator and controller, while Tungsten Fabric is related to SDN overlay like also management of virtual networking in cloud environments.

5 Intra-site Connectivity Services and Enabling Technologies

5.1 Introduction

In this clause an overview of networking infrastructure and virtual networks inside an NFVI-PoP, e.g. a Data Centre (DC), is provided.

Furthermore, the present clause gives a detailed overview of enabling technologies for realizing L2 and L3 VPN services, to support intra-site connectivity. An analysis of virtualized network resource interfaces exposed by VIM considering the documented L2 and L3 VPN Service enabling technologies is also provided. Additionally, the OAM aspects of intra-site connectivity in light of available L2 and L3 VPN Service enabling technologies are analysed.

The term "site" is based on topological considerations rather than geographical [i.23]. This means that two geographically apart NFVI-PoPs can be topologically considered as a single site in case they are linked by leased lines, or they can be considered as separate sites in case they are linked over a L2/L3 VPN service. The focus of this is on the former meaning of site.

In the NFV-MANO system, VIM is the responsible entity for managing virtualized resources, including network resources within an NFVI-PoP. The interfaces exposed by the VIM over the Or-Vi reference point (as specified in ETSI GS NFV-IFA 005 [i.7]), and over the Vi-Vnfm reference point (as specified in ETSI GS NFV-IFA 006 [i.69]), enable the NFVO and VNFM respectively to provision virtualized network resources for different tenants within the same NFVI-PoP.

5.2 Networking Infrastructure and Virtual Networks

DCs employ overlay-based network virtualization approach for provisioning of Virtual Networks (VNs) in order to ensure traffic isolation and address space isolation between different tenants. IETF RFC 7364 [i.30] highlights how overlay-based network virtualization model addresses the issues related to supporting multiple tenants within a DC. IETF RFC 7365 [i.31], on the other hand, defines a reference model along with logical components that is required for implementing a Network Virtualization over Layer 3 (NVO3) in a DC.

As per [i.31], the NVO3 network forms an overlay network, which is a VN topology that operates over an IP (L3) underlay transport network and provides L2 and/or L3 service to tenant systems. The overlay network and the underlay network can use different protocols. The addressing scope of the underlay network is different from the overlay VNs, while there is traffic isolation between the VNs. Tunnelling is used in the underlay network for aggregating traffic from the VMs connected to the overlay network, thus hiding VMs' addresses from the underlay network, and reducing the amount of forwarding states in the underlay network.

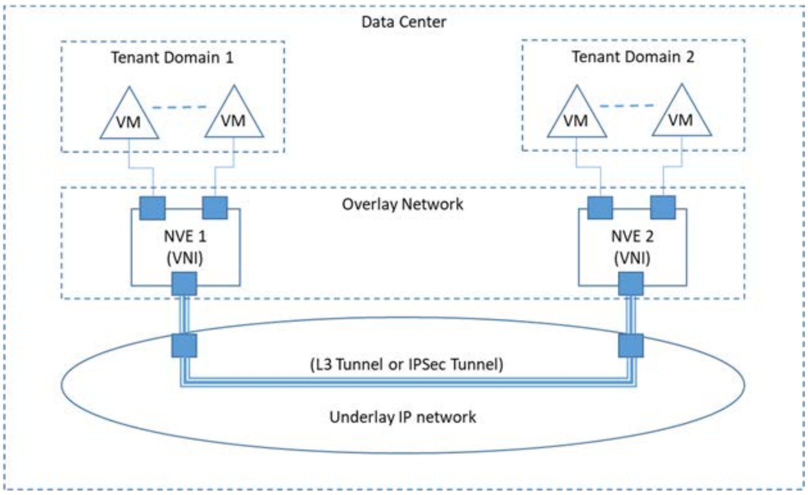

Figure 5.2-1: Overview of the Intra-DC Virtual Network Connectivity Infrastructure

As depicted in Figure 5.2-1, a Network Virtualization Edge (NVE) entity connects the overlay VN to the underlay network, while the underlay network is used to interconnect the NVE nodes. The NVE implements L2 and/or L3 network virtualization functions. The NVE sends and receives Ethernet frames from the VMs on the NBI, while it aggregates and tunnels tenant frames to and from other NVEs over the SBI connecting it to the underlay L3 network. The VMs and NVE can either be co-located or on different physical servers. For example, the VMs can be on physical servers while the NVE function is implemented in the connected ToR switch.

IETF RFC 7365 [i.31] presents the generic NVE reference model, an adaptation of which is depicted in Figure 5.2-2. The NVEs connect the VMs to the corresponding VN instance via Virtual Access Points (VAPs), which is a logical connection point on the NVE. Multiple VN instances can be instantiated on an NVE, whereas a VN instance defines a forwarding context containing reachability information and policies [i.31].

A Virtual Network Context (VNC) identifier is a field in the overlay encapsulation header for the NVE that identifies the VN instance the packet belongs to. Tenant identification and traffic de-multiplexing are based on the VNC identifier, in order to deliver the packets to the correct VN. The identifier scope can be local or global. Upon sending a packet, an NVE encodes the VNC information for the destination NVE, and the L3 tunnelling information such as the IP addresses of the source and the destination NVEs. The overlay module provides tunnelling functions such as tunnel instantiation/termination and/or encapsulation/decapsulation of frames from the VAPs/L3 underlay network.

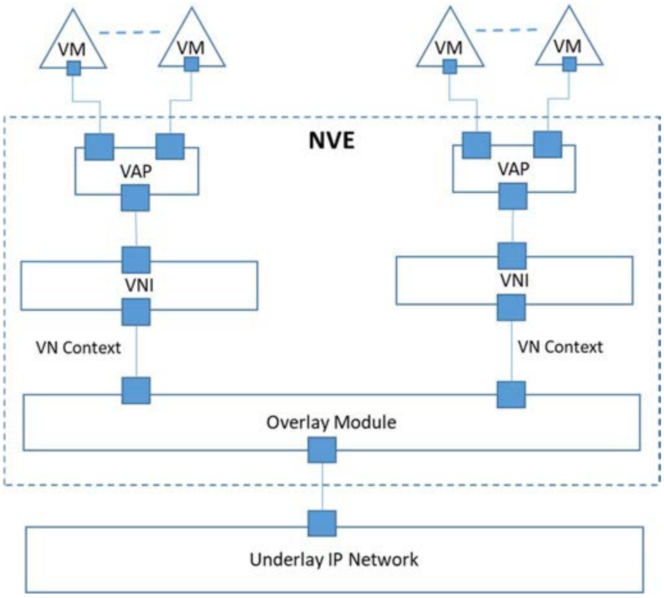

Figure 5.2-2: A Network Virtualization Entity Model

In order to share reachability and forwarding information, the NVEs can either exchange information with each other directly via a control-plane protocol, or they can obtain such information from an external Network Virtualization Authority (NVA) entity. In the latter case, a control-plane protocol is used between the NVA and the NVEs.

5.3 Enabling technologies and network resources

5.3.1 L2 VPN Service Enablers

5.3.1.1 VLAN

VLANs [i.15] allow to logically separate at Layer 2 different traffic flows, which are carried over a single physical LAN. This is done through a VLAN identifier differentiating among the different logical flows of traffic. A network device, which supports and receives a packet from a VLAN, is able to associate the frame to a specific VLAN based on the VLAN identifier carried by the Ethernet frame. Ethernet frames containing a VLAN identifier are referred to as *tagged frames*, whereas frames not having such a VLAN identifier are named as *untagged frames*.

IEEE 802.1Q [i.15] defines the frame format of a VLAN frame, as depicted in Figure 5.3.1.1-1. For tagged frames, a 4-bytes additional field named VLAN tag is included between the destination and source MAC addresses and the Length/Type field. This VLAN Tag is further composed of:

- A 2-bytes *Tag Protocol Identifier (TPID)*, which indicates the type of the frame; in particular, the value of this field is equal to 0x8100 to indicate a 802.1Q type frame that carries a single VLAN tag in the Ethernet frame. When the value is 0x88a8, it indicates an Ethernet frame with multiple VLAN tags, a technique known as *provider bridging* or *stacked VLANs*, and more commonly as QinQ [i.15].

- A 3-bit *Priority Code Point (PCP)*, having 8 priority level ranging from 0 to 7 (with 7 being the highest priority).

- A 1-bit *Drop Eligible Indicator (DEI)*, indicating the possibility to ignore the frame in case of congestion.

- A 12-bit *VLAN ID (VID)* effectively used to identify the VLAN to which the frame belongs. This value ranges from 0 to 4 095, but 0 and 4095 values are reserved; therefore, the available VLAN IDs are 4 094.

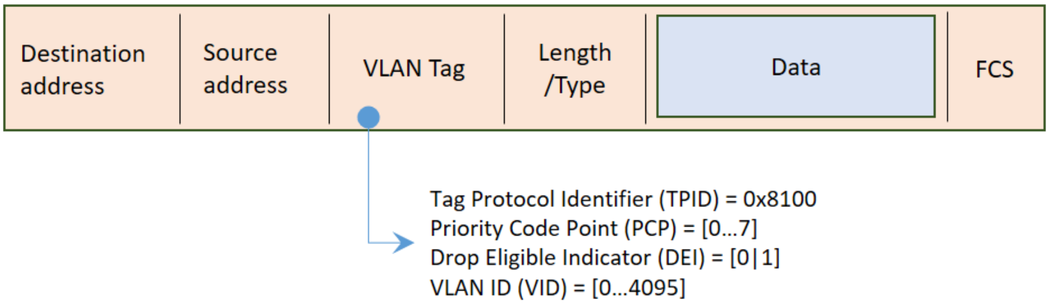
Destination address
Source address
VLAN Tag
Length /Type
Data
FCS
Tag Protocol Identifier (TPID) = 0x8100
Priority Code Point (PCP) = [0...7]
Drop Eligible Indicator (DEI) = [0|1]
VLAN ID (VID) = [0...4095]

Figure 5.3.1.1-1: VLAN frame format

With reference to Figure 5.2-1, the VLAN can be used to connect VMs belonging to multiple tenant domains by tagging the frames with the appropriate VLAN ID. The tagging of the frames can in this case be the responsibility of the NVE. According to IEEE 802.1Q [i.15], the configuration of VLAN tags can happen either dynamically or statically. In the former case, this happens through the use of the layer 2 *Multiple VLAN Registration Protocol (MVRP)* [i.15], allowing to automatically configure VLAN information on switches and similar devices. In the latter case, the VLAN ID is statically assigned through the use of management mechanisms. The tagging of the frame can also happen in a hybrid mode, with some VLAN IDs configured via management mechanisms and others through the use of MVRP.

IEEE 802.1Q [i.15] also specifies *QinQ*, allowing a single frame to have more than one VLAN Tag field. This can be used by ISPs to have their own VLANs internally while carrying traffic from clients that is already VLAN tagged. In this case, the outer tag is named *Service Tag (S-Tag)* and comes before the inner *Customer Tag (C-Tag)* that identifies the VLAN to which a user belongs. The S-Tag's TPID field is equal to 0x88a8, while the C-Tag's TPID has value equal to 0x8100.

5.3.1.2 Virtual Extensible Local Area Network

IETF RFC 7348 [i.17] describes the Virtual Extensible Local Area Network (VXLAN) framework for supporting multitenancy in a DC and enabling intra-site connectivity services. A VXLAN extends the traditional VLAN technology by supporting as many as 16 million VXLANs as opposed to 4 096 VLANs, by using a 24-bit VXLAN Network Identifier (VNI) instead of the 12-bit VLAN identifier in a traditional IEEE 802.1Q standard. This makes VXLAN suitable for instantiating more isolated VN instances in a virtualized infrastructure that includes many VMs.

VXLAN offers a tunnelling scheme in order to have an overlay virtualized L2 network over an underlying L3 network. It is used to connect in a same L2 network virtualized hosts (i.e. VMs) spread across multiple racks, where each rack may be in a different L3 network. With reference to Figure 5.2-1, the VMs in different tenant domains can be connected by using VXLAN. Several VXLAN networks/segments can be created over a common L3 underlay network, where each segment/network is identified by the VNI.

VXLANs use a device called Virtual Tunnel Endpoint (VTEP), which encapsulates and decapsulates Ethernet frames whenever it enters and leaves the VXLAN tunnel. In other words, a VTEP originates and terminates VXLAN tunnels. Figure 5.3.1.2-1 shows a VXLAN packet illustrating the various headers encapsulating the original Ethernet frame.

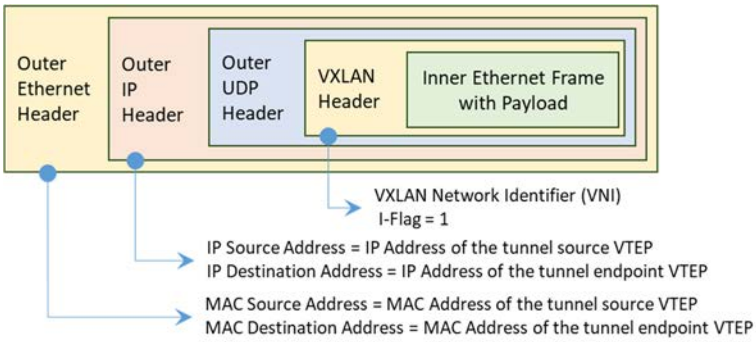
Outer Ethernet Header
Outer IP Header
Outer UDP Header
VXLAN Header
Inner Ethernet Frame with Payload
VXLAN Network Identifier (VNI)
I-Flag = 1
IP Source Address = IP Address of the tunnel source VTEP
IP Destination Address = IP Address of the tunnel endpoint VTEP
MAC Source Address = MAC Address of the tunnel source VTEP
MAC Destination Address = MAC Address of the tunnel endpoint VTEP

Figure 5.3.1.2-1: Overview of the VXLAN frame format with various header encapsulations

When a VM is sending an Ethernet frame to another VM in the same VXLAN network, the VTEP device/function encapsulates the original Ethernet frame with an 8-byte VXLAN header containing the VNI. For transport over the underlay L3 network, the VTEP further encapsulates it with a UDP header, an IP header and finally with an outer Ethernet header. In the UDP header, the source port is dynamically assigned by the VTEP while the destination port is by default 4 789. In the IP header, the source address is the IP address of the originating VTEP's IP interface, while the destination address is the IP address of destination VTEP's IP interface. Similarly, in the outer Ethernet header, the source MAC address is the MAC address of the originating VTEP's interface while the destination MAC address may be the address of the target VTEP or of an intermediate next-hop device (e.g. router).

The VTEP device/function maintains a table for mapping the inner MAC address to the VTEP IP address and the VNI. With reference to Figure 5.2-1, the VTEP function can be part of the NVE device.

5.3.1.3 Network Virtualization using GRE

IETF RFC 7637 [i.18] specifies the usage of the Generic Routing Encapsulation protocol for Network Virtualization (NVGRE) in multi-tenant DCs. NVGRE can be used to enable an overlay virtualized L2 network over an underlying L3 network with an aim to extend VLANs to solve the problem of limited number of VLANs in large DCs and stretch network segments across different domains. With reference to Figure 5.2-1, NVGRE can therefore be used to connect different VMs deployed in different tenant domains.

NVGRE endpoints are the ingress/egress points of the NVGRE tunnel, and they are mainly responsible for encapsulating and decapsulating Ethernet data frames. NVGRE endpoints can correspond to NVE elements in Figure 5.2-1, as they encapsulate Ethernet frames in order to have them routed over the underlying IP network.

Figure 5.3.1.3-1 depicts a NVGRE packet, with the additional headers encapsulating the original Ethernet frame: first a GRE header, then an Outer IP header and an Outer Ethernet header to have the packet routed in the underlay L3 network. A 24-bit Virtual Subnet ID (VSID) is used to identify different NVGRE tunnels over a common L3 underlay network. An additional 8-bit Flow ID field can be used to provide per-flow entropy for flows in the same VSID.

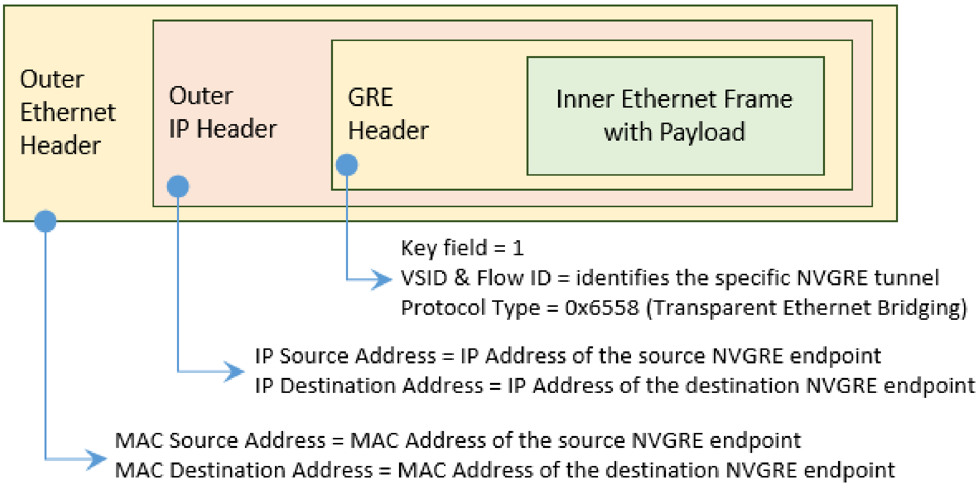
Outer Ethernet Header
Outer IP Header
GRE Header
Inner Ethernet Frame with Payload
Key field = 1
VSID & Flow ID = identifies the specific NVGRE tunnel
Protocol Type = 0x6558 (Transparent Ethernet Bridging)
IP Source Address = IP Address of the source NVGRE endpoint
IP Destination Address = IP Address of the destination NVGRE endpoint
MAC Source Address = MAC Address of the source NVGRE endpoint
MAC Destination Address = MAC Address of the destination NVGRE endpoint

Figure 5.3.1.3-1: NVGRE packet format

The MAC and IP addresses in the outer Ethernet and IP Headers correspond respectively to the MAC and IP addresses of the NVGRE endpoints. NVGRE uses an extension of the GRE header in order to realize the 24-bit VSID and the 8-bit Flow ID fields. To enable said fields, the GRE header key field bit is set to 1, as specified in IETF RFC 2890 [i.35]. The Protocol Type of the GRE Header should be set to 0x6558 (Transparent Ethernet Bridging) for Network Virtualization usage [i.18].

5.3.2 L3 VPN Service Enablers

5.3.2.1 IPsec

IPsec is a protocol designed to make IP connections secure by adding authentication, integrity, and encryption features at Layer 3 [i.41]. IPsec uses a Security Association (SA) between the two IPsec endpoints, in which shared security attributes (e.g. encryption algorithms attributes, shared keys) are established, as well as other options for the connection (e.g. enabling a NAT-traversal feature, if one of the parties involved is behind a NAT). Each SA is identified by a 32-bit Security Parameter Index (SPI). The IKEv2 [i.42] protocol is used to establish a SA.

With reference to Figure 5.2-1, an IPsec tunnel can be established between two NVEs, enabling for a secure connectivity over the underlay IP network. With reference to Figure 5.2-2, the NVE's Overlay Module can be the endpoint for said IPsec tunnel.

IPsec allows for two different configuration modes: the *transport mode* and the *tunnel mode*. The transport mode is used for end-to-end communications where the IP payload is protected using authentication protocols. The tunnel mode, on the other hand, is the default mode, and can be used to realize a VPN-like service between two locations (gateway-to-gateway). The tunnel mode therefore can be used as either an intra-site or an inter-site connectivity service enabler [i.41].

IPsec has a choice between using two protocols, namely the *Authentication Header (AH)* [i.43] and the *Encapsulation Security Payload (ESP*) [i.44], the first used for integrity and data origin authentication, and the latter for confidentiality and partial sequence integrity. In both cases, when the tunnel mode is used, the AH/ESP headers encapsulate the entire original L3 packet. Figure 5.3.2.1-1 depicts both the case of a AH and of an ESP packet format in tunnel mode. In the AH case, the AH header carries as authentication parameter an Integrity Check Value (ICV) in order to authenticate the entire packet, except fields which mutate (e.g. the TTL). In the ESP case, the internal IP packet is entirely encrypted, and the final ESP authentication value authenticates the entire inner packet, excluding the external IP header. In both the AH and the ESP cases, the new IP header which is added contains the IP addresses of the IPsec tunnel source and destination endpoints. In the transport mode, the original IP header with minor changes is used; therefore, the original IP header is not encrypted.

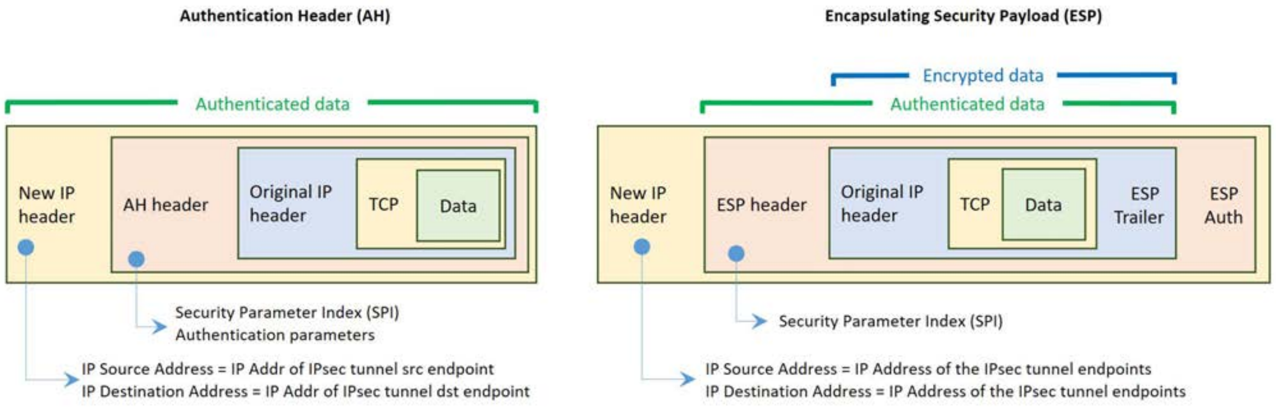
Authentication Header (AH)
Authenticated data
New IP header
AH header
Original IP header
TCP
Data
Security Parameter Index (SPI)
Authentication parameters
IP Source Address = IP Addr of IPsec tunnel src endpoint
IP Destination Address = IP Addr of IPsec tunnel dst endpoint
Encapsulating Security Payload (ESP)
Encrypted data
Authenticated data
New IP header
ESP header
Original IP header
TCP
Data
ESP Trailer
ESP Auth
Security Parameter Index (SPI)
IP Source Address = IP Address of the IPsec tunnel endpoints
IP Destination Address = IP Address of the IPsec tunnel endpoints

Figure 5.3.2.1-1: IPsec AH and ESP packet formats in tunnel mode

5.3.2.2 Interior BGP

The BGP protocol [i.11] can be used either among multiple ASs or inside one AS, as introduced in clause 4.4.1. In the latter case, it is known as *interior BGP* (iBGP) and iBGP messages are exchanged only among routers belonging to the AS, which are named *iBGP peers*. Each iBGP peer maintains routing information in a BGP Routing Information Base (RIB) [i.49], and updates such information based on the updates received by neighbour iBGP peers.

Figure 5.3.2.2-1 shows an example of three AS: AS1, AS2 and AS3. Assuming R1 in AS1 wants to advertise a route for 1.1.1.0/24 in AS1 to AS3, the advertisement needs to go through AS2. AS2 is also known as a *transit AS* [i.50]. To exchange routing information between R1 and R5, *exterior BGP* (eBGP) can be used: R1 uses eBGP to advertise the route for 1.1.1.0/24 to R2. Similarly, eBGP can be also used to advertise the same route between R4 and R5. However, in order for the route information to reach R4 from R2, these needs to be carried inside AS2 first. And to exchange information inside AS2, iBGP is used.

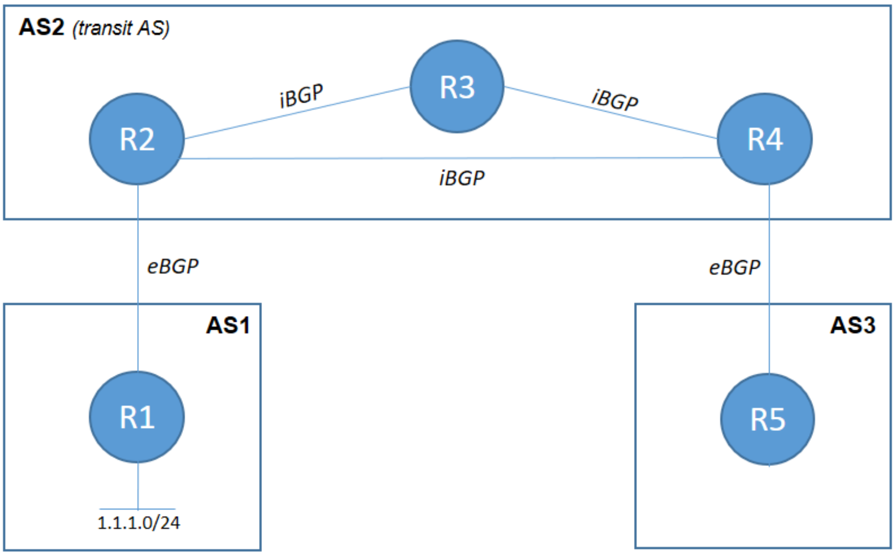

Figure 5.3.2.2-1: Example of the use of iBGP and eBGP

iBGP is not the only possibility to carry routing information inside an AS. With reference to Figure 5.3.2.2-1, the routing information in AS2 can alternatively be carried through the use of Interior Gateway Protocols (IGP), such as Open Shortest Path First (OSPF). However, IGP poses a scalability issue when dealing with a high number of routes, therefore iBGP is preferable when the dimension of the ASs and routes to be exchanged is consistently big.

In iBGP, every iBGP speaker inside one AS has a unique BGP session with all the other iBGP peers in the same AS, i.e. with all other BGP routers having the same BGP AS number. Therefore, *n* iBGP speakers belonging to the same AS are fully meshed, thus managing *n(n-1)/2* unique iBGP sessions. This poses a scaling problem in large networks, when the number of sessions established by each router with its neighbours can cause a degradation of the router performances (e.g. because of a lack of memory and/or high CPU processing requirements). The use of Route Reflectors [i.51] or BGP Confederation [i.52] has been specified to overcome such scaling issue when using iBGP.

Route Reflectors (RR) [i.51] simplify the logical topology of the network reducing the total number of BGP sessions required: a BGP RR is used as a central node, maintaining a session with all others iBGP peers, namely BGP RR Clients. A RR and all its RR Clients form a RR Cluster. Figure 5.3.2.2-2 shows an example of a RR topology with one RR Cluster, composed of a RR and three RR clients (R1, R2 and R3). As shown, the number of sessions needed is reduced from 6 (in a full-mesh topology) to 3 (each RR Client having one session with the RR). This solution therefore allows to have one or multiple RR Clusters instead of a full-mesh topology of iBGP peers, thereby reducing the complexity in bigger networks.

BGP Confederation [i.52] divides one AS into two or multiple member-AS. Each member-AS has a smaller number of iBGP peers which can be configured in a full-mesh or with a RR topology. Each member-AS has its own member-AS number that it is used internally in the AS Confederation, while the AS number, which is common among all member-AS composing one BGP Confederation, is used to exchange information with external ASs. As an example, Figure 5.3.2.2-3 depicts one AS Confederation (AS1) composed of two member-AS (member-AS 1 and member-AS 2), the former with a full-mesh topology and the latter with a RR Cluster. The iBGP peers use iBGP and the member-AS number when exchanging information within the AS Confederation, while they use eBGP and the AS1 number when exchanging messages with AS2.

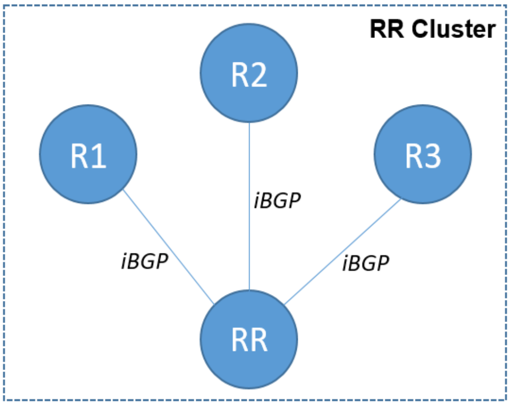

Figure 5.3.2.2-2: BGP Route Reflector Topology

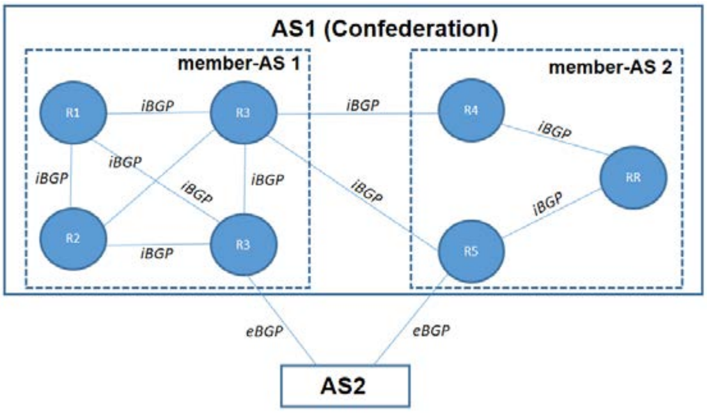

Figure 5.3.2.2-3: BGP Confederation Topology

The BGP specification [i.11] allows the advertisement of only one path to reach an address prefix. If for instance in Figure 5.3.2.2-1 multiple paths from R5 to R1 are available, only the best path is advertised among BGP routers. IETF RFC 7911 [i.53] specifies an additional identifier (i.e. Path Identifier) allowing for the advertisement of multiple paths in BGP to the same address prefix. Each path advertised does not replace the previous ones. This technique leverages the ADD-PATH Capability specified as an additional BGP capability (with Capability Code 69), in IETF RFC 5492 [i.54].

5.4 Analysis

5.4.1 Reference points and interfaces considerations

5.4.1.1 Considerations for L2VPN service enablers

L2VPN technologies described in clause 5.3.1 enable connectivity services inside a single NFVI-PoP (i.e. intra-site). Clause 5.3.1 describes some of the most common L2VPN technologies, such as VLAN, VXLAN and NVGRE. To enable intra-site connectivity services, some configuration parameters are mandatory (such as, an identifier to uniquely identify frames belonging to the same virtual L2 network), while others can be optional and used for specific purposes (such as, in the case when BGP protocol is used to learn reachability information dynamically and autonomously about remote NVEs in the same NVI [i.55]). The configuration parameters are used for configuring the virtualized network resources in the NFVI, which in turn are managed by the VIM.

ETSI GS NFV-IFA 005 [i.7] specifies Virtualised Network Management interfaces and related Information Elements (IE) on the Or-Vi reference point to "perform operations on virtualised network resources", including operations for allocating, querying, updating, and terminating virtualised network resources. For instance, the *Allocate Virtualised Network Resource* operation defined for the Virtualised Network Resources Management Interface enables the NFVO to request the VIM the allocation of virtualised network resources. The parameters and the attributes of the IEs specified for this operation can be used to configure either a VLAN, VXLAN or NVGRE connectivity services.

In particular, among the IEs specified by the ETSI GS NFV-IFA 005 [i.7], a *VirtualTrunkData* IE can be used to provide information about a virtual trunk to be created, which can contain multiple *TrunkSubport*, each with a *segmentationId* and a *segmentationType* attribute. The *segmentationId* attribute can be used to carry a VID, VNI or VSID to identify to which VLAN, VXLAN or NVGRE segment that specific *TrunkSubport* belongs to, respectively.

In the case of the NVGRE, there is an additional challenge to be addressed: apart from the segment identifier (i.e. VSID), the GRE header used by NVGRE defines an additional field, which is used for the identification of the segment in conjunction with the VSID, i.e. the Flow ID field. However, the *TrunkSubport* specified in ETSI GS NFV-IFA 005 [i.3] does not contain an additional attribute to the cited *segmentationId*, which could be used for the Flow ID to be carried.

The *segmentationType* of the *TrunkSubport* IE specified in the ETSI GS NFV-IFA 005 [i.7] supports only "VLAN" or "INHERIT" values, with the latter meaning that the type is inherited from the *VirtualNetworkData* IE to which the *TrunkSubport* is associated. Thus, ETSI GS NFV-IFA 005 [i.7] does not consider for the *segmentationType* values such as "VXLAN" and "NVGRE". To allow for the corresponding services to be selected for the segment, it is needed either (a) the "VXLAN" and "NVGRE" values for the *segmentationType* to be available, or (b) some mechanism to leverage the "INHERIT" value to ensure the use of the *NetworkType* from the *VirtualNetworkData*, when this is configured as "VXLAN" or "NVGRE".

Moreover, the *TrunkSubport* has a *portQoS* attribute, which can be used to carry QoS parameters for the specific *TrunkSubport* identified by the attribute *subportId*.

Finally, the ETSI GS NFV IFA 005 does not contain an attribute, either in the above cited *TrunkSubport* IE, or in other specified IEs, indicating a specific routing protocol to be used, e.g. iBGP in the case in which this is used to discover the NVEs belonging to a specific VXLAN [i.55].

5.4.1.2 Considerations for L3VPN service enablers

In clause 5.3.2 two L3 VPN services that can be used to enable connectivity services within an NFVI-PoP, namely IPSec and iBGP, are summarized.

IPSec can be configured in different modes and, depending on the desired feature, can contain an AH or an ESP header, as described in clause 5.3.2.1. In any of the above-mentioned configurations, IPsec contains in its header a Security Parameter Index (SPI), which is mandatory to differentiate the different flows of packets. Each IPsec tunnel may then be associated with additional authentication parameters, which are also carried to configure properly a virtualized network.

An iBGP session to be established inside a single NFVI-PoP also carries some mandatory parameters, such as the AS number which is used by all the routers for the session.

However, ETSI GS NFV-IFA 005 [i.7] does not specify any attribute in any information element through which the upon parameters for IPsec and iBGP could be carried.

5.4.2 OAM considerations

5.4.2.1 Overview

A baseline architecture for the management of the intra-site (i.e. within one NFVI-PoP) network is depicted in Figure 5.4.2.1-1. The main managed entities are the following:

- NFVI Node: it includes the physical and logical NFVI Node network components (e.g. NIC, vSwitch, SR-IOV, Linux-bridging, etc.) and the VM network components (e.g. vNIC).
- Network fabric: network elements (e.g. switches and routers, including ToR switches) of the NFVI-PoP used to connect NFVI Nodes with other NFVI Nodes or systems (e.g. DNS servers, load balancers, etc.).
- NFVI-PoP network gateway: network element which is used to connect the internal NFVI-PoP network with the WAN network (see ETSI GS NFV-IFA 032 [i.34]). This device is typically referred to as CE gateway.

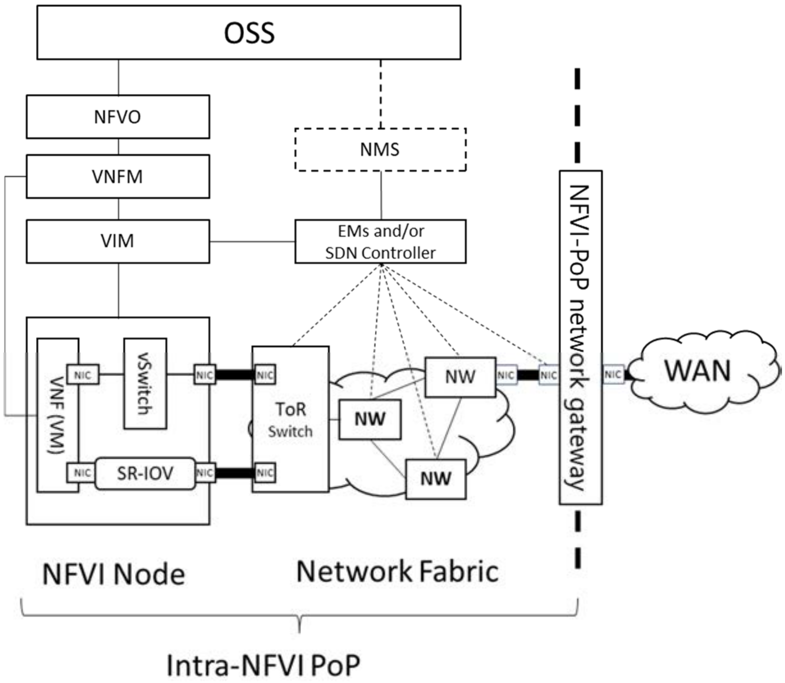

Figure 5.4.2.1-1: Intra-NFVI-PoP network

From a management perspective, a Network Management System (NMS) can be used to provide a holistic view of the entire network fabric. A NMS can provide OAM&P (OAM & provisioning) of the network through interaction with Element Managers (EMs) dedicated to managing one or more specific network devices/equipment composing the network fabric. The NMS typically has management components (like event handlers, performance correlation mechanisms, etc.), provides northbound APIs used to interact with OSS or other NMS, GUI components, and consume southbound APIs to interact with EMs. Information exchange between EMs and network devices is tailored to specific network technologies and protocols, either proprietary or also based on standardized data models (e.g. for device management, ports, interfaces, protocol management, etc.).

Regarding EM operations, an SDN controller resembles properties of EMs and can be used to manage a plurality of network devices and provide overlay network connectivity capabilities.

In a virtualized environment, OAM&P is relevant to both underlay and overlay network connectivity aspects.

5.4.2.2 Challenges

From a management point of view, ETSI NFV specifications have addressed aspects related to OAM&P at the NS, VNF and virtualised resource/containerized workload level. For instance, ETSI GS NFV-IFA 005 [i.7] specifies the interfaces exposed by VIM towards the NFVO, ETSI GS NFV-IFA 008 [i.68] specifies the interfaces exposed on the Ve-Vnfm reference point between the VNFM, VNFs and EMs, and the ETSI GR NFV-EVE 022 [i.64] studies various VNF configuration solutions based on the available interfaces, information modelling and other capabilities supported by other referenced ETSI NFV specifications.

Nevertheless, ETSI NFV specifications lack guidelines and/or relevant specification to cover OAM&P aspects in the case of the following areas of network connectivity:

- NFVI Node device level management, such as:
    - Network interface devices, including management of SR-IOV, NICs, drivers, etc.
    - Connectivity to underlay network, including management of ports, Ethernet parameters like MTU, time protocols, etc.
- Intra-NFVI-PoP network fabric management, such as:
    - Network device.

    - Management of underlay and overlay network.

- NFVI-PoP network gateway, such as:
    - Device level operations (including operations like creation of vRouters).
    - Underlay network management (e.g. port level management).

Challenges regarding OAM&P in an NFVI-PoP network environment are grouped based on the three different areas of network connectivity of concern as listed above. In the proposed solutions, the NFVI-PoP network gateway is considered as a special type of intra-NFVI network device.

5.4.2.3 Proposed Solutions

5.4.2.3.1 Solutions for Challenge #1: NFVI Node device level management

5.4.2.3.1.1 Introduction

This challenge is related to the lack of NFV-defined standard mechanisms enabling the management of NFVI Node related network devices. In the current version of the NFV-MANO architectural framework (see ETSI GS NFV 006 [i.60]), there is no identifiable management entity responsible for the management of the NFVI Node level network devices. Furthermore, no references are present in ETSI NFV specifications regarding which protocols and data models can be used for the management of such devices.

5.4.2.3.1.2 Solution SOL-1.1: BMC-based management via PIM

This solution envisions the use of a Physical Infrastructure Management (PIM) function that interacts with one or more Baseboard Management Controllers (BMC) that implement a standard interface enabling remote management capabilities of various physical subsystems in the NFVI, including network interface cards/adapters of NFVI Nodes. An example of such type of service and interface is DMTF's Redfish DSP0266 [i.67].

Relevant data schemas from DMTF's Redfish DSP2046 [i.66] are listed below:

- *EthernetInterface*: it represents a single, logical Ethernet interface or Network Interface Controller/Card (NIC). It enables the configuration of DHCP (both for IPv4 and IPv6), reading information about configured IP, addresses, MAC address, current interface speed, VLAN support on the interface, MTU size, etc.
- *NetworkAdapter*: it represents a physical network adapter connecting to a computer network, including support for adapters for Ethernet and other network technologies. It provides information about maximum number of physical functions available on the adapter, the number of physical ports on the adapter, whether the network adapter supports some form of virtualization offload, such as SR-IOV, etc.
- *NetworkDeviceFunction*: it represents a logical interface exposed by a network adapter. It enables to read and write various sets of information and configuration related to Ethernet capabilities of the network device function, including configured MAC address, MTU size, and others.
- *PCIeDevice*: it represents and describes properties of a PCIe device that is attached to a system, e.g. computer system.

DMTF's Redfish DSP2046 [i.66] also defines metrics associated to many of the managed objects, such as metrics associated to *NetworkAdapter, NetworkDeviceFunction,* and *Port*. Therefore, it can cover part of necessary monitoring functionality of performance and fault.

5.4.2.3.1.3 Solution SOL-1.2: OS-based management by VIM through standard protocols

This solution considers the VIM as means to perform relevant configuration and monitoring of NFVI Node network devices. The solution assumes that interactions to the devices are performed via the OS (or an agent in it) running on the compute system of the NFVI Node controlled by the VIM. These interactions can be realized by means of standardized management protocols, such as Simple Network Management Protocol (SNMP). SNMP provides both the capability to get information/values and set, in some cases, the values of certain variables. Depending on the hierarchy of Management Information Base (MIB) defined for the device to be managed, the Object Identifier (OID) relates to variables that can either provide configuration information or serve for monitoring purposes.

5.4.2.3.1.4 Solution SOL-1.3: Integration of device management tools by VIM

This solution considers the VIM as means to perform relevant configuration and monitoring of NFVI Node network devices. The solution assumes that interactions to the devices are performed via the OS (or an agent in it) running on the compute system of the NFVI Node controlled by the VIM. These interactions can be realized by means of one or various specialized management tools such as:

- For configuration, an agent-based management system like Progress Chef®, or agentless provisioning system such as Ansible®.
- For monitoring, tools such as Nagios® Core (supporting both agentless and agent-based monitoring) and Zabbix®.

NOTE 1: Ansible® is a trademark of Red Hat, Inc. in the United States and other countries.

NOTE 2: Progress Chef® is a registered trademark of Progress Software Corporation and/or one of its subsidiaries or affiliates in the U.S. and/or other countries.

NOTE 3: Zabbix® is a registered trademark of Zabbix LLC.

NOTE 4: Nagios® is a registered trademark owned by Nagios Enterprises.

5.4.2.3.2 Solutions for Challenge #2: Intra-NFVI-PoP network fabric management

5.4.2.3.2.1 Introduction

This challenge is related to the lack of NFV-defined standard mechanisms enabling the management of Intra-NFVI-PoP network fabric. In the current version of the NFV-MANO architectural framework (see ETSI GS NFV 006 [i.60]), there is no identifiable management entity responsible for the management of the Intra-NFVI-PoP network fabric. Furthermore, no references are present in ETSI NFV specifications regarding which protocols and data models can be used for the management of the network devices conforming the NFVI-PoP network fabric.

5.4.2.3.1.2 Solution SOL-2.1: BMC-based management via PIM

This solution envisions the use of a PIM function that interacts with one or more BMC that implement a standard interface enabling remote management capabilities of various physical subsystems in the NFVI, in this case, the network devices conforming the NFVI-PoP network fabric. An example of such type of service and interface is DMTF's Redfish DSP0266 [i.67].

Relevant data schemas from DMTF's Redfish DSP2046 [i.66] are listed below:

- *Fabric*: it represents a simple network fabric consisting of one or more switches and zero or more endpoints.
- *Port*: it provides properties to describe a port of a switch, controller or any other device that could be connected to another entity. It provides information about current speed (in Gbps) of the port, the Ethernet properties of the port, including MAC addresses information, LLDP, the link speed configuration, the link state, etc.
- *Switch*: it represents and provides properties to describe a fabric switch. It provides information such as the current internal bandwidth of the switch, an indication as of whether the switch is in a managed or unmanaged state, a reference to the collection of ports that conform the switch, the supported protocols and type of switch (e.g. Ethernet).

DMTF's Redfish DSP2046 [i.66] also defines metrics associated to many of the managed objects, such as metrics associated to *Port* and *Switch*. Therefore, it can cover part of necessary monitoring functionality of performance and fault.

5.4.2.3.2.2 Solution SOL-2-2: SDN controller exposing standard interface(s)

This solution considers one or multiple SDN controllers to be part of the NFVI. The SDN controller provides the OAM&P functionality of the NFVI-PoP network fabric and exposes standard interfaces towards other entities like a PIM, for network fabric and underlay network management, and to the VIM, for overlay network management.

An example of standard northbound API for SDN is the ONF TAPI [i.65]. Even though TAPI is more focused towards transport networks, it supports technology-specific interface profiles for Carrier Ethernet (L2).

The consuming managing function can be in these cases:

- VIM: for the management of overlay network.
- PIM: for the management of the underlay network and the network fabric.

5.4.2.3.2.3 Solution SOL-2-3: plug-in based

This solution considers the use of specific plug-ins as form of integration of the network fabric management towards its consumers, such as the VIM.

For instance, the OpenStack®, as realization of VIM's functionality according to the NFV-MANO architectural framework, offers a plug-in mechanism enabling the networking service of OpenStack to offer its capabilities by interacting with respective network devices from third-party providers. The plug-ins define the basic networking building blocks of networking services, and the third-party providers can "translate" such basic building block into a form of realization according to the offered network devices and technologies. OpenStack Neutron plug-ins are categorized in two groups: core plug-ins, which implement the core API of the networking service and consists of the elementary building blocks such as port, subnet and network; and service plug-ins, which implement the networking service API extensions for additional services like L3 router, firewall as a service, VPN as a service, etc.

NOTE: This solution is compatible also with the use of SOL-2-1 with regards to usage of SDN controller, but in this case, no standard interface is defined between the SDN controller and the consuming managing function.

5.4.2.3.2.4 Solution SOL-2-4: device-exposed standard protocol(s)

This solution considers that the NFVI-PoP network fabric devices expose standard interfaces and data models tailored to network configuration.

For instance, NETCONF is a standard protocol specified in IETF RFC 6241 [i.8] that provides the transport to communicate with YANG formatted configuration or operational data on request from an application that runs on a centralized management platform. NETCONF uses a simple Remote Procedure Call (RPC) based mechanism. NETCONF offers the capability to Create, Read, Update and Delete (CRUD) configuration data. In addition, it also offers other powerful operations for managing configurations such as: commits, copy configuration, locking configuration, etc. The client of the exposed interfaces can be an SDN controller or another control and management entity.

5.4.2.3.3 Solutions for Challenge #3: NFVI-PoP network gateway management

5.4.2.3.3.1 Introduction

This challenge is related to the lack of NFV-defined standard mechanisms enabling the management of the NFVI-PoP network gateway. In the current version of the NFV-MANO architectural framework (see ETSI GS NFV 006 [i.60]), there is no identifiable management entity responsible for the management of the NFVI-PoP network gateway. According to Annex E of ETSI GS NFV-SOL 005 [i.33], it is assumed that interactions towards the NFVI-PoP network gateway can be performed by different entities depending on the management boundaries between NFVI-PoP networking and transport networking.

Furthermore, no references are present in ETSI NFV specifications regarding which protocols and data models can be used for the management of the network devices conforming the NFVI-PoP network fabric.

5.4.2.3.3.2 Solution SOL-3-1: device-exposed standard protocol(s)

This solution considers that the NFVI-PoP network gateway exposes standard interfaces and data models tailored to network configuration.

Refer to description of relevant standard protocols in clause 5.4.2.3.2.4.

5.4.2.4 Solutions Evaluation

Table 5.4.2.4-1 provides pros/cons analysis of the solutions described for intra-site OAM management.

Table 5.4.2.4-1: Solutions evaluation for intra-site OAM management

Challenge	Solution	Pros	Cons	Comment
#1: NFVI Node device level management	SOL-1.1: BMC-based management via PIM	• Technology-ready solutions can be exploited regarding BMC operations and exposed interfaces (e.g. Redfish). • Both system/device aspects and network aspects of the Node can be managed from a single entity inside the NFVI-Pop thus simplifying the management architecture.	None	See also ETSI GR NFV-IFA 046 [i.71] clause 5.2.1 regarding potential PIM placement and relationship with NFV-MANO and ETSI GS NFV-IFA 053 [i.72] for requirements and interface specification for PIM.
	SOL-1.2: OS-based management by VIM through standard protocols	• Simplified architecture, no new entities need to be incorporated in NFV-MANO • Correlation of physical and virtual infrastructure can be performed without involving other management entities.	Different protocols to be used to manage device level (e.g. power on the device), but also network level aspects (e.g. creation of virtual interfaces). Increased complexity inside VIM to support certain physical infrastructure management aspects.	
	SOL-1.3: Integration of device management tools by VIM	• Technology-ready solutions can be exploited. • Simplified architecture, no new entities need to be incorporated in NFV-MANO. • Correlation of physical and virtual infrastructure can be performed without having to involve other management entities.	• Increased complexity inside VIM to support certain physical infrastructure management aspects. • No standard interoperability mechanisms, which can impact solution integration and its long-term maintenance.	See also ETSI GR NFV-EVE 022 [i.64] for an analysis on the use of tools like Ansible in ETSI NFV, (e.g. how VNFM can configure a VNF instance via Ansible).

Challenge	Solution	Pros	Cons	Comment
#2: Intra-NFVI-PoP network fabric management	SOL-2.1: BMC-based management via PIM	Like in SOL-1.1.	• Besides device management aspects, it is difficult to support a large set of network configuration and management capabilities (L1/L2/L3, underlay, overlay, etc.).	
	SOL-2.2: SDN controller exposing standard interface(s)	• Decoupling of control plane and data plane. • Leverage the use of open APIs in the northbound of the SDN controller to facilitate integration and maintenance. • Plug-in plus standard interfaces architecture used in the SDN controller southbound allows interaction with many different technologies. • Logically centralized, physically distributed control plane is scalable.	• Tight integration between the SDN controller and VIM leads to more complex maintenance operations (e.g. upgrade).	Can also work in conjunction with PIM by delegating physical device management to PIM.
	SOL-2.3: plug-in based	• Ease of integration with existing NFV-MANO systems (e.g. OpenStack-based), assuming readily available solutions. • Ease of integration with SDN controllers, assuming readily available solutions.	• Different project/plugin is used to manage device level aspects and different one the network aspects. • Plug-in framework is coupled to the specific NFV-MANO solution offering it, and therefore not portable to another implementation. • Not easy upgrade/maintenance, due to the coupling of the integration to specific solution.	

Challenge	Solution	Pros	Cons	Comment
	SOL-2.4: device-exposed standard protocol(s)	• No need to consider new management entities in NFV MANO. • Leverage the use of standard protocols in the northbound of the device to facilitate integration and maintenance.	• Increased complexity inside VIM to support both physical infrastructure management aspects and network management • In case the solution is not combined with the use of an SDN controller, increased complexity in VIM to perform network management without common network modelling and configuration abstractions, as configuration might change on a per-device basis.	
#3: NFVI-PoP network gateway management	SOL-3.1: device-exposed standard protocol(s)	Like in SOL-2.4.	Like in SOL-2.4.	

6 Inter-site Connectivity Services and Enabling Technologies

6.1 Introduction

Clause 4.3 of ETSI GS NFV-IFA 032 [i.34] provides a high-level overview of the multi-site connectivity network framework and resources identifying and defining the different resources involved in enabling multi-site network connectivity. In this context two managed objects are identified, namely MSCS and Multi-Site Network Connection (MSNC), that are exposed by the WIM concerning network connectivity. ETSI GS NFV-IFA 032 [i.34] specifies interfaces exposed by the WIM in order to enable the consumers (e.g. NFVO) for the LCM of multi-site network services in terms of MSCS and MSNC objects.

The present clause provides a detailed overview of the enablers of the inter-site connectivity services between different sites i.e. NFVI-PoPs, in terms of protocols and network resources. It should be noted that the definition of "site" is based on topological considerations rather than geographical [i.23]. This means that two geographically apart NFVI-PoPs can be topologically considered as a single site in case they are linked by leased lines, or they can be considered as separate sites in case they are linked over a L2/L3 VPN service. The focus of this clause is on the latter meaning of site where multiple NFVI-PoPs are linked over a VPN service. An analysis will then be carried out to identify their impacts on the OAM aspect of the NFV-MANO system in general and its reference points in particular.

6.2 Networking Infrastructure and Virtual Networks

Figure 6.2-1 represents a network infrastructure, highlighting the relevant entities and links, for enabling connectivity of composite network service components deployed across multiple sites. This figure is used as a reference to describe the scope, applicability of relevant protocols and the key design features of the relevant network resources.

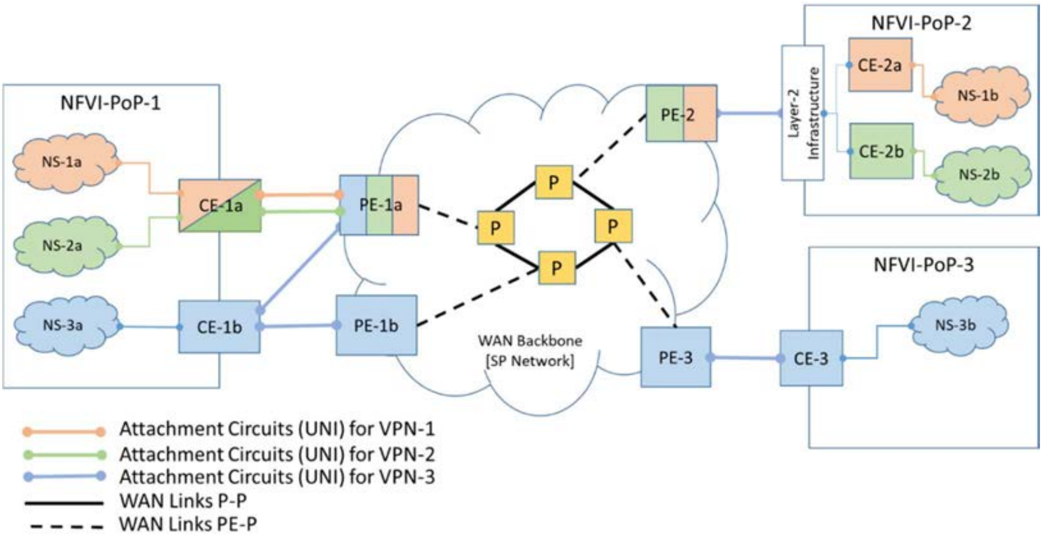

Figure 6.2-1: An inter-site network connectivity framework

Figure 6.2-1 depicts the scenario where three composite network services namely, NS-1, NS-2 and NS-3, are deployed across three sites i.e. NFVI-PoPs, interconnected over the WAN infrastructure. Each composite NS is composed of two nested NS components, where each nested NS component is deployed in two different NFVI-PoPs. For example, the nested NS components NS-1a and NS-1b of composite NS-1 are deployed in NFVI-PoP-1 and NFVI-PoP-2 respectively. Similarly, nested NS instances NS-3a and NS-3b of composite NS-3 are deployed across NFVI-PoP-1 and NFVI-PoP-3. Although in separate NFVI-PoPs, these nested NS components are inter-connected over the WAN infrastructure to realize a composite NS. The nested NS instances can belong to a tenant system and be part of a VNI as described in clause 5.

As shown in Figure 6.2-1, the NS components that are expected to communicate to other entities outside the NFVI-PoP can do so through the NFVI-PoP's network fabric connected to the CE device, which in turn is connected to a PE router. The information about the internal network fabric of an NFVI-PoP is provided in clause 5. According to IETF RFC 4364 [i.23], a CE device can be a host or a router, and typically each site, an NFVI-PoP in this case, contains one or more CE devices. A CE device that connects to the PE router is referred to as a CE router. In case a non-routing CE device connects to a PE device (not shown in Figure 6.2-1), then it is called a CE host. The CE router is connected to *one or more* PE router via an *attachment circuit*, which can be a data-link or it can be a tunnel of some sort enabling the routers to connect to each other in order to establish L3 peering.

Typically, the CE router communicates with the PE router through the Layer 2 switch. If the Layer 2 infrastructure provides a multi-point service then multiple CE routers can communicate with the PE router over a single attachment circuit, as shown in the case of NFVI-PoP-2 in Figure 6.2-1. For robustness, a CE router may attach to multiple PE routers as in the case of NFVI-PoP-1 where CE-1b is connected to PE-1a and PE-1b via two attachment circuits shown in Figure 6.2-1.

The interconnectivity between NFVI-PoPs is realized via the associated PE routers, which are inter-connected via the WAN network fabric characterized by the Provider (P) routers. The PE and P routers typically belong to a Service Provider's (SP) WAN infrastructure.

Constituents of NS instance can connect to the CE router, and the connectivity aspects are managed based on intra-site connectivity protocols as described in clause 5. According to ETSI NFV GS-IFA 032 [i.34], the CE router's ports/interfaces over which the attachment circuit is enabled is referred to as *connectivity service end-point* that characterizes the endpoint of the connectivity fulfilled by the MSCS and represents the UNI between the NFVI-PoP and the external network that interconnects the multiple sites. The connectivity service endpoint represents the shared context information of the attachment circuit connecting the NFVI-PoP and the external network. ETSI GS NFV-IFA 005 [i.7] defines the *ConnectivityServiceEndpoint* information element in clause 8.10.4, which is a VIM's managed object containing UNI (or attachment circuit) data between NFVI-PoP and the external network (e.g. the WAN). Moreover, ETSI GS NFV-SOL 005 [i.33] defines the attribute *connectivityServiceEndpointConfigDatas* that enables the API consumer to provide information relevant for the configuration of the connectivity service endpoints in order to establish the MSCS endpoint from the NFVI-PoP network gateway perspective.

An L2/L3 VPN service enables the interconnectivity among nested NS components in different sites. With reference to Figure 6.2-1, there are three VPN services to instantiate composite NSs, namely NS-1, NS-2 and NS-3, respectively.

These VPN services are referred to as VPN-1, VPN-2 and VPN-3 respectively, and colour coded in Figure 6.2-1 to differentiate between them. The CE and PE routers that are used for the connectivity of nested NS components are part of the respective VPN service. As illustrated in Figure 6.2-1, the CE and PE routers can be part of multiple VPN services. For example, CE-1a supports VPN-1 and VPN-2, while PE-1a is configured to support all three VPN services. A single attachment circuit may support multiple VPN services, as is the case with PE-2, or separate attachment circuits may be configured between the CE and PE routers, as is the case between CE-1a and PE-1a routers.

The CE routers of different NFVI-PoPs participating in a common VPN are not connected directly but via the PE routers. The PE routers know about the VPN service(s) of its corresponding CE routers, and thus exchange routing information with other related PE routers using a routing protocol, such as BGP. The PE router(s) then provide the learned routes to its corresponding CEs. The PE routers maintain routing information about supported VPN service(s). For example, PE-1a contains route information about VPN-1, VPN-2 and VPN-3, while PE-1b and PE-3 contain route information about VPN-3 only. This routing information in a PE router is maintained inside a VPN Routing and Forwarding (VRF) table. According to IETF RFC 4364 [i.23], a PE-CE attachment circuit can be associated with exactly one VRF table. With reference to Figure 6.2-1, PE-1a maintains three VRF tables for each of the three VPN service it is associated with, while the other PE routers have one VRF table. IETF RFC 4364 [i.23] also dictates that a single attachment circuit can be associated with multiple VRF tables in the case when a single VPN service is divided into several sub-VPNs.

When an IP packet is received over a particular attachment circuit, its destination IP address is looked up in the associated VRF table. The result of that lookup determines how to route the packet. If an IP packet arrives over an attachment circuit that is not associated with any VRF table, the packet's destination address is looked up in the *default forwarding table*, and the packet is routed accordingly. Thus, a PE router has one default routing table containing public routes, and as many VRF tables as the number of attachment circuits it supports, containing private routes.

With reference to Figure 6.2-1, for the construction of the underlay network, the segment routing technology can be considered to support intra-site communication (within the NFVI-PoP) but also inter-site communication between NFVI-PoPs. In this regard, the relevant network devices (CE, PE, P) are expected to be able to support segment routing functionality. From a management perspective, PE and P elements are managed by WIM as defined by ETSI GS NFV-IFA 032 [i.34].

Note that the relevant interface is not about providing details on the relevant configuration that enables Segment routing (e.g. BGP SID distribution) rather than the MSCS related protocol aspects (e.g. *mscsLayerProtocol* attribute of the MscsData information element is used to define that the type of MSCS is L3VPN). It is up to the network control mechanism in effect (e.g. SDN controller) to actually realize the MSCS and consequently configure the necessary SR aspects. Similarly, within the NFVI-PoP, the relevant connectivity aspects are handled by the VIM. Regarding the NFVI-PoP gateway configuration for the end-to-end connectivity establishment the different management models and network demarcation described in ETSI GS NFV-SOL 005 [i.33] are applicable (i.e. the CE is NFVI-PoP network provider managed or WAN provider managed or co-managed).

In the above description segment routing is used as an example of an underlay technology, similar reasoning can be considered for other underlay technologies (e.g. MPLS based underlay, etc.).

6.3 Enabling technologies and network resources

6.3.1 L2 VPN Service Enablers

6.3.1.1 Virtual Private LAN Service (VPLS)

VPLS is a L2VPN technology that provides Ethernet based point to multi-point (P2MP) connectivity service over IP or MPLS networks using L2 label switching technology. Typically, it is used to interconnect NFVI-PoPs via PE devices over Ethernet pseudowires (PWs) in a common Ethernet broadcast domain as if the participating PE devices were connected by a LAN. An Ethernet PW is used to carry Ethernet/802.3 Protocol Data Units (PDUs) over an MPLS network, and operates either in raw mode or tagged mode [i.40]. In a VPLS, the Local Area Network (LAN) at each site is extended to the edge of the provider network. The provider network then connects all the customer LANs by emulating a switch or bridge to create a single bridged LAN, thereby enabling full mesh connectivity. VPLS is also known as Transparent LAN Service (TLS) and Virtual Private Switched Network service.

The LAN segments, including VLAN and VXLAN, in different NFVI-PoPs are interconnected over a common VPLS domain. A VPLS domain is characterized by the PE devices and u-PE devices to which the CE devices connect via Attachment Circuits (ACs).

NOTE: The AC can be an Ethernet interface, an Ethernet trunk or a VLANIF. The VLANIF is a VLAN-based L3 logical network configured with an IP address to enable inter-VLAN communication.

According to IETF RFC 4761 [i.38], a u-PE is essentially a L2 PE device used for L2 aggregation. In the context of VPLS, the PE device is also referred to as a VPLS edge (VE) device [i.38]. CE devices are VPLS-unaware.

All the PEs that are participating in a VPLS domain are assumed to be fully meshed in the data plane. This means that there is a bidirectional PW between every pair of PE that is part of the same VPLS domain, enabling PEs to send VPLS packets to each other directly without the need of any intermediate PE device.

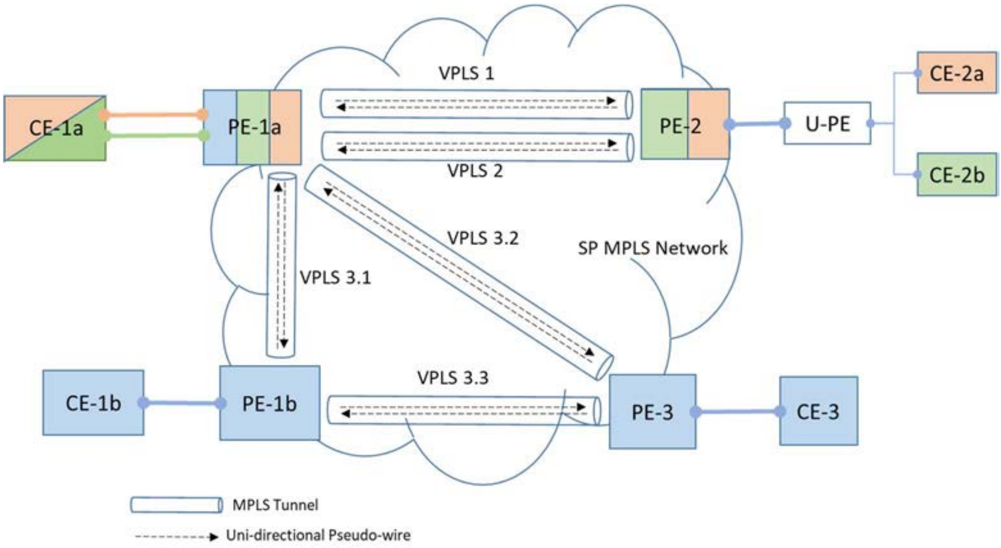

Figure 6.3.1.1-1: An example scenario with three VPLS domains

Figure 6.3.1.1-1 shows inter-site connectivity enabled by the VPLS L2VPN technology with reference to the scenario depicted in Figure 6.2-1. As shown in Figure 6.3.1.1-1, three VPLS domains (VPLS 1, VPLS 2, VPLS 3) are realized in order to interconnect the three composite network service instances deployed in different NFVI-PoPs via PE routers over an MPLS IP backbone network. VPLS 1 and VPLS 2 are independent and isolated point-to-point VPLS connectivity services realized between PE-1a and PE-2, whereas the nested components of the composite NS-1 and NS-2 connect to these PEs via their respective CEs (see Figure 6.2-1). On the other hand, since the CEs linked to the nested NS components of composite NS-3 connect to PE-1a, PE-1b and PE3, these three PEs are linked to each other in a fully meshed VPLS domain represented by VPLS 3. In case the SP network is an MPLS network, a VPLS domain is realized by an MPLS tunnel between any two PEs that is part of the same VPLS domain over which data can be transmitted transparently between the PEs. The MPLS tunnel consists of bidirectional PWs, which is actually a pair of unidirectional MPLS virtual circuits in opposite directions. In case of VPLS 3 domain, three MPLS tunnel instances, namely VPLS 3.1, VPLS 3.2 and VPLS 3.3, are realized between PE-1a, PE-1b and PE-3 to realize a full-mesh VPLS-based connectivity service.

NOTE: The tunnels between the PE devices in a VPLS domain can also be IP tunnels, such as GRE.

A VPLS network consists of a control plane and the forwarding plane.

The VPLS control plane has two primary functions;

- Auto-discovery of PEs in a VPLS domain by means of BGP, where the PEs in a VPLS domain tell other PEs in the same VPLS domain that it is also a member of that domain.
- Signaling mechanism for the establishment, maintenance and tearing down of PWs between each auto-discovered PE in a VPLS domain.

IETF has specified two different methods of enabling VPLS with both methods offering the same function. IETF RFC 4761 [i.38] specifies VPLS using BGP for auto-discovery and signalling, while IETF RFC 4762 [i.39] specifies VPLS using Label Distribution Protocol (LDP) signalling.

VPLS incorporates MAC address learning, flooding, and forwarding functions in the context of PWs that connect these individual LANs across the packet switched network.

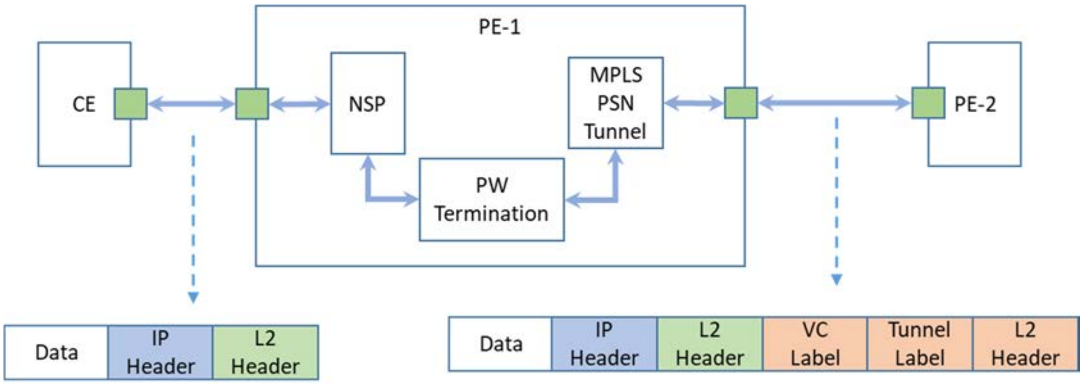

Figure 6.3.1.1-2: Overview of VPLS Packet Encapsulation over PW

The VPLS forwarding plane has three main functions as described below:

1) Encapsulation, whereby an ingress PE encapsulates the Ethernet frames it receives from a CE over an AC with a forwarding header relevant to the Packet Switched Network (PSN), which in our case is the MPLS network.

2) Forwarding, whereby the ingress PE device will forward the encapsulated VPLS packet over the MPLS tunnel. Packets are forwarded in the context of the service instance based on the destination MAC address.

3) Decapsulation, whereby the egress PE device will decapsulate the received packet and forward the inner Ethernet frame over the AC to the CE.

The encapsulation/decapsulation is carried out as described in IETF RFC 4448 [i.40]. Figure 6.3.1.1-2 illustrates the encapsulation and forwarding process of Ethernet frames over the VPLS network. An Ethernet frame is received by the ingress PE device (PE-1) from the CE over the AC. The Native Service Processing (NSP) function in the PE processes the Ethernet frame before forwarding it to the PW Termination Point.

The NSP functions can include stripping, overwriting or adding VLAN tags, physical port multiplexing and demultiplexing, PW-PW bridging, L2 encapsulation, shaping, policing, etc. [i.38]. The PW Termination Point on the other hand is responsible for operations for setting up and maintaining the PW, and for encapsulating the Ethernet frames as necessary to transmit them across the MPLS network and decapsulate the frames received from the MPLS network [i.38]. The MPLS PSN Tunnel is used for managing the tunnel containing the PWs terminating between two PEs. As shown in Figure 6.3.1.1-2 the PE-1, with PW Termination Point and MPLS PSN Tunnel function encapsulates the Ethernet frame with MPLS labels, whereas the inner label identifies the MPLS VC while the outer header is the MPLS tunnel label. This outer header is decapsulated by the egress PE device (i.e. PE-2), which then forwards the original Ethernet frame to the corresponding CE over the AC.

6.3.1.2 Ethernet Virtual Private Network

Similar to VPLS technology, Ethernet VPN (EVPN) is a L2VPN technology for interconnecting L2 and L3 networks spanning large datacentres over an IP or IP/MPLS backbone network. Moreover, it can carry L2VPN services thus reducing protocol complexity. EVPN is a BGP MPLS-based solution that overcomes several limitations of the VPLS technology in terms of meeting several requirements for datacentre deployments described on IETF RFC 7209 [i.45]. EVPN uses several building blocks from existing MPLS technology while it requires extensions to IP/MPLS protocols as specified in IETF RFC 7432 [i.46].

In EVPN, unlike VPLS where MAC learning takes place via flooding, the MAC learning between PE routers occurs in the *control plane* using Multiprotocol BGP (MP-BGP) [i.47]. The EVPN instances (EVI) are configured on PE routers, which exchange reachability information with each other using MP-BGP [i.47] extensions to transmit L2 and L3 reachability information, thus separating the control plane from the forwarding plane. This feature of EVPN can support VNF migration from one NFVI-PoP to another, which is also known as MAC mobility. When a VNF is moved, it transmits a gratuitous ARP, thereby updating the MAC-VRF table at the destination PE router. This PE device then sends a MAC route update message to all peering PE routers, which in turn update their respective MAC-VRF tables. Since the MAC address learning is handled at the control plane, EVPN can thus support different types of data plane encapsulation techniques between PE routers e.g. MPLS.

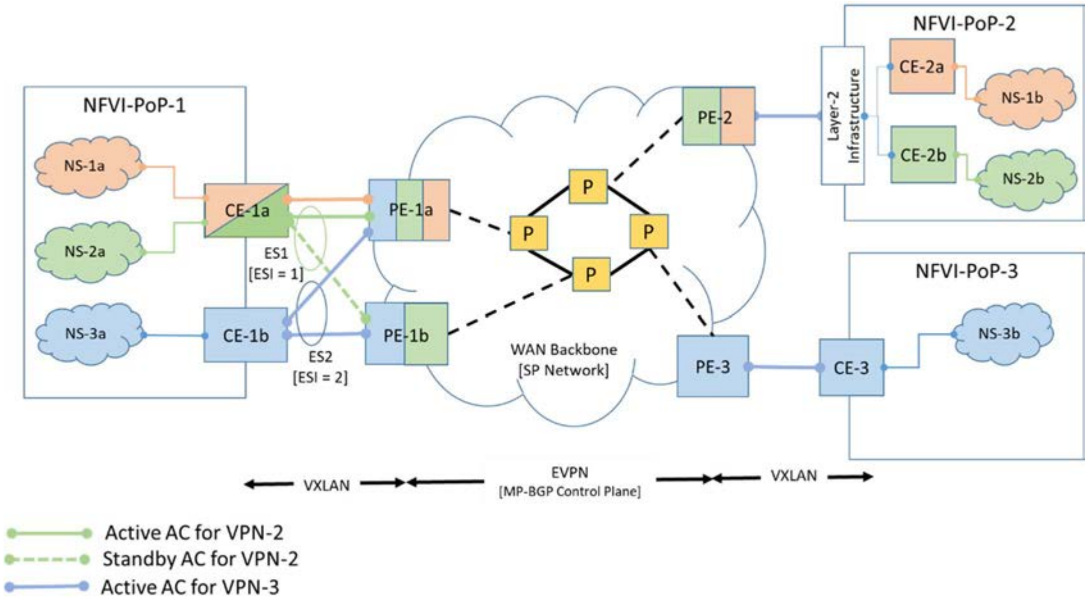

Figure 6.3.1.2-1: Example scenario of EVPN network interconnecting VXLAN segments

Besides MAC mobility, EVPN also supports multihoming thereby enabling a CE device to load-balance between multiple PE routers, and provides redundancy in case of the failure of an Attachment Circuit (AC) or a PE router. Moreover, the EVPN multihoming feature enables a remote PE router to load-balance traffic towards multihomed PE routers. This is depicted in Figure 6.3.1.2-1, which is based on Figure 6.2-1 showing an example scenario where VXLAN segments in different NFVI-PoPs are interconnected using EVPN technology. As illustrated, the CE-1a device multihomes the AC for VPN-2 to PE-1a and PE-1b routers in an *active-standby* configuration enabling redundancy. The CE-1b device multihomes the AC for VPN-3 to PE-1a and PE-1b routers in an *active-active* configuration enabling load-balancing. The Ethernet links that multihome a CE device to two or more PE routers are collectively referred to as an Ethernet Segment (ES), and each ES is identified by an Ethernet Segment Identifier (ESI). Figure 6.3.1.2-1 the two ES instances (i.e. ES1 and ES2) are uniquely identified as ESI = 1 and ESI = 2 respectively.

An ESI = 0 refers to a single-homed CE. The ESs are used to extend/connect the VXLAN segments to the EVIs, configured on the PE routers. Each VPN customer (i.e. VXLAN) is identified by an EVI and a PE router can have multiple EVIs. Each EVI has an independent MAC-VRF table, which stores MAC addresses learned by an EVI through MP-BGP. The PE routers perform data-plane learning on the traffic received over the VXLAN tunnels.

6.3.2 L3 VPN Service Enablers

6.3.2.1 MPLS-based L3 VPN Service

As described in clause 4.4.2, MPLS can be used to establish L3VPN service. Such a network is also referred to as a VPRN and routing is done at the service provider's routers. An IP backbone is required to enable a MPLS-based L3 VPN service interconnecting VPN sites (i.e. NFVI-PoPs) in different ASs. In case MPLS is not being used as the tunnelling technology, filtering is done to ensure that an MPLS-in-IP or MPLS-in-GRE packet can be accepted into the backbone [i.23].

IETF RFC 4364 [i.23] describes details of setting up the MPLS L3VPN between VPN sites, and it describes the key role of EBGP for setting up routes in VRF tables by means of route distribution, which is a key step for establishing MPLS L3VPN service. The present clause summarizes how the packets are forwarded on an MPLS L3VPN connection between VPN sites with reference to Figure 6.2-1. It is assumed that the PE routers and the WAN backbone support MPLS.

When a PE router receives an IP packet belonging to a VPN from a CE device over an ingress AC, it converts the IP address into a VPN-IPv4 address. The VPN-Ipv4 address is a 12-byte quantity, beginning with an 8-byte RD and ending with a 4-byte Ipv4 address [i.23]. An RD is simply a number without any inherent information, but it is used to distinguish VPNs that use the same Ipv4 address prefix. In this way distinct paths and policy can be computed/installed by EBGP for different VPNs having same address prefix. IETF RFC 4364 [i.23] describes the encoding of RD. Using BGP, the PE routers distribute the VPN-IPv4 routes between the CE routers that belong to the same VPN but associated to different VPN-sites (or NFVI-PoPs). Thus, the CE routers do not peer with each other and the P routers inside the WAN are not aware of the VPNs or how packets are to be forwarded from one VPN-site to another.

Thus, upon the reception of an IP over an AC from CE, the PE will select a particular VRF, whereas the choice of VRF selection is based on the ingress AC the packet is received from. Based on the destination address the next hop router is determined. In case the next hop is reachable directly over a VRF AC from the same PE i.e. the ingress AC and the egress AC are on the same PE, then the packet is forwarded on the egress AC and no MPLS labels are pushed on the packet's label stack.

If the packet's next hop is not reachable through any AC associated with a VRF, then the packet travels at least one hop through the SP's IP-based WAN backbone (see Figure 6.2-1). With reference to Figure 6.2-1, it is assumed that the packet from NS-1a belonging to VPN-1 is received by PE-1a via CE-1a over the AC for VPN-1. The PE-1a selects the VRF associated with this AC and determines the BGP Next Hop router for the destination address. In this example, the BGP Next Hop is PE-2, which is associated with VPN-site NFVI-PoP-2 hosting NS-1b also belonging to VPN-1. The BGP Next Hop (i.e. PE-2 in this example) is *assigned an MPLS label, which is referred to as "VPN Route Label",* for the route that best matches the packets destination address. Thus, the IP packet received from the CE over the AC is converted into a MPLS packet with the VPN Route Label as the only label on the packet's label stack.

The PE routers insert the /32 address prefixes into the IGP routing tables of the P-routers in the backbone; thereby enabling MPLS at each P-router in the WAN to assign a label corresponding to the route to each PE-router that is part of the sites belonging to the VPN to which the packet belongs. For this purpose, LDP is expected to be supported for setting up LSPs across the backbone. The packet is then tunnelled to the BGP Next Hop (i.e. PE-2 in our example) *by adding another label to the MPLS label stack referred to as the "Tunnel Label".* Figure 6.3.2.1-1 depicts the MPLS packet format with the relevant labels.

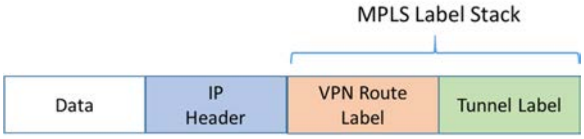

Figure 6.3.2.1-1: MPLS Packet Format with MPLS Labels

The Tunnel Label is used for tunnelling the packets along the LSPs in the IP backbone until it reaches the BGP Next Hop, which in our example is PE-2 router. At the BGP Next Hop router, the Tunnel Label is removed and the VPN Route Label is examined and processed. Based on the VPN Route Label, the PE device determines the egress over which the packet is transmitted towards the CE device. In this example, the PE-2 router forwards the packet towards CE-2a device over the AC for VPN-1. The fact that the Tunnel Label is used to tunnel the packets with the VPN Route Label through the IP backbone makes it possible to keep the VPN routes out of the P-routers, and thus ensures the scalability of the scheme.

In case the SP WAN does not support MPLS, then the MPLS packet carries only the VPN Route Label and the Tunnel Label is replaced with the IP Header or the GRE Header. This results in an MPLS-in-IP or MPLS-in-GRE encapsulation and is described in IETF RFC 4023 [i.19]. In case of MPLS-in-IP, the source address and destination address of the outer IP header are set to addresses of the encapsulating and decapsulating LSRs (i.e. PE1 and PE-2 in this example scenario), and the Protocol Number field is set to 137 indicating an MPLS unicast packet. In the case of MPLS-in-GRE, the packet then consists of an IP header (either IPv4 or IPv6), followed by a GRE header, and followed by the MPLS packet (with the VPN Route Label). This encapsulation causes MPLS packet to be sent through "GRE tunnels". In the present example, the GRE tunnel end-points can be on the ingress and egress PE routers of the VPN-1 (i.e. PE-1 and PE-2). The egress PE router will decapsulate the MPLS packet by removing the IP and the GRE headers, exposing the MPLS packet with the VPN Route Label. It then forwards this packet to the relevant CE device as explained above.

6.3.2.2 Multiprotocol BGP (MP-BGP) for L3 VPN Service

As mentioned in clause 4.4.1, MP-BGP [i.47] is an extension of BGP to enable the carrying of routing information for multiple network layer protocols beyond just IPv4. MP-BGP is thus widely used in the case of MPLS L3 VPNs (see clause 6.3.2.1) where the CE router sends the IPv4 addresses, which is translated into *labelled* VPN-IPv4 addresses by the PE router using a configured RD. The MP-BGP is then used by PE routers to distribute the MPLS labels and the VPN-IPv4 routes between the CE routers that belong to the same VPN but associated to different VPN-sites (see clause 6.3.2.1 for more details).

The MP-BGP encodes the VPN-IPv4 address in the Network Layer Reachability Information (NLRI) in the BGP UPDATE message. For this purpose, if the Address Family Identifier (AFI) is set to 1, the Subsequent Address Family Identifier is set to 128 and the capability code is set to 1, then the NLRI is an MPLS-labelled VPN-IPv4 address. The PE distributes the exact set of routes that appears in the VRF table.

A PE learns about a particular VPN's IP routes from a peering CE router that is in that particular VPN. The routes thus learned over the AC are then installed in the VRF associated with that AC. The PE and CE routers may be RIP, OSPF or BGP peers in order for the PE router to learn the set of IP routes of VPN(s) that the CE is part of.

A PE router *exports* the VPN-IPv4 to BGP, which distributes it to all other PEs that need to know about it. The receiving PEs convert the VPN-IPv4 routes into IPv4 routes and *import* them into one or more VRF tables. This is achieved by associating a VPN-IPv4 route with one or more Route Target (RT) attributes, which are then carried in BGP as attributes of the route. A VPN-IPv4 route has one RD but it can have multiple RTs. An RT identifies a set of sites, or more precisely a set of VRF tables, and the VPN-IPv4 route associated with a RT will be placed in the VRF tables that are used for routing traffic that is received from the corresponding sites.

In the case the CE and PE are BGP peers, then the SP may allow the customer to specify how its routes are to be distributed by agreeing in advance with the set of RTs that the customer is allowed to attach to its VPN routes. The CE then attaches the one or more of those RTs to each IP route that it distributes to the PE.

With reference to Figure 6.2-1, a set of RTs that the PE-1a router attaches to a route received from the site, say NFVI-PoP-1, are called *Export Targets,* whereas the set of RTs that PE-1a router uses to determine whether a route received from another PE router(s) (e.g. PE-2 or PE-3) could be placed in the VRF table associated with NFVI-PoP-1 are called *Import Targets.*

Besides RTs, a PE router may also assign a *Site of Origin* attribute to the route, which is encoded as a Route Origin Extended Community attribute of the BGP. This attribute uniquely identifies the set of routes learned from a particular site, e.g. NFVI-PoP-1 in Figure 6.2-1, to ensure that a route learned from a particular site via a particular PE/CE connection (e.g. CE-1a and PE-1a in Figure 6.2-1) is not distributed back to the same site through a different PE/CE connection (e.g. CE-1b and PE-1b in Figure 6.2-1).

6.4 Analysis

6.4.1 Reference points and interfaces considerations

6.4.1.1 Considerations for L2VPN service enablers

The inter-site connectivity between NFVI PoPs spread across multiple sites can be achieved using various L2/L3 VPN service enabling technologies. An L2 or L3 VPN service for inter-site connectivity essentially extends the VPN from one site (NFVI PoP) to other sites via a WAN backbone. Clause 6.3.1 provides an overview of two L2VPN technologies, VPLS and EVPN, that enable inter-site connectivity. The WIM manages the inter-site connectivity through the interfaces and managed objects that it exposes towards the NFV-MANO. WIM provides the means to abstract the underlying network connectivity and offers a way to establish MSCS between multiple MSCS endpoints belonging to different NFVI-PoPs.

ETSI GS NFV-IFA 032 [i.34] defines interfaces and information elements that are exposed by WIM. Through its MSCS Management Interface, WIM offers the consumers to manage the LCM of inter-site connections using MSCS and MSNC information elements. These information elements allow the API consumer to configure the multi-site network connections (i.e. MSNCs) in between different NFVI PoPs and the network connectivity details of the MSCS, that is realized by one or more MSNCs. The analysis in this clause identifies the key components and parameters of L2VPN enabling technologies, particularly VPLS and EVPN, for establishing inter-site connectivity between multiple NFVI-PoPs.

The MSCS Management Interface specified in ETSI GS NFV-IFA 032 [i.34] enables an authorized consumer to perform operations related to MSCS, such as, Create MSCS operation, Query MSCS operation, Update MSCS operation, Terminate MSCS operation, etc. to manage the lifecycle of MSCS with the support of relevant information elements. The *mscsLayerProtocol* attribute in the *MscsData* information element can be configured to support the applicable L2VPN technology, e.g. VPLS, EVPN etc. Other attributes in the *MscsData* information element such as *mscsName* and *mscsDescription* can be used to carry the name and description of the corresponding L2VPN service respectively.

ETSI GS NFV-IFA 032 [i.34] specifies additional information elements such as *MscsProfile* and *MscsEndpointData*, which can be used to carry further VPN configuration parameters, particularly on the mode of connectivity, directionality, link aggregation, network addressing, QoS, bandwidth etc. The *connectivityMode* attribute in *MscsProfile* represents whether the connection is point-to-point or multipoint, which can be configured appropriately for VPLS and EVPN. Additionally, QoS and bandwidth requirements can be configured using the *qosMetric*, *bandwidthIn* and *bandwidthOut* attributes of the *MscsProfile* information element. As reported in ETSI GR NFV-SOL 017 [i.22], the *bandwidthIn* attribute corresponds to the inbound bandwidth for the link between the PE and the CE, i.e. from PE to CE. Similarly, *bandwidthOut* attribute represents the bandwidth of the link in the outbound direction, i.e. from CE to the PE. The *directionality* attributes in *MscsEndpointData* and *MscsProfile* information elements can be used to configure the topology and direction of traffic flow for the MSCS endpoint. The *lag* attribute in *MscsEndpointData* provides a way to configure link aggregation from the NFVI-PoP toward a PE in WAN, as described in the VPLS example scenario shown in Figure 6.3.1.1-1.

NOTE: ETSI GR NFV-SOL 017 [i.22] profiles the ETSI GS NFV-IFA 032 [i.34] information elements with the IETF Layer 2 VPN YANG Model [i.25]. At the time of the work performed under ETSI GR NFV-SOL 017 [i.22], the IETF Layer 2 VPN YANG model was under 'draft' status, and it contained attributes *svc-input-bandwidth* and *svc-output-bandwidth* which are now changed to *svc-pe-to-ce-bandwidth* and *svc-ce-to-pe-bandwidth* respectively in the published IETF RFC 9291 [i.25].

Consequently, WIM uses *Mscs* information element to provide information back to the consumer about the established connectivity upon successful MSCS creation. It contains the appropriate identifier for the created MSCS, and information related to the MSCS profile, MSCS endpoints and the MSNCs realizing the MSCS. The *msnc* attribute in the *Mscs* information element contains information related to the instantiated MSNCs. This includes information related to:

a) the protocol aspects for specific layers of the established MSNC, such as EVPN, VPLS, etc.; and

b) identifiers for virtual private network segments, e.g. RDs for the MPLS backbone. This information is available in the *msncLayerProtocol* attribute of the *Msnc* information element.

As described in clause 6.3.1.2, EVPN enables CEs to provide "multihoming" feature where a CE is connected to multiple PEs in the WAN, providing redundancy, and enabling load balancing. Information about these Ethernet links realizing the multihoming feature can be provided using the *lag* attribute of *MscsEndpointData* information element. Additionally, information about the links being in *active* or *standby* configuration could also be included in the *lag* attribute, to further specify whether the link aggregation is being used to offer load-balancing or redundancy. The multihomed Ethernet links collectively make up an Ethernet Segment (ES), which is identified by an Ethernet Segment Identifier (ESI). This ESI correspond to the *connectivityServiceEndpointId* attribute of the *MscsEndpointInfo* information element in ETSI GS NFV-IFA 032 [i.34].

6.4.1.2 Considerations for L3VPN service enablers

In the context of inter-site connectivity, MPLS and MP-BGP technologies are described in clauses 6.3.2.1 and 6.3.2.2 respectively as enablers for providing L3VPN services. The present clause analyses the existing interfaces and IEs specified in ETSI GS NFV-IFA 032 [i.34] considering the aforementioned enabling technologies. The MSCS Management interface, specified in ETSI GS NFV-IFA 032 [i.34], allows an authorized consumer to configure multi-site connectivity service between multiple sites (NFVI-PoPs) through different operations such as Create MSCS, Update MSCS, Terminate MSCS, etc.

On the MSCS level, the *mscsLayerProtocol* attribute in the *MscsData* information element, specified in ETSI GS NFV-IFA 032 [i.34], can be configured to support the applicable L3VPN technology, e.g. l3vpn with specific protocol technologies such as MPLS, BGP, etc. Additional information elements such as *MscsProfile* and *MscsEndpointData* can be used to carry further VPN configuration parameters, particularly on the mode of connectivity, directionality, link aggregation, network addressing, QoS, bandwidth etc. Additionally, QoS and bandwidth requirements can be configured using the *qosMetric*, *bandwidthIn* and *bandwidthOut* attributes of the *MscsProfile* information element.

networkAddressing attribute in *MscsEndpointData* information element carries information about the network addressing configuration applicable to the MSCS endpoint. For instance, it can include information about identifiers of the virtual private network segments, e.g. RDs in case of MPLS.

The L3VPN service enablers, described in clause 6.3, are closely linked with the technologies realizing the WAN backbone between the different sites. A MSCS abstracts the details of network connectivity that is realized by one or more MSNCs. On the MSNC level, the *Msnc* information element, specified in ETSI GS NFV-IFA 032 [i.34], encapsulates the information specific to the multi-site network connection realizing the MSCS, particularly the *msncLayerProtocol* can be used to describe protocol aspects for particular layer of the MSNC, such as MPLS, etc. Information about encapsulation technologies, such as MPLS-in-IP, MPLS-in-GRE, etc., may also be provided in the *msncLayerProtocol* attribute if encapsulation is required as discussed in clause 6.3.2.1.

MP-BGP, described in clause 6.3.2.2, is an important routing technology for L3 VPNs realizing inter-site connectivity. MP-BGP is mostly involved inside the Provider Network for distribution of MPLS labels and the VPN-IPv4 routes among the CE routers belonging to the same VPN. Interfaces and information elements specified in ETSI GS NFV-IFA 032 [i.34] do not necessarily cover the configuration aspects of BGP or MP-BGP inside the provider network. However, if the CE and PE are BGP peers, a set of RTs may be provided to the PE in order to configure its *Export Targets,* discussed in clause 6.3.2.2, using the MSCS Management interface of ETSI GS NFV-IFA 032 [i.34]. Currently, the relevant information elements in ETSI GS NFV-IFA 032 [i.34] do not support this configuration.

6.4.2 OAM considerations

6.4.2.1 Overview

Regarding inter-site connectivity management, the network elements to be managed are NFVI-PoP network gateway (e.g. CE), and network elements in the transport/WAN, typically referred as Provide Edge (PE) and Provider nodes (P) devices. Connectivity services created over the transport/WAN are managed by the WIM, while connectivity over the NFVI-PoP network gateway can be managed either by the WIM or the VIM; this depends on the network demarcation point considered (see also ETSI GS NFV-SOL 005 [i.33], Annex E for a relevant analysis). In case the NFVI-PoP network gateway is managed by NFV MANO, the solutions described in clause 5.4.2.3.3 are applicable.

Figure 6.4.2.1-1 depicts a visual representation of the end-to-end network.

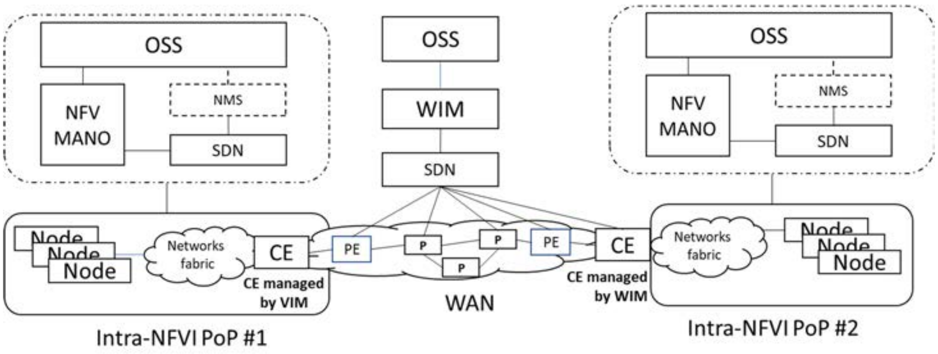

Figure 6.4.2.1-1: Managed and management entities for the case of Inter-NFVI PoP connectivity

WIM can reside inside or outside NFV MANO. In Figure 6.4.2.1-1 the latter case is depicted. In case WIM is part of NFV MANO, then the NFVO interfaces with the WIM for the management of MSCS connectivity on the WAN. In case WIM is not part of NFV MANO, OSS/BSS interfaces with the WIM for the management of MSCS connectivity on the WAN.

6.4.2.2 Challenges

Like in the case of intra-site connectivity, the details on how the WIM communicates in the southbound with the different network management systems (e.g. SDN controllers) and how these can support and apply the relevant OAM operations, are not covered by referenced ETSI NFV specifications.

ETSI NFV specifications lack guidelines and/or relevant specification to cover OAM&P aspects in the case of the following areas of network connectivity:

- Provide Edge (PE) and Provider node (P) management:
    - Network device management.
    - Management of underlay and overlay networks.
- NFVI-PoP network gateway (e.g. CE) management:
    - Network device.
    - Management of underlay and overlay networks.

Network device management refers to operations like device power management, device level FM/PM, configuration, and SW management like backup, restore, upgrade operations, etc. Management of underlay and overlay networks refers to protocol configuration for connectivity, network level FM/PM, connectivity testing operations, etc.

6.4.2.3 Proposed solutions

6.4.2.3.1 Solutions for Challenge #1: Provider Edge (PE) and Provider (P) management

6.4.2.3.1.1 Introduction

Regarding the underlay and overlay network OAM management for inter-site, ETSI GS NFV-IFA 032 [i.34] specifies the MSCS management interface, the capacity management interface, the fault management interface, and the performance management interface. ETSI GR NFV-SOL 017 [i.22] describes the relevant operations exposed by the WIM and investigates whether stage 3 solutions like ONF TAPI and IETF ACTN can provide support of the requirements specified in ETSI GS NFV-IFA 032 [i.34]. How the WIM communicates with the different network management systems (e.g. SDN controllers) and how these can support and apply the relevant OAM operations, are not covered by the referenced ETSI NFV specifications.

6.4.2.3.1.2 Solution SOL-1.1: SDN controller exposing standard interface(s)

Like Solution SOL-2-2 described in clause 5.4.2.3.2.2, WIM interacts with SDN control mechanisms which are responsible for the management of the overlay connectivity realizing the different MSCSs and the corresponding MSNCs. SDN controllers interact with WIM over standardized interfaces.

6.4.2.3.1.3 Solution SOL-1.2: device-exposed standard protocol(s)

Like Solution SOL-2-4 described in clause 5.4.2.3.2.4, devices expose standard interfaces and data models tailored to network configuration. The client of the exposed interfaces can be an SDN controller or another control and management entity.

6.4.2.3.1.3 Solution SOL-1-3: BMC-based management via PIM

This solution is similar to Solution SOL-2.1: BMC-based management via PIM described in clause 5.4.2.3.1.2. The Fabric, Port and Switch managed objects from the DMTF's Redfish DSP2046 [i.66] are reused for network devices conforming the transport/WAN.

6.4.2.3.2 Solutions for Challenge #2: NFVI-PoP network gateway management

6.4.2.3.2.1 Introduction

In case the NFVI-PoP network gateway is managed by NFV-MANO of functional blocks/functions other than the WIM, then the solutions described in clause 5.4.2.3.3 are applicable. In case the NFVI-PoP network gateway is managed by WIM, either as part of or external to NFV-MANO, the following solutions can be considered.

6.4.2.3.3.2 Solution SOL-2.1: SDN controller exposing standard interface(s)

This solution is similar to Solution SOL-1.1 described in clause 6.4.2.3.1.2.

6.4.2.3.3.3 Solution SOL-2.2: PIM

This solution is similar to Solution SOL-1.3: BMC-based management via PIM described in clause 6.4.2.3.1.3.

6.4.2.3.3.4 Solution SOL-2.3: device-exposed standard protocol(s)

This solution is similar to Solution SOL-2.4 described in clause 5.4.2.3.2.4, where devices expose standard interfaces and data models tailored to network configuration. The client of the exposed interfaces can be an SDN controller or another control and management entity.

6.4.2.4 Solutions Evaluation

Table 6.4.2.4-1 provides pros/cons analysis of the solutions described for inter-site OAM management.

Table 6.4.2.4-1: Solutions evaluation for Intra-site OAM management

Challenge	Solution	Pros	Cons	Comment
#1: Provider Edge (PE) and Provider (P) management	SOL-1.1: SDN controller exposing standard interface(s)	Like in SOL-2.2 in clause 5.4.2.4.	Like in SOL-2.2 in clause 5.4.2.4.	Can also work in conjunction with PIM by delegating physical device management to PIM.
	SOL-1.2: device-exposed standard protocol(s)	• Leverage the use of standard protocols in the northbound of the device to facilitate integration and maintenance.	Like in SOL-2.4 in clause 5.4.2.4, with the descriptions tailored to WIM operations instead the VIM.	
	SOL-1.3: BMC-based management via PIM	• Technology-ready solutions can be exploited regarding BMC operations and exposed interfaces (e.g. Redfish). • Both system/device aspects and network aspects of the P and PE elements can be managed from a single entity thus simplifying the management architecture.	• Like in SOL-2.1 in clause 5.4.2.4. • BMC-based solutions are more tailored for compute, storage and network resources which are part data centers and cloud environments.	
#2: NFVI-PoP network gateway management	SOL-2.1: SDN controller exposing standard interface(s)	Like in SOL-2.2 in clause 5.4.2.4.	Like in SOL-2.2 in clause 5.4.2.4.	Can also work in conjunction with PIM by delegating physical device management to PIM.
	SOL-2.2: PIM	Like in SOL-1.1 in clause 5.4.2.4 tailored to the Gateway system operations.	Like in SOL-1.1 in clause 5.4.2.4.	
	SOL-2.3: device-exposed standard protocol(s)	Like in SOL-2.4 in clause 5.4.2.4.	Like in SOL-2.4 in clause 5.4.2.4, with the descriptions tailored to WIM operations instead the VIM.	

7 Recommendations

7.1 Introduction

The present clause documents recommendations about potential enhancements or modifications to existing ETSI NFV specifications. The recommendations are derived based on the analysis performed for intra-site and inter-site connectivity services in clause 5.4 and clause 6.4 respectively. Recommendations related to different aspects of ETSI NFV framework and specifications are categorized below:

- recommendations related to NFV architectural framework (refer to clause 7.2);
- recommendations related to functional aspects (refer to clause 7.3);
- recommendations related to NFV descriptors and other artefacts (refer to clause 7.4);
- recommendations related to interfaces and associated information elements (refer to clause 7.5); and
- recommendations related to OAM solutions (refer to clause 7.6).

7.2 Recommendations related to the NFV architectural framework

Based on the analysis provided in the previous clauses, no new architectural element (i.e. function or functional block) is recommended to be introduced to support functionalities related to network management aspects for both the intra-NFVI-PoP and inter-NFVI-PoP network connectivity types:

- For the Intra-NFVI-PoP case, depending on the solution selected, existing NFV-MANO management entities like the VIM and the PIM can be used to support OAM operations considering Node level, network fabric level and NFVI-PoP network gateway management aspects. See clause 5.4.2 for a description of the solutions related to intra-NFVI-PoP OAM aspects.
- For the Inter-NFVI-PoP case, depending on the solution selected, existing NFV-MANO management entities like the WIM can be used to support OAM operations considering management aspects for PE and P devices. Capability of WIM to manage connectivity services and configuration up to the NFVI-PoP network gateway is acknowledged by the ETSI GS NFV-SOL 005 [i.33]. See clause 6.4.2 for a description of the solutions related to inter-NFVI-PoP OAM aspects.

Other control entities like SDN controllers which are part of the OAM solutions described in clauses 5.4.2 and 6.4.2, are not considered to be part of the NFV-MANO framework itself. However, it is recommended that the NFV-MANO architectural framework describes and specifies the relationship between NFV-MANO and SDN controllers that become part of the NFVI, both intra-NFVI-PoP and inter-NFVI-PoP.

7.3 Recommendations related to functional aspects

The present clause provides recommendations related to functional aspects of the different NFV-MANO entities like functions and functional blocks.

Table 7.3-1 provides the recommendations related to functional aspects.

Table 7.3-1: Recommendations related to functional aspects

Identifier	Recommendation description	Comments
nfv-conn.func.001	It is recommended that a requirement be specified for the PIM to support the management of networking capabilities of the NFVI Nodes.	Management interactions can be supported via BMCs.
nfv-conn.func.002	It is recommended that a requirement be specified for the PIM to support the management of the NFVI-PoP network elements.	Management interactions can be supported via BMCs for device-level management of the network elements comprising the NFVI-PoP network fabric, and via exposed standard interfaces by SDN controllers for the network protocol management aspects of the network fabric.
nfv-conn.func.003	It is recommended that a requirement be specified for the VIM to support the management of virtual networks instantiated over the NFVI-PoP network fabric by consuming standard interfaces exposed by other functions.	As specified in referenced ETSI NFV specifications, the VIM already supports the management of virtual networks within the NFVI-PoP, to be also realized via standard interfaces exposed by SDN controllers responsible for the network protocol management aspects of network overlays over the network fabric.
nfv-conn.func.004	It is recommended that a requirement be specified for the PIM to support device-level management of the NFVI-PoP network gateway. See notes 1 and 2.	This recommendation covers the device-level management of the NFVI-PoP network gateway and considers the case when NFVI-PoP network gateway is managed or co-managed from the NFVI-PoP demarcation according to Annex E of ETSI GS NFV-SOL 005 [i.33]. See nfv-conn.func.005 for network related aspects.
nfv-conn.func.005	It is recommended that a requirement be specified for the VIM to support the management of networking capabilities of the NFVI-PoP network gateway via standard protocols exposed by the network gateway. See note 2.	The NFVI-PoP network gateway standard protocols can also be consumed directly by the NFVO as documented in Annex E of ETSI GS NFV-SOL 005 [i.33].
nfv-conn.func.006	It is recommended that a requirement be specified for the NFV architectural framework to support interactions between the PIM and VIM, and between PIM and CCM regarding the management of NFVI nodes. See notes 3 and 4.	See ETSI GS NFV-IFA 053 [i.72] for the specification of the interface and interactions between PIM and VIM.
nfv-conn.func.007	It is recommended that a requirement be specified for the NFV architectural framework to support interactions between the PIM and VIM, and between PIM and CCM regarding the management of NFVI-PoP network elements. See notes 3 and 4.	See ETSI GS NFV-IFA 053 [i.72] for the specification of the interface and interactions between PIM and VIM.
nfv-conn.func.008	It is recommended that a requirement be specified for the NFV architectural framework to support interactions between the PIM and VIM, and between PIM and CCM regarding the management of NFVI-PoP network gateway. See notes 3 and 4.	See ETSI NFV-IFA053 [i.72] for the specification of the interface and interactions between PIM and VIM.
nfv-conn.func.009	It is recommended that a requirement be specified for the WIM to support device-level management of the NFVI-PoP network gateway. See notes 1 and 2.	Management interaction can be performed via the PIM (see nfv_conn.func.004).
nfv-conn.func.010	It is recommended that a requirement be specified for the WIM to support the management of networking capabilities of the NFVI-PoP network gateway via standard protocols exposed by the NFVI-PoP network gateway. See note 2.	Management interactions are supported either via SDN controllers exposing standard interfaces or interfaces exposed directly by the NFVI-PoP network gateway.
nfv-conn.func.011	It is recommended that a requirement be specified for the WIM to support the management and configuration of PE and P network elements.	Management interactions are supported via SDN controllers exposing standard interfaces.

Identifier	Recommendation description	Comments
NOTE 1: In a PIM-based solution the NFVI-PoP network gateway can be managed either by WIM/PIM or VIM/PIM or PIM/OSS. NOTE 2: The management system for the Intra-NFVI-PoP network configuration of the NFVI-PoP network gateway can be different from the management system used for the Inter-NFVI-PoP network configuration of the same NFVI-PoP network gateway system. NOTE 3: PIM services can also be consumed by other management functions, e.g. comprising the OSS. NOTE 4: In containerized environments CCM can interact with PIM for the management of the networking aspects within the NFVI-PoP for the case of bare-metal CIS clusters.		

7.4 Recommendations related to NFV descriptors and other artifacts

The present clause provides recommendations related to NFV descriptors and other artifacts. Table 7.4-1 provides the relevant recommendations.

Table 7.4-1: Recommendations related to NFV descriptors and other artifacts

Identifier	Description	Applicability	Comments
nfv-conn.des.001	It is recommended that the information elements related to the VNF Descriptors (e.g. *the Subport* information element) specified in ETSI GS NFV-IFA 011 [i.70] be enhanced to support identification of unique traffic flows in case the NVGRE connectivity service is used for intra-site connectivity.	Intra-site	See analysis in clause 5.4.1.1 and recommendation nfv-conn.iie.001 in Table 7.5-1.
nfv-conn.des.002	It is recommended that the parameters in the data model for the descriptor-based *Virtualised Network Management* specified in ETS GS NGV-IFA 005 [i.7] be enhanced to support a more granular description of the network characteristics (e.g. for IPSec as an L3VPN enabler technology).	Intra-site and Inter-site	See analysis in clause 5.4.1.2 and recommendation nfv-conn.iie.004 in Table 7.5-1 (see note).
NOTE: The case of IPsec for inter-site connectivity has not been analysed separately, however the same analysis and principles of operation are applicable like the ones presented in clause 5.4.1.2.			

7.5 Recommendations related to interfaces and information elements

The present clause provides recommendations related to relevant interfaces and information elements specified in IFA and SOL documents, based on the analysis provided in clauses 5.4.1 and 6.4.1. The recommendations cover both the cases of connectivity services used within the same NFVI-PoP (intra-site) and between two or multiple NFVI-PoPs (inter-site).

Table 7.5-1 provides the recommendations related to interfaces and information elements, also indicating the applicability (intra-site or inter-site, or both) for each recommendation.

Table 7.5-1: Recommendations related to interfaces and information elements

Identifier	Description	Applicability	Comments
nfv-conn.iie.001	It is recommended that the *TrunkSubport* IE specified in ETSI GS NFV-IFA 005 [i.7] be enhanced to support identification of unique traffic flows in case the NVGRE connectivity service is used for intra-site connectivity.	Intra-site	Reference points and interfaces considerations for L2VPN service enablers of intra-site connectivity are described in clause 5.4.1.1.
nfv-conn.iie.002	It is recommended that the *segmentationType* attribute of the *TrunkSubport* IE specified in ETSI GS NFV-IFA 005 [i.7] be extended to allow additional values such as "VXLAN" and "NVGRE".	Intra-site	Reference points and interfaces considerations for L2VPN service enablers of intra-site connectivity are described in clause 5.4.1.1.
nfv-conn.iie.003	It is recommended that an additional attribute is introduced (e.g. in the *TrunkSubport* IE specified in ETSI GS NFV-IFA 005 [i.7]), to indicate if the discovery of NVEs happens automatically (e.g. through the use of a routing protocol, such as iBGP).	Intra-site	Reference points and interfaces considerations for L2VPN service enablers of intra-site connectivity are described in clause 5.4.1.1.
nfv-conn.iie.004	It is recommended that *the Virtualised Network Resource Management* interface specified in ETSI GS NFV-IFA 005 [i.7] and the *MSCS Management interface* specified in ETSI GS NFV-IFA 032 [i.34] be enhanced to support a more granular description of the network characteristics (e.g. for IPSec as an L3VPN enabler technology) for intra-site and inter-site connectivity.	Intra-site and Inter-site	Reference points and interfaces considerations for L3VPN service enablers of intra-site connectivity are described in clause 5.4.1.2. See note 1.
nfv-conn.iie.005	It is recommended that the *lag* attribute described in ETSI GS NFV-IFA 032 [i.34] be enhanced to provide additional information about link aggregation, i.e. whether the link aggregation is being used to offer load-balancing or redundancy.	Inter-site	Reference points and interfaces considerations for L3VPN service enablers of inter-site connectivity are described in clause 6.4.1.2.
nfv-conn.iie.006	It is recommended that the relevant information elements in ETSI GS NFV-IFA 032 [i.34] be enhanced to support providing set of RTs to the PE for configuring its Export Targets.	Inter-site	Reference points and interfaces considerations for L3VPN service enablers of inter-site connectivity are described in clause 6.4.1.2. See note 2.
NOTE 1: This could mean introduction of additional information elements and attributes to carry information related to encryption, authentication, etc. NOTE 2: This type of configuration is applicable if the CE and PE are BGP peers.			

7.6 Recommendations for OAM solutions

The present clause documents recommendations for OAM solutions proposed in clause 5.4.2.3 for the case of intra-site connectivity OAM management and clause 6.4.2.3 for the case of Inter-site connectivity OAM management.

The recommendations are based on the evaluation analysis provided in clause 5.4.2.4 and clause 6.4.2.4 respectively.

Table 7.6-1 provides recommendations for Intra-site OAM management solutions to be further considered for normative work.

Table 7.6-1: Recommendations for Intra-site OAM management solutions

Challenge	Solution Identifier	Description	Recommended	Comment
#1: NFVI Node device level management	SOL-1.1	BMC-based management via PIM	YES	
	SOL-1.2	OS-based management by VIM through standard protocols	NO	
	SOL-1.3	Integration of device management tools by VIM	NO	
#2: Intra-NFVI-PoP network fabric management	SOL-2.1	BMC-based management via PIM	Partially	Recommended, in case it is working in conjunction with an SDN controller-based solution (SOL-2.2). PIM will be responsible for device level management and the SDN controller responsible for network protocol management aspects. See note.
	SOL-2.2	SDN controller exposing standard interface(s)	YES	See note.
	SOL-2.3	Plug-in based	NO	
	SOL-2.4	Device-exposed standard protocol(s)	Partially	Recommended, in case it is working in conjunction with an SDN controller-based solution (SOL-2.2). See note.
#3: NFVI-PoP network gateway management	SOL-3.1:	device-exposed standard protocol(s)	YES	Considering that no other alternatives are analysed.
NOTE: SOL-2.1 can work in parallel with SOL-2.2, and SOL-2.4. In such a scenario PIM (SOL-2.1) can be responsible for device level management, while the SDN controller (SOL-2.2) can be used to manage network level aspects through interaction with the gateway system over standardized interfaces (SOL-2.4).				

Table 7.6-2 provides recommendations for Inter-site OAM management solutions to be further considered for normative work.

Table 7.6-2: Recommendations for the solutions for Inter-site OAM management

Challenge	Solution Identifier	Description	Recommended	Comment
#1: Provider Edge (PE) and Provider (P) management	SOL-1.1	SDN controller exposing standard interface(s)	YES	
	SOL-1.2	Device-exposed standard protocol(s)	Partially	Recommended, in case it is working in conjunction with an SDN controller-based solution (SOL-1.1). See note 1.
	SOL-1.3	BMC-based management via PIM	NO	
#2: NFVI-PoP network gateway management	SOL-2.1	SDN controller exposing standard interface(s)	YES	
	SOL-2.2	PIM	Partially	Recommended, in case it is working in conjunction with an SDN controller-based solution (SOL-2.2). PIM will be responsible for device level management and the SDN controller responsible for network protocol management aspects. See note 2.
	SOL-2.3	device-exposed standard protocol(s)	YES	See note 2.
NOTE 1: For the case of PE and P management SOL-1.1 can work in parallel with SOL-1.2. In such a scenario device-exposed standard protocols can be responsible for device level management, while the SDN controller can be used to manage network level aspects through interaction with the gateway system over standardized interfaces. NOTE 2: Similar case as note 1 but applies to the NFVI-PoP network gateway management instead of PE and P management.				

8 Conclusion

The present document studies the network connectivity aspects within NFV framework by reporting on different networking technologies that enable intra-site and inter-site connectivity in NFV. L2 and L3 VPN technologies are analysed with regards to the relevant NFV reference points and interfaces. With respect to network connectivity integration and operationalization, OAM&P aspects of network resources are also analysed in the present document. The present document identifies challenges related to OAM&P of network resources for both intra-site and inter-site network fabric and proposes potential solutions which are further evaluated. The analyses performed in the present document are:

- Analysis related to intra-site connectivity services and enabling technologies (refer to clause 5.4).
- Analysis related to inter-site connectivity services and enabling technologies (refer to clause 6.4).

Based on the analyses, a set of recommendations is derived in the present document. These recommendations indicate the need to perform additional normative specification work to enhance the capabilities of the NFV architectural framework to better support network connectivity and operationalization aspects. The areas for which additional normative work is identified are:

- Recommendations related to NFV architectural framework (refer to clause 7.2).
- Recommendations related to functional aspects (refer to clause 7.3).
- Recommendations related to NFV descriptors and other artefacts (refer to clause 7.4).
- Recommendations related to interfaces and associated information elements (refer to clause 7.5).
- Recommendations related to OAM solutions (refer to clause 7.6).

Annex A: Examples about Network infrastructure architecture and virtual networks

A.1 Networking infrastructure in OpenStack without SDN support

In the basic OpenStack deployment, networking services are executed on the OpenStack controller and compute node. The controller node runs network management, network plugins, networking agents (e.g. L3 agent, DHCP agent, etc.). The compute and controller nodes run networking agents (e.g. Open vSwitch agent, Linux bridge agent, SRIOV NIC Switch agent, MacVTap agent, etc.) to manage virtual networks connecting the instantiated Virtual Machines (VMs). As shown in Figure A.1-1, the controller and compute nodes are attached to the switch fabric of the leaf and spine architecture [i.6]. The DC gateway node works as a border router connecting to external networking infrastructure (e.g. Wide Area Network (WAN), Internet, VPN, etc.).

There are two types of virtualised network, provider and tenant networks, supported in OpenStack. The provider network is a Layer 2 network instantiated by the administrator. The provider network is mapped on to the physical network and publicly routable to and from other external network through the DC gateway node. On the other hand, tenant network is a private Layer 2 network instantiated by tenant users. Each tenant network is isolated from other provider and tenant networks. By attaching to the OpenStack Neutron routers, the tenant networks become reachable to each other.

A provider network can be configured as a flat network configuration or VLAN network configuration. As shown in Figure A.1-1, in the flat configuration, a single Layer 2 network is shared between multiple virtual machines and the Layer 2 network is terminated by the routing instance (e.g. VRF) in the DC gateway node. In this case, all attached VMs share the same broadcast domain.

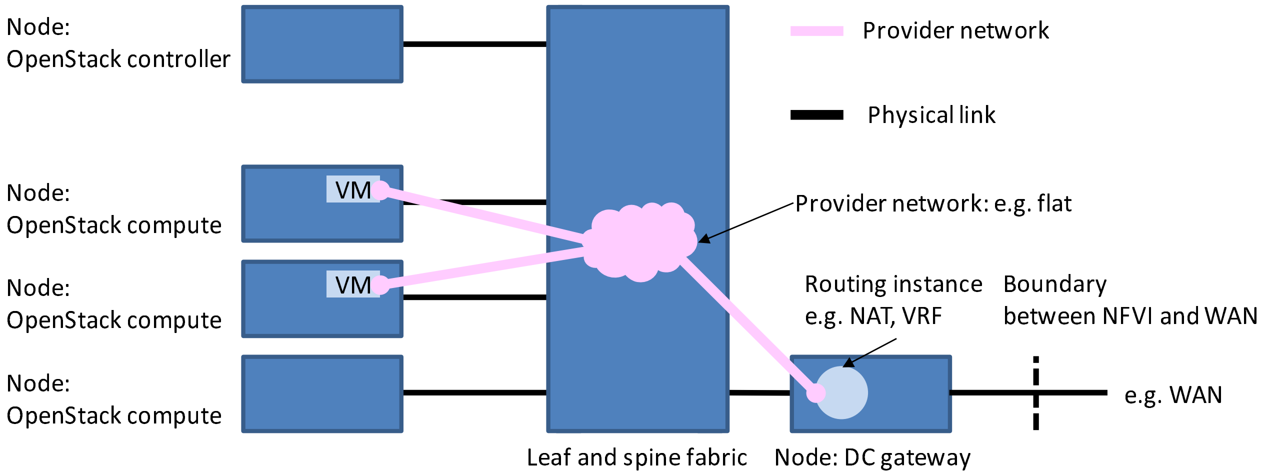

Figure A.1-1: Provider network with flat configurations

Figure A.1-2 shows another case where multiple provider networks are realized based on VLAN configurations. Those Layer 2 networks are logically isolated on the leaf and spine switch fabric, and are terminated by individual routing instances on the DC gateway node. Traffic from multiple VLANs can traverse over a single physical link between the compute node and a leaf switch if switch ports are configured as trunks.

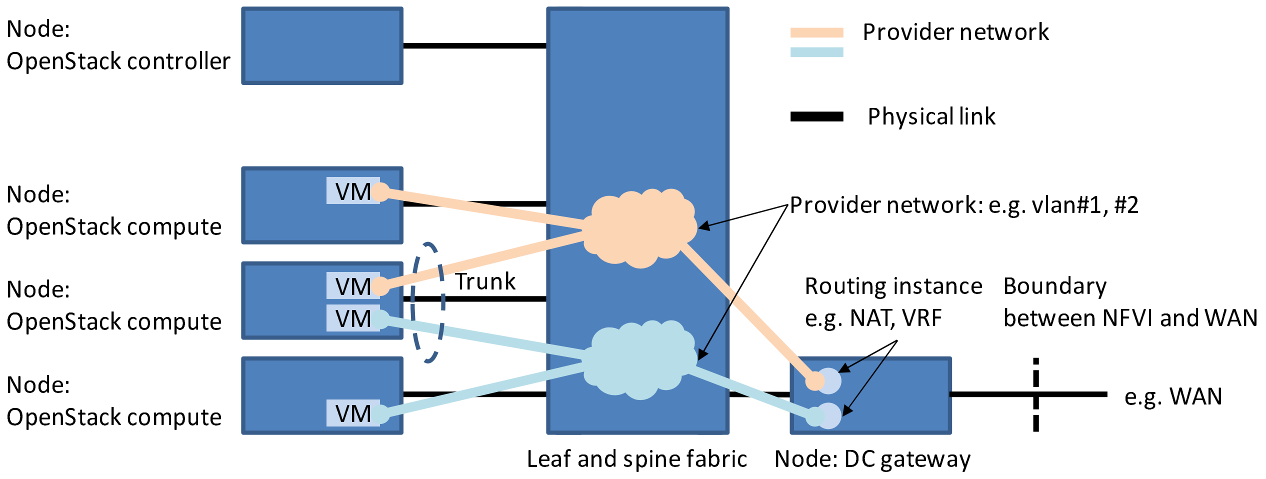

Figure A.1-2: Provider network with VLAN configurations

Figure A.1-3 shows another connectivity model where a tenant network is realized with overlay networking technologies such as GRE or VXLAN. As described, the tenant network is connected with a Neutron router to enable reachability to other provider and tenant networks. The Neutron router to which the tenant network attaches is configured as a gateway between the tenant and provider networks.

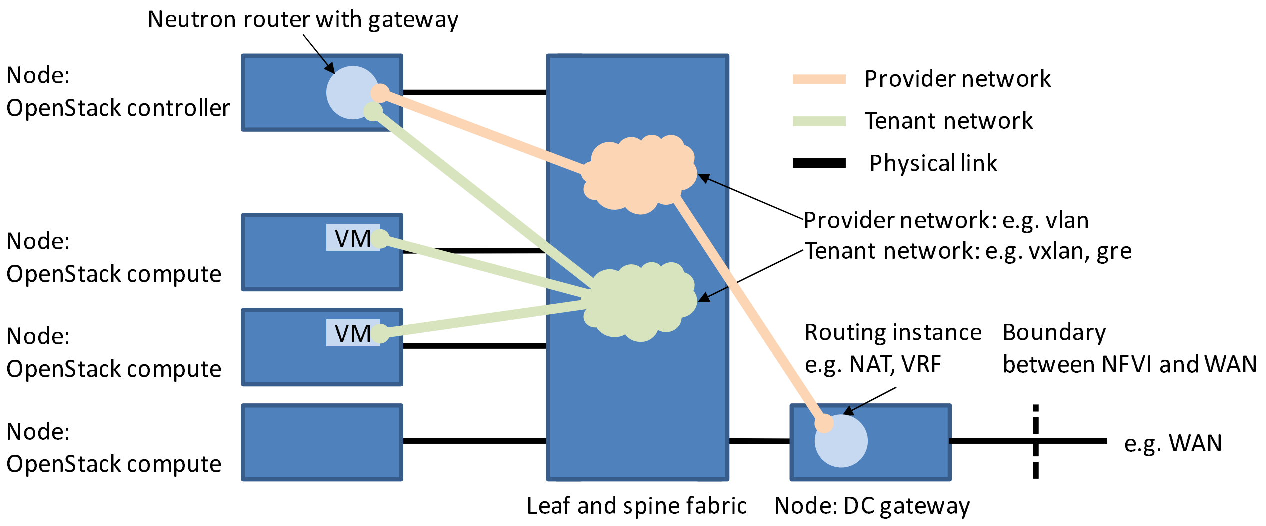

Figure A.1-3: Provider and tenant network with Neutron router

A.2 Networking infrastructure in OpenStack with SDN support

In an OpenStack environment, SDN controllers are used to provide simplified and centralized network management for distributed network devices. The SDN basically brings generic operations and management for virtualised networks, subnets and ports. In addition, the SDN brings features to configure and manage physical network devices with routing, VPN, traffic engineering features, etc. The additional features complement the network management in OpenStack.

As depicted in Figure A.2-1, the networking service in OpenStack is provided by the Neutron service. The OpenStack implementation of the Neutron service has a plugin that allows to support a wide variety of underlying network agents. In the case of OpenStack Neutron Service with SDN support, the SDN controllers interact with the Neutron server through the plugins and the northbound interfaces of the SDN controllers. The SDN controllers in turn interact with the network agents via their southbound interfaces. In this case the SDN controller (e.g. OpenDaylight [i.13]) complements the Neutron service and works below the plugin mechanism. With other implementations (e.g. Tungsten Fabric [i.14]) the SDN controller substitutes the OpenStack Neutron service by implementing an alternative network management service that could provide additional network features that the OpenStack Neutron service might not implement.

In general, the SDN controllers offer RESTful APIs on their northbound interface. Thus, simplified and centralized network management interfaces can be exposed as GUIs for network management staff as well as APIs towards external applications and orchestrators. On the other hand, SDN controllers support variety of protocols over the southbound interfaces to configure and manage agents running in OpenStack compute/ network nodes, and physical network devices with routing, VPN and traffic engineering features. Some of those protocols are:

- NETCONF [i.8] provides mechanisms to install, manipulate, and delete the configuration of network devices.
- OVSDB [i.9] provides Open vSwitch Database Management Protocol. The Open vSwitch is an open-source software switch designed to be used as a vSwitch (virtual switch) in virtualized server environments.
- OpenFlow [i.10] is used to control and manage OpenFlow switches from a remote OpenFlow controller feature.
- BGP [i.11] is an inter-autonomous system routing protocol to exchange network reachability information with other BGP systems.
- PCEP [i.12] is Path Computation Element (PCE) Communication Protocol for communications between a Path Computation Client (PCC) and a PCE, or between two PCEs.

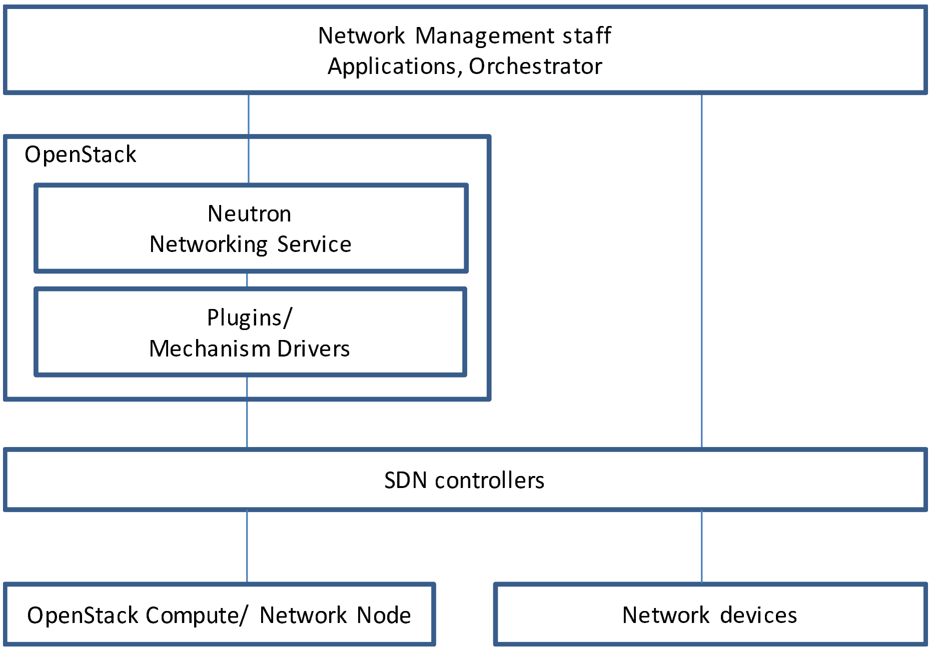

Figure A.2-1: SDN Controller in the perspective of OpenStack Environment

Annex B:
Change History

Date	Version	Information about changes
V0.1.0	12.2019	• NFVIFA(19)000918r2 - IFA035 - GR Structure and ToC • NFVIFA(19)000956 - IFA035 Scope Statement
V0.2.0	04.2020	• NFVIFA(20)000207r1 - IFA035: Network infrastructure architecture and virtual networks • NFVIFA(20)000208 - IFA035: Annex X (informative): Examples about network infrastructure architecture and virtual networks
V0.3.0	07.2020	• NFVIFA(20)000335r2 - IFA035: Summary of Intra-site Connectivity Technologies • NFVIFA(20)000292 - IFA035: A.x Networking infrastructure in OpenStack with SDN support
V0.4.0	11.2020	• NFVIFA(20)000541r1 - IFA035 Network isolation and gateway • NFVIFA(20)000613r2 - IFA035 Connectivity service for L3VPN Service • NFVIFA(20)000676r1 - IFA035 Connectivity service for L2VPN Service
V0.5.0	01.2021	• NFVIFA(20)000827 - IFA035 Overview of Carrier Ethernet Service
V0.6.0	09.2021	• NFVIFA(21)000788r1 - IFA035 - Clause 5.2 Network infrastructure and VNs in DC
V0.7.0	11.2021	• NFVIFA(21)000900r1 - IFA035-Clause 4 Carrier Ethernet Service types • NFVIFA(21)000911r2 - IFA035 - Clause 6 intro inter-site Network infrastructure and VNs
V0.8.0	01.2022	• NFVIFA(21)0001019 - IFA035 - Clause 5.3 intra-site L2 VPN service enabler VXLAN • NFVIFA(21)0001020r1 - IFA035 - Clause 5.3 intra-site L2 VPN service enabler NVGRE • NFVIFA(21)0001021r1 - IFA035 - Clause 4.4 MPLS VPN connectivity services
V0.9.0	03.2022	• NFVIFA(22)000122r2 - IFA035 - Clause 5.3 inter-site L2 VPN service enabler VPLS
V0.10.0	06.2022	• NFVIFA(22)000321r1 - IFA035 Clause 4.3.4 - MPLS Packet Format • NFVIFA(22)000349 - IFA035 Clause 5.3.2 - IPsec • NFVIFA(22)000366r1 - IFA035 Clause 6.3.1 - EVPN Overview • NFVIFA(22)000367 - IFA035 Clause 4.4.1 BGP for connectivity technologies
V0.11.0	07.2022	• NFVIFA(22)000397r1 - IFA035 - Clause 5.3 VLAN
V0.12.0	08.2022	• NFVIFA(22)000521r1_IFA035-Clause_5_3_2_iBGP • NFVIFA(22)000533r1_IFA035_-_Clause_5_3_1_NVGRE_clause_update
V0.13.0	10.2022	• NFVIFA(22)000671 - IFA035 - Clause 6 - MPLS Based L3 VPN • NFVIFA(22)000672 - IFA035-Clause 5.4.1 Reference Points and Interfaces Considerations for L2VPN
V0.14.0	01.2023	• NFVIFA(22)000969r1 - IFA035 Clause 4 restructuring and adding links to IETF RFCs • NFVIFA(22)000870 - IFA035-Clause 6.4.1 Reference Points and Interfaces Considerations for L2VPN • NFVIFA(22)000919r2 - IFA035-Clause 5.4.1 Reference Points and Interfaces Considerations for L3VPN in intra-site connectivity service • NFVIFA(22)000925r1 - FEAT19 IFA035 segment routing • NFVIFA(22)000945 - FEAT19 IFA035 recommendations section structure
V0.15.0	03.2023	• NFVIFA(23)000157 - FEAT19 IFA035 SDN Overview • NFVIFA(23)000171 - IFA035 - MPBGP • NFVIFA(23)000130 - Clause 7.5 - Recommendations related to interfaces and information elements
V0.16.0	05.2023	• NFVIFA(23)000158r1 - FEAT19b IFA035 Intra-site networking OAM • NFVIFA(23)000159r1 - FEAT19b IFA035 Inter-site networking OAM • NFVIFA(23)000248r1 - IFA035 Clause 5 Introduction • NFVIFA(23)000249r1 - IFA035 Clause 6 L3VPN Analysis for Inter-site Connectivity • NFVIFA(23)000250r1 - IFA035-Clause 3 Definitions Symbols and Abbreviation • NFVIFA(23)000251r1 - IFA035-Clause 4 Introduction and VPN overview • NFVIFA(23)000252r1 - IFA035-Clause 4 Editor s Notes • NFVIFA(23)000253 - IFA035-Clause 6.3 Editor s note

Date	Version	Information about changes
V0.17.0	06.2023	• NFVIFA(23)000254 - IFA035-Clause 7.4 Recommendations related to NFV descriptors and other artefacts • NFVIFA(23)000261r1 - IFA035 Clause 8 Conclusion • NFVIFA(23)000262r1 - IFA035 Clause 7 Recommendations Introduction • NFVIFA(23)000382r1 - FEAT19 IFA035 OAM additional solution for Inter-site connectivity • NFVIFA(23)000383 - FEAT19 IFA035 OAM solution evaluation for Intra-site connectivity • NFVIFA(23)000384 - FEAT19 IFA035 OAM solution evaluation for Inter-site connectivity • NFVIFA(23)000387 - IFA035-VRF term use in the document • NFVIFA(23)000420r1 - FEAT19 IFA035 OAM recommendations for solutions • NFVIFA(23)000431r1 - FEAT19 IFA035 recommendations related to the NFV architectural framework • NFVIFA(23)000432 - FEAT19 IFA035 recommendations related to the NFV functional aspects • NFVIFA(23)000433 - FEAT19 IFA035 other recommendations • NFVIFA(23)000513r1 - IFA035-Removing editor's note for inter-site IPsec in recommendation Rapporteur actions: - Added OAM&P as an abbreviation in clause 3.3 - Formatting in Table 7.5-1 (conn.iie.003 Description) - change blue font colour to black - Removed duplicate reference of EVE022 in contribution 383 - Changed the heading number from 7.2 to 7.6 while implementing NFVIFA(23)000420r1 - Changed the heading title from 'Recommendations for Solutions' to 'Recommendations for OAM Solutions' for better clarity - Updated the last bullet point in NFVIFA(23)000262r1 as 'Other Recommendations' are removed by NFVIFA(23)000433, instead put a pointer in the last bullet to Recommendations for OAM solutions introduced by NFVIFA(23)000420r1 - Removed duplicate reference of IFA053 in NFVIFA(23)000432 as the reference already existed as [i.72], used the latest referred version v0.1.0 in [i.72] - Updated the last bullet in Clause 8 Conclusion (NFVIFA(23)261r1) for consistency, changed 'OAM&P' to 'OAM' - Removed the placeholder Editor's Note in 7.3 as the recommendations have now been provided by contribution NFVIFA(23)000432 - Added generic text above Table 7.4-1 for consistency: *"The present clause provides recommendations related to NFV descriptors and other artefacts. Table 7.4-1 provides the relevant recommendations."* - Updated table of contents.
V0.17.1	07.2023	• NFVIFA(23)000541r1 - IFA035 WG review clauses 1-2 • NFVIFA(23)000542r1 - IFA035 WG review clauses 3-4 • NFVIFA(23)000564r3 - IFA035 WG review clause 5 • NFVIFA(23)000563r3 - IFA035 WG review clause 6 • NFVIFA(23)000585r2 - IFA035 WG review Clauses 7, 8 Rapporteur actions: - Adjusted font colour of the caption text for Figure 4.4.3.3-1 in clause 4.4.3.3 - Assigned appropriate heading styles to multiple sections - Changed font style and size in tables 7.6-1 and 7.6-2 for uniformity - Updated table of contents.
V0.18.0	08.2023	• NFVIFA(23)000608r1 - FEAT19 IFA035 scope update • NFVIFA(23)000609r2 - FEAT19 IFA035 revision of stable draft Rapporteur actions: - Capitalized the letters while introducing abbreviations in clause 5.4.2.3

History

Document history		
V5.1.1	October 2023	Publication