TECHNICAL REPORT

TR-384
Cloud Central Office Reference Architectural Framework

Issue: 1
Issue Date: January 2018

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Technical Report has been approved by members of the Forum. This Technical Report is subject to change. This Technical Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Technical Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Technical Report, or use of any software code normatively referenced in this Technical Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Technical Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Technical Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Technical Report except to the extent this Technical Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH

RIGHTS ARE DESCRIBED ON THE TECHNICAL REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

In addition to the notation above, the Forum draws attention to the fact that it is claimed that compliance with this Specification may involve the use of a patent ("IPR") concerning sections 4.3.6.2 and 4.3.6.3. The Forum takes no position concerning the evidence, validity or scope of this IPR.

The holder of this IPR has assured the Forum that it is willing to License all IPR it owns and any third party IPR it has the right to sublicense which might be infringed by any implementation of this Specification to the Forum and those Licensees (Members and non-Members alike) desiring to implement this Specification. Information may be obtained from:

Trinity College Dublin
College Green
Dublin 02 Ireland

Attention is also drawn to the possibility that some of the elements of this Specification may be the subject of IPR other than those identified above. The Forum shall not be responsible for identifying any or all such IPR.

The text of this notice must be included in all copies of this Technical Report.

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	17 January 2018	31 January 2018	Georgios Karagiannis, Huawei Technologies	Original

Comments or questions about this Broadband Forum Technical Report should be directed to help@broadband-forum.org.

Editors:
Georgios Karagiannis, Huawei Technologies
Ding Hai, China Unicom

SDN and NFV Work Area Director(s):
George Dobrowski, Huawei Technologies
Chris Croot, BT

CloudCO Project Stream Leader(s):
Yves Hertoghs, VMWare
Ning Zong, Huawei Technologies

TABLE OF CONTENTS

List of Figures

List of Tables

Executive Summary

This Cloud Central Office (CloudCO) document helps drive a key element of the BBF's strategy in enabling new revenue generating-services, which is articulated in the Forum's Broadband 20/20 vision. It encompasses the enabling work required for dramatically faster and more efficient provisioning of these services and their more efficient operation.

Operators want their networks to be adaptable, agile, scalable and dynamic, while reducing costs. Besides reducing Capital Expenses (CAPEX) and Operational Expenses (OPEX), shortened time-to-market is much more desired. Software Defined Networking (SDN) and Network Function Virtualization (NFV) in conjunction with general purpose hardware are two powerful tools that are exploited in the industry to optimize current networks.

BBF is defining Broadband networks and has great opportunities in providing such a common cloud platform that can serve both wireline and wireless networks. In order to reuse network connectivity and provide maximum convergence, the Central Office can be an appropriate location to provide such a common cloud platform. Operators will have the opportunity to run a single network with all varieties of access technologies, and flexibly deploy innovative services. With the flexibility of a cloud platform, time-to-market is expected to be shortened as well. New business models can be created as well, where the Operator can offer Anything-as-a-Service to 3rd parties. 3rd parties can be easily inserted into the CloudCO platform, enabling new business models. It is anticipated that market-paced migration to NFV/SDN enabled services will co-exist with installed systems with long-term development for non-virtualized implementations of an MSBN creating a hybrid legacy/CloudCO network. The CloudCO project will provide a valuable element of the migration process for large and regional Operators alike.

This specification is the first phase in a multiphase project. As such it defines the high level architectural framework. Additional work is underway to include migration, detailed interface definitions, test cases & application notes, use cases & scenarios and reference implementations.

1 Purpose and Scope

1.1 Purpose

1.1.1 Overview

The application of Network Function Virtualization (NFV) and Software Defined Networking (SDN) techniques dramatically change the way broadband networks can be designed and deployed. NFV replaces the special-purpose, dedicated nodes used in legacy networks with general-purpose, off-the-shelf resources for computing, storage and switching wherever possible, implementing network functions as software elements that can be orchestrated and deployed on the general-purpose resources as needed. SDN enables centralized control of traffic flows across the NFV environment, facilitating automated provisioning, fine-grained control of subscriber flows, and faster deployment of new services, among other benefits.

The architectural framework describes the interrelation of network functions where those network functions, or a chain of them, can all be consumed ‘as-a-service’, thereby allowing the rapid development of new subscriber services.

A Cloud Central Office (CloudCO) is a recasting of the Central Office (CO) hosting infrastructure that utilizes SDN, NFV and Cloud technologies to support network functions and is aligned with the BBF 20/20 vision.
In doing so, it radically redefines the architectures of the access and aggregation networks that have developed incrementally in previous Broadband Forum specifications such as TR-101 [2] and TR-178 [5]. This document defines the new, cloud-based architecture and sets the stage for further development.

1.1.2 Purpose of this Document

This specification defines the reference architectural framework of the CloudCO. Functional Modules are defined, but not the details of their internal operation. This framework starts from basic Northbound API capabilities so that function composition, initialization, and management of the CloudCO resources can be supported and created as a byproduct of end-to-end service composition.

It provides interface descriptions, which allows the interconnection of these functional modules in an interoperable manner, and allows the consumption of the CloudCO functionality through the Northbound Application Programming Interface (API).

1.2 Relationship to other Work

There are several projects in the BBF that are related to SDN and virtualization in MSBN. This project will consider inputs, as appropriate, from TR-317 (Networked Enhanced Residential Gateway) [8], TR-328 (Virtual Business Gateway) [12], TR-345 (Broadband Network Gateway and Network Function Virtualization) [9], TR-359 (Framework for Virtualization) [11] and TR-370 (Fixed Access Network Sharing) [13] when architecting the CloudCO.

Significant work on SDN and NFV based architectures is ongoing in industry outside the BBF. ETSI's NFV Industry Specification Group (ISG) has generated a great deal of material on NFV from which this document draws. A large number of Open Source software projects are under development to provide both SDN and NFV components and systems for integration into networks.

Further, the CloudCO is a core part of the BBF Open Broadband Initiative which is a collaborative space for integration, interoperability, testing of open source software, vendor provided software and standards-based implementations.

1.3 Scope

CloudCO re-architects the broadband network using SDN and NFV technologies. The CloudCO's functionality can be accessed through a Northbound API, allowing Operators, or 3rd parties, to consume its functionality, while hiding how the functionality is achieved from the API consumer. In order to achieve this, SDN and NFV techniques will be leveraged, running on a cloud-like infrastructure deployed at Central Offices. In this way the Northbound API offers a Platform-as-a-Service (PaaS) style API.

'Cloud-like' means that the CloudCO architecture would typically leverage Data Center (DC) style equipment, i.e., generalized network switches and generalized Compute Nodes where applicable. These switches, which can be seen as Physical Network Functions (PNFs), enable traffic forwarding at L2 or L3 from access functions (where the subscriber line terminates) towards virtualized network functions and/or towards the Broadband Core Network. This Project does not preclude the use of existing network hardware equipment to provide these forwarding capabilities as long as it exposes the defined interfaces to the rest of the functional modules in the CloudCO architecture framework.
This specification will refer to the term 'fabric' to identify the set of PNFs that interconnect the Compute nodes, access-facing PNFs, and network-facing PNFs. The Compute Nodes will be running virtualization software in order to host these aforementioned Virtual Network Functions (VNFs) as Virtual Machines (VMs) or Containers. Infrastructure Operators can offer these VNFs, or a service chain of VNFs to 3rd parties, enabling the offering of the resulting functionality, as-a-service. Third party VNFs can be onboarded onto the platform, allowing to create Value Added Services.

Note that the term 'CloudCO' is used quite liberally, i.e., the CloudCO can be extended across more than one physical location if necessary. It is not the aim to turn the entire national network into one large CloudCO though. Similarly, many smaller or subtending COs may not have any compute infrastructure, but instead use compute infrastructure in another CO.

The network is simplified by centralizing the selected control plane functions through SDN control of both physical and VNFs, and SDN control of the appropriate Service Graphs through these functions, reducing redundant functions and optimizing service processing flows.

2 References and Terminology

2.1 Conventions

There are no requirements in this document which provide an architectural framework.
The reference figures and associated text describe the key elements expected to be part of the CloudCO implementations.

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document		Title	Source	Year
[1]	BBF TR-069, Amendment 5	*CPE WAN Management Protocol, Amendment 5*	BBF	2013
[2]	BBF TR 101, Issue 2	*Migration to Ethernet-Based Broadband Aggregation; Issue 2*	BBF	2011
[3]	BBF TR-156, Issue 3	*Using PON Access in the context of TR-101, Issue 3*	BBF	2012
[4]	BBF TR-167, Issue 2	*PON-fed TR-101 Ethernet Access Node, Issue 2*	BBF	2010
[5]	BBF TR-178 Issue 1	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[6]	BBF TR-197, Issue 2	*DQS: DSL Quality Management Techniques and Nomenclature, Issue 2*	BBF	2014
[7]	BBF TR-301, Issue 1	*TR-301 Architecture and Requirements for Fiber to the Distribution Point, Issue 1*	BBF	2015
[8]	BBF TR-317	*Network Enhanced Residential Gateway*	BBF	2016
[9]	BBF TR-345	*Broadband Network Gateway and Network Function Virtualization*	BBF	2016
[10]	BBF TR-355	*YANG Modules for FTTdp Management*	BBF	2016
[11]	BBF TR-359, Issue 1	*A Framework for Virtualization, Issue 1*	BBF	2016
[12]	BBF TR-328	*Virtual Business Gateway (vBG)*	BBF	2017
[13]	BBF TR-370	*Fixed Access Network Sharing - Architecture and Nodal Requirements*	BBF	2017

[14]	ETSI GS NFV-MAN 001	*Network Functions Virtualisation (NFV); Management and Orchestration*	ETSI ISG NFV	2014
[15]	ETSI GS NFV-INF 005	*Network Functions Virtualisation (NFV); Infrastructure; Network Domain*	ETSI ISG NFV	2014
[16]	ETSI GS NFV-IFA 005 V2.1.1	*Network Functions Virtualisation (NFV); Management and Orchestration; Or-Vi reference point - Interface and Information Model Specification*	ETSI ISG NFV	2016
[17]	ETSI GS NFV-IFA 006 V2.1.1	*Network Functions Virtualisation (NFV); Management and Orchestration; Vi-Vnfm reference point - Interface and Information Model Specification*	ETSI ISG NFV	2016
[18]	ETSI GS NFV-IFA 007 V2.1.1	*Network Functions Virtualisation (NFV); Management and Orchestration; Or-Vnfm reference point - Interface and Information Model Specification*	ETSI ISG NFV	2016
[19]	ETSI GS NFV-IFA 008 V2.1.1	*Network Functions Virtualisation (NFV); Management and Orchestration; Ve-Vnfm reference point - Interface and Information Model Specification*	ETSI ISG NFV	2016
[20]	ETSI GS NFV-IFA 013 V2.1.1	*Network Functions Virtualisation (NFV); Management and Orchestration; Os-Ma-Nfvo reference point - Interface and Information Model Specification*	ETSI ISG NFV	2016
[21]	RFC 7365	*Framework for Data Center (DC) Network Virtualization*	IETF	2014
[22]	ITU G.988	*Recommendation G.988 and its amendments: ONU management and control interface specification (OMCI)*	ITU	2016

2.3 Definitions

The following terminology is used throughout this Technical Report.

CloudCO	Central Office (CO) Domain that is (1) leveraging SDN and NFV techniques, (2) running on a cloud-like infrastructure deployed at Central Offices and (3) that is accessed through a Northbound API, allowing Operators, or 3rd parties, to consume its functionality, while hiding how the functionality is achieved from the API consumer.
CloudCO Domain	One or more CloudCO Macro-Nodes, orchestrated by a single CloudCO Domain Orchestrator and sharing a common, uniquely addressable CloudCO Northbound Interface.
CloudCO Domain Orchestrator	Manages, controls and orchestrates each CloudCO Domain.

DC SDN Manager & Controller	Directly accesses the Network Function Virtualisation Infrastructure (NFVI) networking resources to implement functions (e.g., L3 routes in the switch fabric), via configuration of the underlying physical network infrastructure.
CloudCO Macro-Node	The ensemble of network, compute, storage, and application components that work together to deliver networking services, located in a single network site (this may comprise remotely located access functions whose backhauling is terminated on that site).
Container	An instance of operating-system level virtualization where the operating system kernel allows the existence of multiple, isolated user-space instances.
eOAM	Ethernet operations, administration and maintenance is the set of protocols for installing, monitoring and troubleshooting Ethernet used in LANs, MANs and WANs.
Hypervisor	Computer software, firmware or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine, or 'host', and each virtual machine is called a guest machine, or 'guest'. The technology involved with this process is referred to as 'full system virtualization'.
Logical Subscriber Link (LSL)	A logical point to point L2 connection between the physical Business Gateway (pBG) and the virtual Business Gateway (vBG).
Management Control Orchestration (MCO) Engine	Component of the CloudCO Domain Orchestrator that expresses a continuum of Management, Control and Orchestration (MCO) tasks as well as CloudCO state transitions and supervision tasks.
NFV Orchestrator (NFVO)	Component of the CloudCO Domain Orchestrator. It has two main responsibilities: (1) the orchestration of NFVI resources across multiple Virtualized Infrastructure Managers (VIMs) and (2) the lifecycle management of network services. For a complete list of NFVO capabilities refer to section 5.4 of [14].
Physical Business Gateway	The equipment located at the business customer premises that contains all hardware-dependent Business Gateway functions that must be performed at the customer premises. It may have a built-in NFVI.
PNF and VNF SDN Managers&Controllers	Responsible for Fault, Configuration, Accounting, Performance, Security (FCAPS) and Flow Control management functionalities respectively for PNFs and VNFs.
Retailer (or Virtual) ISP	An Internet Service Provider (ISP) who purchases resources from a wholesaler in order to offer ISP services to a customer. The retailer ISP does not own the infrastructure directly, but differentiates themselves on the service itself.
Service Instance	One instantiation of a Service on a CloudCO Domain.

Shared UNI	A Shared UNI is a network interconnect between a user and the network spanning the switch fabric through which more than one user side service endpoint is connected to a single network side service attachment point. The aggregate of traffic from the service endpoints is policed at the service attachment point to a traffic contract as if it was a single service. The single network-side service attachment point polices the aggregate of the user service endpoint traffic to a single traffic contract.
User Plane	Defines the part of the router architecture that decides what to do with packets arriving on an inbound interface. In routing, the user plane is sometimes called the data plane.
Virtual Business Gateway	A virtual entity located at the network and/or at the customer site, serving one or more pBG entities, supporting some network and service functions such as IP routing.
VBG system	A system that includes the pBG component at the customer site and the vBG component at the CloudCO and/or at the customer site. It also includes their connection over the LSL as well as their interfaces and management system.
Virtualized infrastructure manager (VIM)	Responsible for controlling and managing the NFVI compute, storage and network resources, usually within one Operator's Infrastructure Domain. For a complete list of VIM capabilities refer to section 5.4 of [14].
Virtual Machine (VM)	Emulation of a computer system, providing the full functionality of a physical computer to the applications running on it.
VNF Manager (VNFM):	Responsible for the lifecycle management of VNF instances. For a complete list of VNFM capabilities refer to section 5.4 of [14].
Wholesaler ISP	A third-party provider handles all of the needs of the end user but is invisible to the end user who only sees the retailer (virtual) ISP.
Workload	An instance of a process or processes running on a VM or container.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3GPP	Third Generation Partnership Project
AAA	Authentication, Authorization, Accounting
ABR	Adaptive Bit Rate
ACS	Auto-Configuration Server
AN	Access Node
API	Application Programming Interface
ARP	Address Resolution Protocol
BAA	Broadband Access Abstraction

BNG	Broadband Network Gateway
BPCF	Broadband Policy Control Framework
BSS	Business Support System
CAPEX	Capital Expenses
CloudCO	Cloud Central Office
COTS	Commercial Off-The-Shelf
CPE	Customer Premises Equipment
CPU	Central Processing Unit
CO	Central Office
CRUD	Create, Read, Update, Delete
C-Tag	Customer -Tag
DBA	Dynamic Bandwidth Allocation
DC	Data Center
DCF	Data Collection Function
DHCP	Dynamic Host Configuration Protocol
DPU	Distribution Point Unit
DSL	Digital Subscriber Line
DSLAM	Digital Subscriber Line Access Multiplexer
DTA	Dynamic Timing Allocation
E2E	End to End
eBNG	Evolved BNG
EMS	Element Management System
eOAM	Ethernet OAM
eNodeB	Evolved Node B
ETSI	European Telecommunications Standards Institute
EVPN	Ethernet VPN
FANS	Fixed Access Network Sharing
FCAPS	Fault, Configuration, Accounting, Performance, Security
FTTA	Fiber To The Antenna
FTTdp	Fiber To The distribution point
GbE	Gigabit Ethernet
GEM	Generic Encapsulation Method
HVAC	Heating, Ventilation and Air Conditioning
HW	Hardware
IGP	Interior Gateway Protocol
I/O	Input/Output
IaaS	Infrastructure as a Service
IGMP	Internet Group Management Protocol

IPoE	Internet Protocol over Ethernet
ISG	Industry Specification Group
ISP	Internet Service Provider
L2	Layer 2
L2VPN	Layer 2 VPN
L3	Layer 3
L3VPN	Layer 3 VPN
L4	Layer 4
LDP	Label Distribution Protocol
LLID	Logical Link identifier
LSL	Logical Subscriber Link
LTE	Long Term Evolution
LAN	Local Access Network
MAC	Media Access Control
MANO	Management and Orchestration
MCO	Management Control Orchestration
MEF	Metro Ethernet Forum
MME	Mobility Management Entity
MP2MP	Multi-point to Multi-point
MPLS	Multi-Protocol Label Switching
mroute	Multicast Route
MSBN	Multi Service Broadband Network
NaaS	Network as a Service
NAT	Network Address Translation
NAPT	Network Address Port Translation
NB	Northbound
NB API	Northbound Application Programming Interface
NBI	Northbound Interface
NERG	Networked Enhanced Residential Gateway
NF	Network Function
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NFVO	NFV Orchestrator
NNI	Network to Network Interface
NT	Network Terminal
NVO3	Network Virtualization Overlays 3
OAM	Operation, Administration, Maintenance
ODN	Optical Distribution Network

OLR	On-Line Reconfiguration
OLT	Optical Line Termination
OMCI	ONU Management and Control Interface
ONU	Optical Network Unit
OPEX	Operational Expenses
OSS	Operations Support System
P2C	Point to Cloud
P2P	Point to Point
PaaS	Platform as a Service
pBG	Physical Business Gateway
PCRF	Policy and Charging Rule Function
PNF	Physical Network Function
P-GW	Packet Data Network Gateway
PHY	Physical
PLOAM	Physical Layer OAM
PON	Passive Optical Network
PPP	Point to Point Protocol
PPPoE	Point to Point Protocol over Ethernet
QoS	Quality of Service
RBAC	Role-Based Access Control
RFS	Resource Facing Service
RG	Residential Gateway
S-GW	Serving Gateway
SBI	Southbound Interface
SDN	Software Defined Networking
SLA	Service Level Agreement
SO	Service Orchestrator
SP	Service Provider
SR	Segment Routing
S-Tag	Service –Tag
TCP	Transmission Control Protocol
ToR	Top of Rack
UID	User Interface Design
UNI	User to Network Interface
VAS	Value Added Service
vBG	virtual Business Gateway
vBNG	virtual BNG
VIM	Virtualized Infrastructure Manager

vEPC	virtual Evolved Packet Core
VM	Virtual Machine
VNF	Virtual Network Function
VNFC	Virtual Network Function Component
VNFM	VNF Manager
VPN	Virtual Private Network
VTEP	VxLAN Tunnel End Point
VxLAN	Virtual Extensible LAN
WAN	Wide Area Network

3 Technical Report Impact

3.1 Energy Efficiency

This Technical Report may impact energy efficiency, as network functions can now be decoupled from existing standalone nodes. Use of generic hardware, as such not optimized for a specific network application, and migration of network functions to more distributed locations could lead to higher energy consumption. However, on-demand allocation of hardware resources and hardware sharing across multiple applications can produce energy gains. This Technical Report does not intend to quantify these opposite effects on energy efficiency.

Regulatory differences related to electrical power, Heating, Ventilation and Air Conditioning (HVAC) and fire protection between traditional COs and datacenters is out-of-scope for this document.

3.2 IPv6

This Technical Report has no impact on IPv6.

3.3 Security

Security provides "a form of protection where a separation is created between the assets and the threat." CloudCO enables the sharing of a common infrastructure between various use cases that may be operated by different departments (e.g., wireline and mobile) or different companies (other service providers, including other network service providers). CloudCO also provides an increased opportunity for Operators to dynamically control the network service behavior, with the use of API's. In addition, it is expected that management and control plane interfaces are protected from security risks (CloudCO relies on an increased separation between the control plane and the forwarding plane).

It is noted that existing threats, safeguards, and enhancements remain applicable to CloudCO deployments, whether in the forwarding or management-control planes. This specification assumes a foundation of current security best practices that have been defined for the existing Multi Service Broadband Network (MSBN). However, some new or amplified concerns also appear and without appropriate precautions, the above conditions could impact a network's security.

The next issue of this specification will address security aspects whenever two entities communicate focusing on the threats between different planes, i.e., the SDN/NFV application plane, control plane, NFV Infrastructure (NFVI) user plane, and management plane Open Source Software (OSS), Virtualized Infrastructure Manager (VIM), orchestrator), and the defined interfaces between these blocks. However, it will not specify the details of security threats and solutions to the sub-components of each plane. In addition, developing the security requirements will also need to consider use cases and scenarios.

The BBF will work in collaboration with other industry organizations to apply and/or adapt protocols for security, for example Internet Engineering Task Force (IETF), European Telecommunications Standards Institute (ETSI) NFV, Open Networking Foundation (ONF), etc.

3.4 Privacy

A multi-tenant CloudCO hosts functionality for a set of actors with potentially competing interests that the CloudCO will be required to isolate from each other. At the same time it is required to enable business interactions between the same set of actors requiring careful design of the points of contact.

A multi-tenant CloudCO is a system of sufficient complexity that it will expose new attack vectors to malicious parties that have access to the CloudCO system. For example, the “black box” steady state functionality of a virtualized system may be identical to a corresponding physical network function implementation, but the elasticity and dynamic behavior a virtualized system is capable of implies significantly different system responses to load will be possible, which can be exploited for malicious purposes if poorly designed or executed.

Privacy involves the need to ensure that information to, from and between customers can only be accessed by those who have the right to do so. Further, privacy requirements can vary by regulatory region. In general, two ways to ensure privacy is recognized:

- Preventing data, from being copied to a non-intended destination.
- Encrypting data, so that it cannot be understood even if it is intercepted.

This document does not define any specific mechanisms.

4 Introduction

This document describes the architectural framework of a CloudCO Domain. A CloudCO Domain makes use of a NFVI, in particular the use of general and its storage general purpose network switches, compute nodes and storage, as well as Management and Orchestration Functionality (MANO). The NFVI hosts VNFs, which implement in software many of the functions requiring special-purpose hardware in legacy networks. PNFs will be attached to the NFVI as well, to allow connecting access Input/Output (I/O) and network facing I/O, as well as to allow connecting any PNF necessary to deliver a service. Hence, a certain amount of SDN functionality will need to be leveraged and orchestrated in tandem with the NFVI to deliver an end-to-end service.

The CloudCO architecture has the following benefits:

1. It enables fine-grained control of system-level network and service design, in part due to the disaggregation of legacy network nodes into separate network functions.
2. It supports the flexibility and scalability offered by virtualization of network functions.
3. It supports automated, rapid deployment and onboarding of services.

This Technical Report describes the functions that are needed to achieve service delivery across the architecture, as well as how those blocks of functionality will need to interact, via its Interfaces.

Section 5 will elaborate on what exactly encompasses a CloudCO and describes how legacy network nodes like a Broadband Network Gateway (BNG), Access Node (AN) or Customer Premises Equipment (CPE) can be potentially disaggregated into various network functions, and will describe how some of these functions can be a candidate for virtualization.

Section 6 will then focus on the architectural framework of the entire system to allow for dynamic placement and service chaining between the various network functions as well as allowing external consumption of that functionality. It describes the supporting hardware and software architecture needed and its functionality as well as how it interacts with the rest of the service provider OSS/Business Support System (BSS) environment in order to build end-to-end services. It will then detail the actual architectural framework, i.e., the various functional blocks as well as the interfaces that connect them. It will describe what the function of every building block is, without describing how that functionality is achieved. It will also describe the interfaces between the building blocks which can be leveraged to create the necessary state across the system to build services. It will also describe the SBI, which is the CloudCO API. The CloudCO Northbound Application Programming Interface (CloudCO NB API) allows the CloudCO Operator to allow 3rd party consumption of CloudCO functionality without having to expose the internals of the CloudCO architecture. This 3rd party consumable API can also be used to on-board and create new services.

Section 7 describes how a CloudCO Domain could be bootstrapped, and also what capabilities the relevant interfaces need to have in order to fulfill the Use Cases as described by the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios.

Finally, Section 8 concludes and discusses the possible next steps.

5 Characteristics of a Cloud Central Office

This section describes a CloudCO from a contextual point of view. It also discusses the principles applied to the potential decomposition of legacy network nodes, and reviews decomposition and virtualization options for several types of nodes, considering the services (both legacy and new) that they should support.

5.1 The Cloud Central Office Context

The architecture defined for CloudCO is nearly completely redefined relative to TR-178 [5], although the two architectures can co-exist. In that sense, the supporting hardware will be deployed inside a Central Office location, next to the existing CO equipment. The supporting hardware includes at minimum:

- General purpose network switches interconnected over a switch fabric.
- General purpose compute hosts connected to the switches
- Access I/O hardware connected to the Switches. These can be an evolution of a legacy access node or a new design
- Network I/O hardware connected to the Switches, to connect the CloudCO Domain to the Service Provider (SP) backbone.

Although most of the CloudCO hardware maybe installed inside a CO, the CloudCO Domain can extend over multiple physical locations, even non-CO locations. It could be envisioned that the supporting hardware is installed across multiple locations to offer a single instance of a CloudCO Domain can be offered that spans multiple locations. Furthermore, it can be envisioned that network functions that are located traditionally at the customer premise will be disaggregated across multiple physical locations (customer premise and CO) and therefore logically the CloudCO Domain can extend to the customer premise as well.

5.2 Decomposition of Central Office Legacy Entities

The existing broadband network architecture as specified by TR-178 [5] defines node types, such as BNG, AN, and CPE, assigns nodal functions, and specifies the interfaces between these nodal functions. While the CloudCO architecture aims to offer new capabilities, scale and agility, as a first step it should realize the same network functionality, as the existing legacy network. The functional partitioning process of existing functionality to the new architecture involves two main topics (1) to disaggregate or not and (2) to virtualize or not.

Disaggregation is the process of separating existing nodal functionality into more modular and granular network functions. Through disaggregation some legacy functions may be eliminated or rendered unnecessary by the new architecture. An example could be the network control plane on Access Nodes that can be centralized into an SDN-like application, and the forwarding plane of a BNG can be distributed across various different nodes.

Virtualization is the softwarization of one or several of the nodal network functions that may be hosted on generic Commercial Off-The-Shelf (COTS) hardware (e.g., VNF in NFV terminology).

Each function in the CloudCO architecture is implemented either as software or as dedicated hardware and embedded software (e.g., PNF) via an abstraction layer towards an SDN control framework.

Components of the network itself may also be virtualized, leveraging a combination of Layer 2 (L2) and Layer 3 (L3) forwarding and filtering functionality within virtual switches inside the hypervisor, along with overlay based technologies, such as the ones defined by IETF Network Virtualization Overlays 3 (NVO3).

Another type of virtualization is network virtualization. Rather than implementing the network function as a workload inside a VM or Container, the network function is leveraging a combination of L2 and L3 forwarding on a PNF (like a L2/L3 switch) on the one hand and filtering functionality and overlay based technologies inside the virtual switch inside the hypervisor the other hand as defined by IETF NVO3. Virtual networks by using overlays are defined in RFC 7365 [21]. In summary, overlay networking leverages virtual networks to create L2 connectivity across hypervisors, and also leverages virtual distributed routers to establish the same default gateway in every hypervisor across the infrastructure. Virtual networks connect to distributed routers. This allows a Workload to be positioned on any compute node, as it will always find the same virtual network and the same default gateway.

To following two sections suggest considerations for disaggregation and virtualization decisions. Note that no specific disaggregation and virtualization technique is assumed. The virtualized software can be arranged as VNFs, VNF Components (VNFCs), VMs, Linux Containers, micro-services, etc. The text refers to the term ‘workload’ to generalize the above listed terms.

5.2.1 Disaggregation Considerations

The main considerations for disaggregation and its granularity are:

1. **Ability to virtualize -** Disaggregation is essential in order to separate functions that can benefit and are candidates to be virtualized from others that cannot, see Section 5.2.2.
2. **Ability to efficiently split the user plane and the control plane –** It is frequently beneficial to split control plane and user plane functions in order to virtualize and centralize the control plane to facilitate SDN control. As control plane functions are compute rather than forwarding intensive, they are clear candidates for virtualization. User plane functions, which may have high packet processing performance requirements, may be more frequently implemented in hardware that is abstracted towards an SDN control layer. Alternatively, some of the forwarding decisions (e.g., switching, routing, service chaining, etc.) could also be implemented with functionality in a virtual switch inside the compute host hypervisor and by leveraging (1) network virtualization techniques as defined by IETF NVO3, and (2) the forwarding capacity of the Top of Rack (ToR) switches as an efficient IP underlay. In addition, separation of the control plane and user plane must also consider the following:
    - The communication channel between the user plane and control plane components will be implemented via a network and will require similar Quality of Service (QoS) to a legacy backplane. Sufficient resilience needs to be built inside the control plane

application (e.g., workload clustering) and inside the NFVI for network link failures and/or compute host failures. Additionally, enough data capacity needs to be allocated to allow control plane updates between the control plane and user plane components.

3. **Reusability** – It may be advantageous to disaggregate components that can be reused in several use cases or at various physical or logical points in the network. For example, an Internet Group Management Protocol (IGMP) snooping or proxy function can be used in different parts of a multicast hierarchy.
4. **Upgrade cycle** – components that are disaggregated and implemented in software can be upgraded more easily than components integrated into monolithic functions or in hardware. Examples include implementations of unstable standards and functions with potential for high innovation.
5. **Interface simplification and standardization** – The set of interfaces and APIs of a component is derived from the functional decomposition. Disaggregation should minimize interface and API complexity. One guideline is to avoid exposing interfaces unnecessarily, especially if they are not standardized. Another guideline is to expose interfaces when it is beneficial from an operational perspective, e.g., for troubleshooting. Additionally, it is important to maintain, where possible, standardized interfaces and reference points.
6. **Performance** – In some cases disaggregation may need to be optimized for performance. It is beneficial to separate components that are resource sensitive. For example, decoupling Central Processing Unit (CPU) intensive components may allow scaling out those components independently or possibly running them on different Hardware (HW) platforms.
7. **Operational consideration** – A few operational considerations need to be taken into account:
   - **Orchestration** – in general the complexity of orchestration is likely to grow if the orchestrated service requires a larger number of network functions.
   - **Availability –** Disaggregation may improve availability as the effect of upgrade or failure of a function is minimized with smaller components because strategies like reboot of a function instances, redundancy and hot swap are more time and resource effective with “lighter” functions.
   - **Diagnostics –** With more granular functions and more exposed interfaces, increases the multiplicity of components thus increasing the diagnostics challenge, and as well it also has the potential to pin point a root cause.

In view of these considerations, one of the key choices to make on disaggregation is granularity. A trade off should be made between the complexity of finer granularity and the benefits of increased modularity. When evaluating this trade off in the context of an actual virtualization environment other factors may need to be considered. For example, higher numbers of VNFs through a forwarding path may add latency because of the increased number of hops or switching actions for packets traversing the path.

5.2.2 Virtualization Considerations

Like disaggregation, virtualization choices should be made considering the benefits vs. the complexity, risk, and performance impact.

Some key characteristics of functions that would most benefit from virtualization are:

- **High rate of change** – Functions that are expected to be frequently modified, for example due to standards evolution, or ever changing requirements such as security threats. The decoupling of those functions from embedded systems allows for an independent and easier lifecycle and upgrade process.
- **Varying scale** – Functions that may frequently change in their resource usage over time, such as a control protocol implementation that needs to process a growing number of transactions over time. Such functions can benefit from cloud scale out techniques.
- **Differentiation** – Functions that may be subject to differentiation between vendors, e.g., in features, scale, or innovation.
- **Specialization –** Functions that require specialized domain knowledge.
- **Interoperability –** Functions that can alleviate interoperability by having a centralized common implementation to reduce variability across multiple vendors.
- **Component reuse –** Functions that are virtualized may be instantiated in different configurations to serve specific needs, e.g., implemented in different network slices. The component stays the same and testing efforts may be reduced.

The characteristics of functions that make them good candidates to remain PNFs are:

- **Performance intensive** – One of the main challenges with virtualization remains user plane performance. There is still a significant gap between Application Specific Integrated Circuits (ASICs)/ network processors and COTS hardware in the achievable throughput and the associated cost and power consumption. At some points in the network, where there is high throughput and standard and stable packet processing and traffic management operations, it is difficult to justify virtualization of the entire user plane. However, overlay networking can combine hardware-based forwarding on COTS switches and more intelligent functions like routing and overlay encapsulations inside COTS hosts, clearly lessening the dependencies on ASICs/network processors.
- **Real Time Sensitive** – Another type of function that may be affected by virtualization is real-time functions, i.e., functions that must process and react to certain events within very strict time limits. Examples of such functions relevant to the CloudCO are protection switching, time-stamping, Dynamic Bandwidth Allocation (DBA) allocation. It is often difficult or impractical to virtualize such functions, due to either the latency involved in sending the event to a VNF, or the actual latency of software processing or both.
- **Deployed Equipment** – Some functions may already be well supported by deployed network equipment, and service providers may want to continue realizing their investment in that equipment for years to come. This does not necessarily mean that they cannot upgrade their network to a CloudCO architecture, but rather that selected functions could be implemented as existing equipment PNFs.

The tradeoffs between VNF and PNF implementations will shift over time as both virtualization technology and the performance of the general-purpose hosts making up the NFVI improve. Therefore, the CloudCO architecture should be flexible enough to support repartitioning and replacement of functions without affecting service performance to customers.

Figure 1 illustrates an example of disaggregation and virtualization. In the existing network, function F is deployed as an embedded function on a network node. F comprises four embedded

sub-functions, Fa, Fb, Fc and Fd, with some internal interfaces between them. These interfaces are implementation specific and not standardized. Fd is partially implemented by HW. An example of the disaggregation and virtualization decisions is described below:

a) Fa and Fb have a complex interface between them. They are interdependent and unlikely to be deployed separately. They have similar resource needs and both are HW independent. They are therefore candidates for virtualization and there is no need to create separate VNFs for them, so a single vFab function fulfills their combined functionality in the new architecture.
b) Fc is a self-contained function that can be reused in other network locations, making it a good candidate for virtualization as a separate VNF.
c) vFab and vFc can be provided by different vendors, and they may have independent upgrades and scale out cycles.
d) Fd is a user plane function that requires high packet processing performance and has a stable standard feature set. Therefore, it makes sense to leave it as a PNF. The PNF may be realized on a subset of an existing system or on a new one.
e) In the new architecture, the interfaces between vFab and vFc and between vFab and Fd PNF are APIs that are standardized or published.

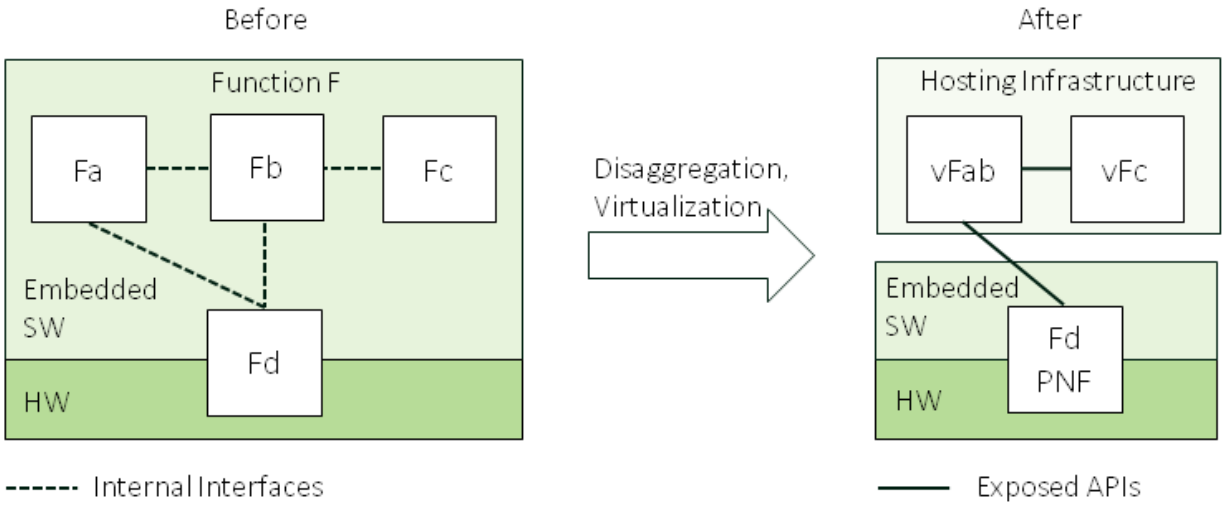

Figure 1: Example of disaggregation and virtualization

Alternatively, point f) can result in another disaggregation and virtualization decisions; see Figure 2:

f) Fd is a User plane function that requires high packet processing performance that can be implemented by dividing the function into two sub-functions:
    1. An underlay forwarding function Fd-UL that is implemented as a PNF and that limits itself to L2 and IP forwarding between various infrastructure HW elements
    2. An overlay forwarding function vFd-OL that is part of the hosting infrastructure which does L2 and L3 lookups for the user plane, and subsequently encapsulates the packet ready to be forwarded by the underlay forwarding function.

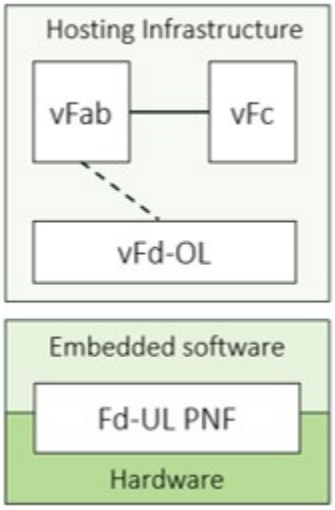

Figure 2: Alternative example of disaggregation and virtualization

5.2.3 PNF Requirements

The CloudCO solution is likely to include both VNFs and PNFs. It is expected that different applications may require a different combination of the two NF types. It is also expected that over time, with improvement in virtualization technology more functions will be virtualized. Therefore, there is a need for PNFs to be as modular and independent as possible to allow such scenarios.

5.2.4 Multi Tenancy Considerations

A tenant is an entity which uses or controls resources in a network owned or operated by a different entity. In the case of broadband access services, a tenant can have different meanings, depending on context. A tenant can be:

- A service provider's customer, either enterprise (a company) or residential (a subscriber). In this case the tenant is the retail customer.
- A network Operator's customer, in case of a wholesale or resale relationship. In this case the tenant is the wholesale customer of the network Operator and the service provider to the retail customer.

Multi-tenancy refers to sharing multiple tenants over a single CloudCO Domain. As a system, the CloudCO is multi-tenant: the same CloudCO Domain is able to serve multiple customers and partners.
While multi-tenancy is not a new concept (web servers, voice servers, BNGs, etc. are multi-tenant), computer and network virtualization techniques provide new ways to achieve multi-tenancy. For example, assume that the MSBN must support wholesale service to 10 Internet Service Providers (ISPs), each with their own subscriber databases, address pools and routing domains. The traditional approach would be to have one BNG instance running 10 Virtual Private Network (VPN) contexts, each with their own Authentication, Authorization, Accounting (AAA) client, address pool and

routing table. With NFV and SDN, an alternative approach is to have one virtual BNG instance per ISP and to leverage network overlays or flows to separate ISPs.

Generalizing this example, multi-tenancy can be achieved by having a single network function supporting multi-tenancy or by having multiple single-tenant network functions. Similar to the disaggregation consideration in Section 5.2.1 there are advantages in each approach. Therefore, the CloudCO needs to simultaneously support single tenant and multi-tenant network functions.

As an example, a Network Enhanced Residential Gateway (NERG) is required to offer Dynamic Host Configuration Protocol (DHCP)-based address management per home, where each home may have overlapping private subnets. An implementation can use a single multi-tenant DHCP server (for example, by adding a home User Interface Design [UID] to the address tables) or to have a contained DHCP server per home (for example, a container for each subscriber, containing a DHCP application).

A service can leverage a combination of multi-tenant and single-tenant network functions. For example, a NERG could combine DHCP-based containers per subscriber and a multi-tenant Network Address Port Translation (NAPT) network function.

Different multi-tenancy models impact different requirements of the CloudCO, in terms of:

- User plane.
- Control plane.
- Management plane.

	Dedicated NF	Multi-tenant NF
User plane	Requires a dedicated forwarding path between the tenant (e.g., RG) and the NF	Same path for multiple tenants
Control plane	Must build the forwarding paths dynamically (e.g., flows, overlays)	Integration between NF and control plane to identify the tenant, for example based on source IP address or subscriber identifier
Management plane	1:1 relationship between NF / management interface / tenant	Per tenant management requires additional abstraction layer between tenant and management interface

Table 1: Multi-tenancy models and requirements

Furthermore multi-tenancy is extended to the service level. Individual single-tenant and/or multi-tenant network functions can be combined into a network service. These network services need to be multi-tenant capable as well, and should provide the necessary isolation between the tenants packet flows, as well as allow every tenant to change the characteristics of its own network service, while not impacting other tenants behavior. Furthermore the level of abstraction that is visible to the service consumer should be such that the underlying network functions are not necessarily exposed.

Network services can be combined/wrapped with other services into another higher layer service, again with the same multi-tenancy and abstraction requirements.

5.2.5 Decomposition of the BNG

As defined in TR101 [2], the BNG is defined an Ethernet-centric IP edge router, and the aggregation point for the user traffic. In TR-178 [5], the MS-BNG is introduced, which can offer services to both residential and business customers as well as allow mobile backhaul deployments. An MS_BNG performs Ethernet aggregation and packet forwarding via IP/MPLS, and supports user management, access protocols termination, QoS and policy management, etc.

5.2.5.1 Deconstruction & Inventory of Functions

The BNG handles both subscriber management services and routing functions. The decomposition of the legacy BNG is shown in Figure 3, which mainly includes the BNG Service / Subscriber Session Control, BNG Service Forwarding Plane, Routing Control, and Routing Forwarding Plane. Figure 3 shows the functional components inside the BNG and the interactions with external entities

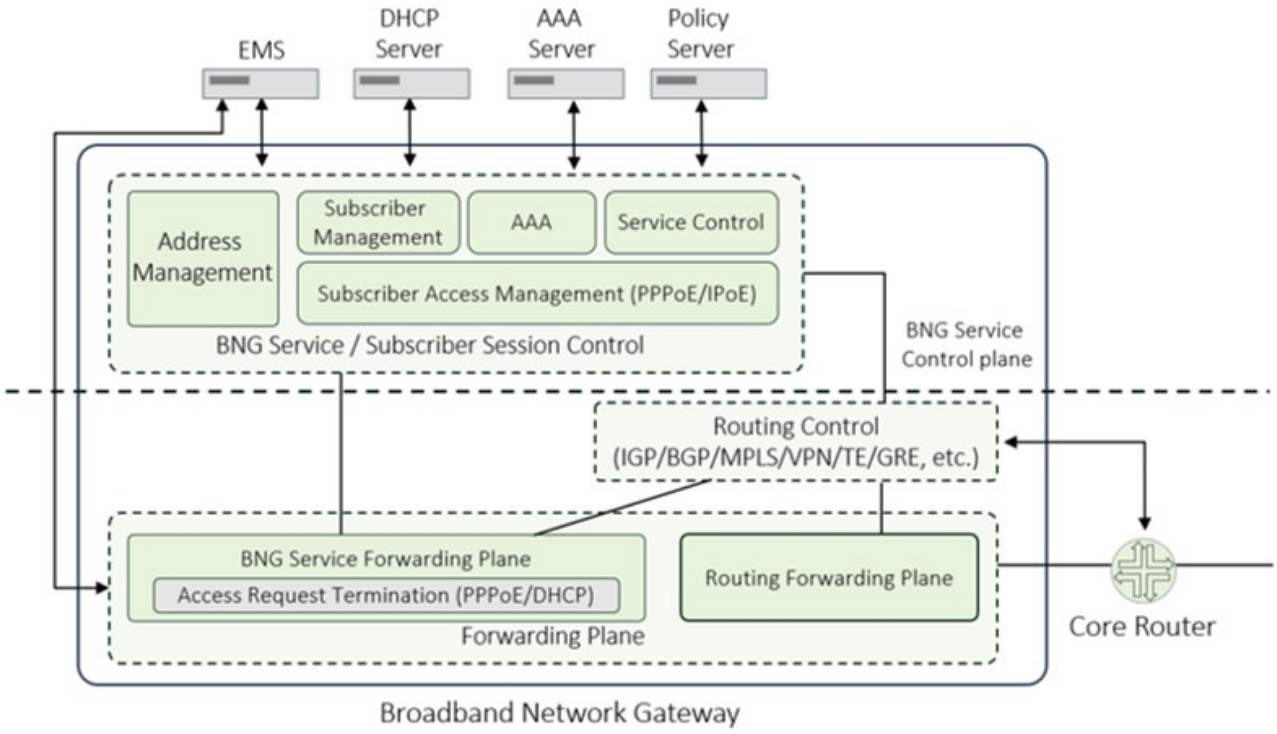

Figure 3: Functional components inside BNG and the interactions with external entities

The BNG Service/Subscriber Session Control includes Address Allocation & Management, Subscriber Access Management, AAA, Subscriber Management, and Service Control:

- **Address Allocation & Management:**

By interacting with the external DHCP server or using local address pool, the Address Allocation & Management component is responsible for allocating IP addresses.

- **Subscriber Access Management:**
  The Subscriber Access Management component is used to terminate the Access-Requests of the subscribers, process the access protocol messages (e.g., Point to Point Protocol over Ethernet (PPPoE), Internet Protocol over Ethernet (IPoE), and DHCP), and obtain the subscribers' credentials for further authentication.
- **Subscriber Management:**
  The Subscriber Management component is responsible for managing subscriber's access to the system and forwarding policy.
- **AAA:**
  This component performs Authentication, Authorization and Accounting, together with Radius, DIAMETER. The BNG communicates with the AAA server to check whether the subscriber who sent an Access-Request has network access authority. Once the subscriber goes online, this component together with the Service Control component implement accounting, data capacity limitation, and QoS enforcement policies.
- **Service Control:**
  By interacting with the Policy Server, this component implements policy enforcement function such as data capacity limitation and enforces QoS policies on services.

Routing Control is responsible for constructing the routing forwarding plane of the BNG for data traffic. It interacts with central router using routing protocols such as Interior Gateway Protocol (IGP)/Boarder Gateway Protocol (BGP)/ Multi-Protocol Label Switching (MPLS) (e.g., Label Distribution Protocol [LDP]).

BNG Service Forwarding Plane forwards the subscriber access control related messages (e.g., PPPoE, IPoE, DHCP) and also performs user plane policy enforcement functions.

Routing Forwarding Plane delivers the data traffic traversing the BNG as constructed by the Routing Control component.

5.2.5.2 Identifying Opportunities for Consolidation & Elimination of Duplication

The above-mentioned decomposed functional components could be consolidated into fewer components in order to ease implementation and management. Such consolidation is dependent on design and performance driven by scaling of the various components:

- **Subscriber & Service Management**:
  The Subscriber Access Management, AAA, Subscriber Management, and Service Control functional components in the BNG Service Control could be further consolidated into a single VNF, named Subscriber & Service Management.
- **Forwarding Plane consolidation**:
  BNG Service Forwarding Plane could be consolidated together with the Routing Forwarding Plane forming a single forwarding plane. Such consolidation is dependent on design and performance driven by scaling of the various components.

5.2.5.3 Deployment Options of a Disaggregated BNG

The disaggregated BNG is shown in Figure 4. The BNG Service Control Plane could be virtualized and centralized, which provides significant benefits such as centralized session management, flexible address allocation, high scalability for subscriber management capacity, and cost-efficient redundancy, etc. The functional components inside the BNG Service Control Plane can be implemented as VNFs and hosted in a NFVI. The User Plane Management module in the BNG control plane centrally manages the distributed BNG user plane (e.g., load balancing), as well as the setup, deletion, and maintenance of interfaces between the control plane and the user plane. The routing control and forwarding plane, i.e., the BNG user plane (local), could be distributed across the infrastructure, shown as Option 1 in Figure 4. They could be left in a physical device as PNFs due to high packet processing performance requirements, or they could use a mix of VNFs and PNFs. The routing control could also be centralized, while the forwarding states are kept in the distributed forwarding planes, shown as option 2 in Figure 4. The centralized routing control for controlling many distributed forwarding planes across a number of CloudCOs adds complexity such that initial deployments are expected to use distributed routing control.

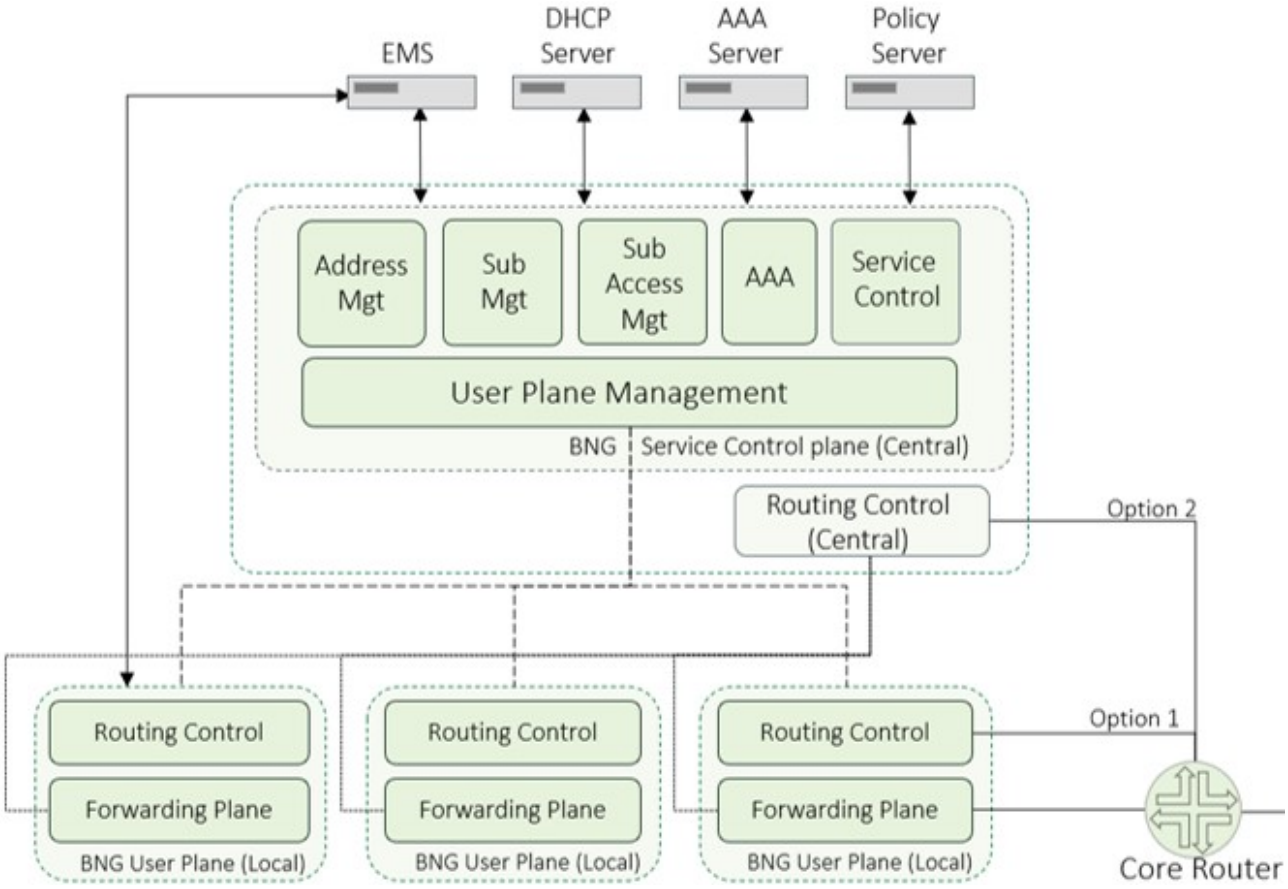

Figure 4: Control plane and user plane separation of the disaggregated BNG

Another deployment option is to virtualize both control plane and a user plane of the BNG in the CloudCO. The previous option requires a switch fabric overlay and must guarantee a certain QoS for control traffic and inter user plane traffic (same requirement if the user plane is a PNF while the

control plane is hosted separately as a VNF). In addition, resource scaling is "coupled" to network overlay planning. For example, scaling up either control plane or user plane requires scaling the overlay network as a prerequisite.

To decouple this "overlay" requirement for virtualization and scaling, as alluded to earlier in section 5.2.1, service providers can host the BNG application on a single VM server. If more (subscriber or throughput) scale is required, there are two ways to scale up. Within a VM, additional control and user plane compute and memory resource can be increased. And if the resources become limited within a VM server, a new VM can be spawned up hosting new BNG(s). Each BNG application is independent, containing its individual control plane, eliminating the need of a network underlay completely for control plane traffic or inter user plane traffic.

To manage all VNF BNGs, a centralized OSS/BSS Management system can be used. The OSS and BSS will allocate addresses, provision the subscriber management service and etc. The diagram shown in Figure 5 illustrates this architecture.

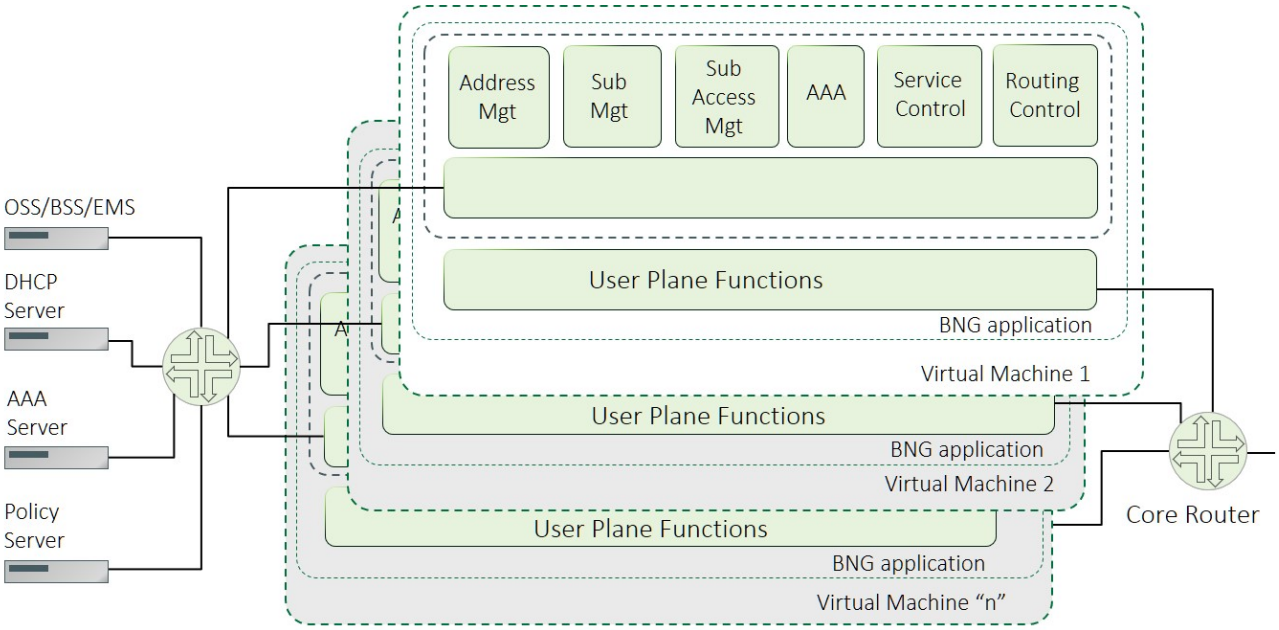

Figure 5: Functional components inside BNG for scenario that both control plane and user plane are virtualized

5.2.6 Decomposition of the Access Node

TR-178 [5] defines three variations of the AN: the Ethernet Access Node, which corresponds to the AN as defined in TR-101 [2]; the MPLS-enabled Access Node, which adds MPLS features to the Ethernet Access Node; and the BNG-embedded Access Node, which adds specific MS-BNG features to the MPLS-enabled Access Node. All three variations base their requirements on the Ethernet Access Node, with the MPLS-enabled Access Node and the BNG-embedded Access Node adding successive layers of requirements on top. The decomposition of the AN in this section is limited to the functions on the subscriber side of the Vc reference point, and therefore based on the AN requirements for the Ethernet Access Node.

Note that when the AN is a Passive Optical Network (PON)-based AN as defined in TR-156 [3], its functions are distributed between the Optical Line Termination (OLT) and the Optical Network Units (ONUs) on the Optical Distribution Network (ODN). The decomposition shown here also considers that distribution. Functions that reside in the ONUs are not affected.

5.2.6.1 Deconstruction & Inventory of Functions

The decomposition of the legacy AN is shown in Figure 6.

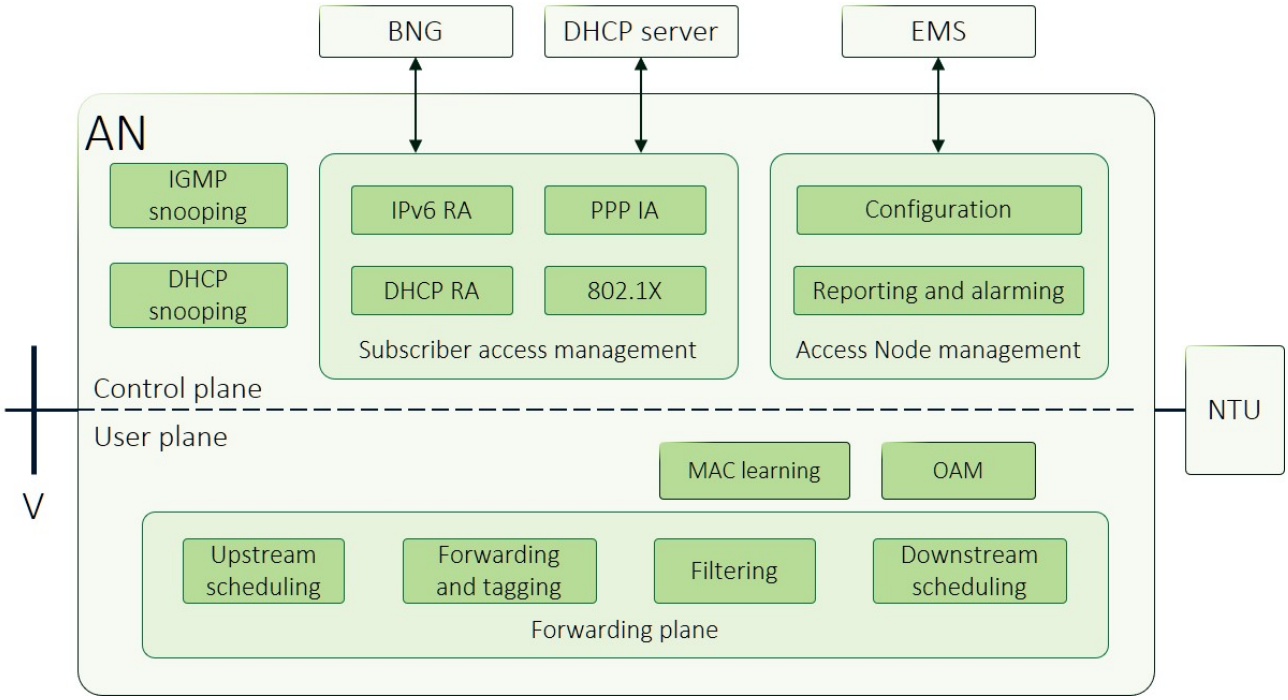

Figure 6: Functional components in an Access Node

The forwarding plane functions include:

- **Forwarding and (packet processing) functions.** Forwarding may be based on Service – Tag (S-Tag) including P-bits, Customer-Tag (C- Tag) including P-bits, Media Access Control (MAC) Address, MPLS/ Segment Routing (SR) Labels.
- **Traffic management.** Upstream classification and scheduling of traffic forwarded to the interface at the V reference point. Downstream scheduling at each access port.
- **Filtering,** including MAC address filtering and anti-spoofing.

Additional functions which are coupled to the user plane include MAC learning and Ethernet layer Operation, Administration and Maintenance (OAM).
In the control and management planes, functions can be grouped into node management and subscriber access management.

- **Node Management** includes the configuration, reporting and alarming functions associated with management of the AN by an Element Management System (EMS).

- **Subscriber access management** includes the Relay Agent and Intermediate Agent functions used for subscriber access by DHCP, IPv6 and Point to Point Protocol (PPP), as well as 802.1X.
- **IGMP snooping and DHCP snooping** observe frames in the user plane for their respective protocols, and use the observed data to control filtering and/or forwarding behavior.

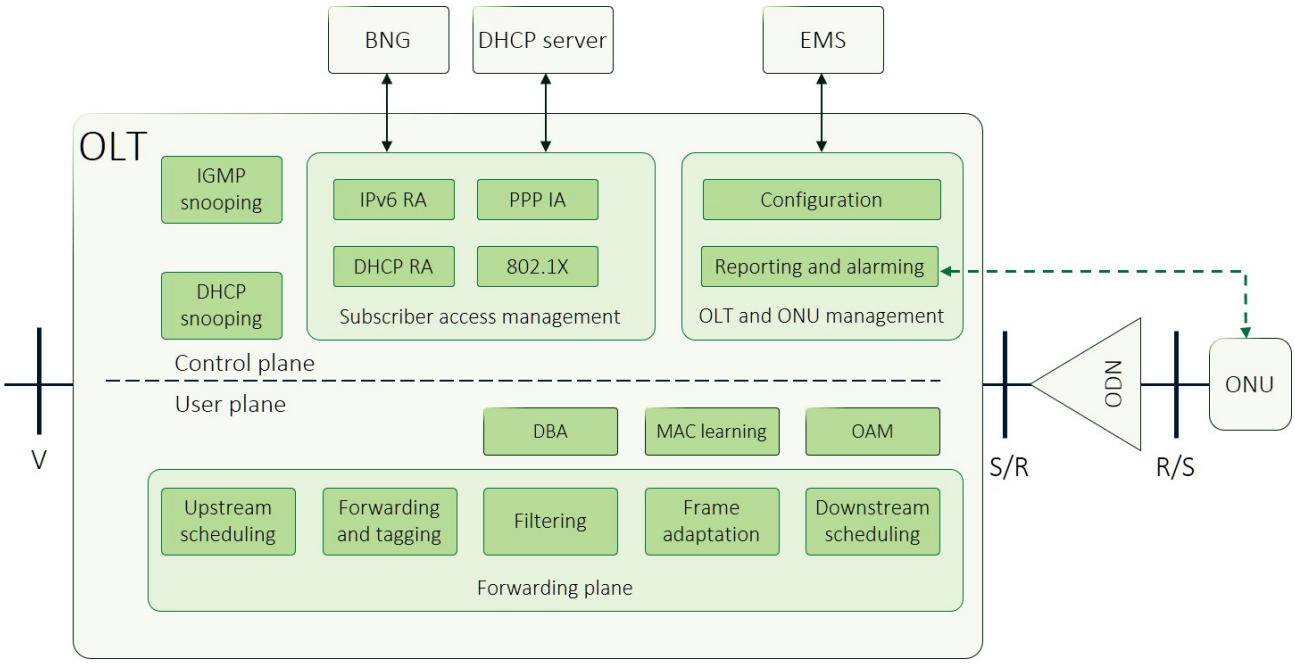

Figure 7: Functional components in OLT

In PON-based ANs, some of the functions described above may be distributed between the OLT and the ONUs. PON-based ANs may also include additional functions as shown in Figure 7. These functions include:

- **Forwarding and tagging functions:** Tag manipulation in the OLT involves S-Tags and/or MPLS labels. C-Tag manipulation typically takes place in the ONUs.
- **Frame adaptation:** Mapping and encapsulation of downstream traffic into Generic Encapsulation Method (GEM) ports or Logical Link Identifiers (LLIDs) based on Virtual Local Access Network (VLAN) tags, traffic class and/or MPLS labels. Decapsulation of upstream traffic.
- **Dynamic Bandwidth Allocation (DBA):** Allocation of upstream grants on the PON.
- **OLT and ONU Management**: In legacy PONs, the OLT and ONUs are managed as a single entity, with the OLT providing the interface to network management and sending control messages to/from ONUs over multiple protocols, including Physical Layer OAM (PLOAM) or Multipoint Control Protocol (for Physical [PHY] layer and management channel control), and ONU Management and Control Interface (OMCI) [22] or Ethernet OAM (eOAM) (for higher layer control).

5.2.6.2 Potential Virtualization Options

Most, but not all, of the control plane functions shown in Figure 6 and Figure 7 are candidates for virtualization:

- The Relay Agent, Intermediate Agent, and 802.1X functions associated with subscriber access management can be externalized to a controller or other virtualized application, as long as the associated traffic is redirected to the controller.
- The control plane functions associated with IGMP snooping and DHCP snooping can both be virtualized, provided that the associated traffic is exposed to the virtualized functions.
- The higher level functions associated with configuration, reporting and alarming for the AN can be virtualized, provided that the appropriate interfaces are included in low-level functions on the AN hardware. These interfaces may be hardware specific.
- When the AN is PON-based:
    - Management of the ONUs above the PHY layer including configuration, reporting and alarming functions can be virtualized. These functions correspond roughly to the management functions communicated over OMCI [22] or eOAM in legacy OLTs.
    - The higher level functions associated with configuration, reporting and alarming for the OLT can be virtualized provided that the appropriate interfaces are included in low-level functions on the OLT hardware. These interfaces may be hardware specific. Note that since configuration of the OLT and ONUs must be consistent, the corresponding functions must be coordinated whether or not they are virtualized.
- When the AN is a Digital Subscriber Line Access Multiplexer (DSLAM), or an Fiber To The distribution point (FTTdp) Distribution Point Unit (DPU) supporting Digital Subscriber Line (DSL) or G.fast:
    - Management of control plane functions including Dynamic Time Assignment, On-Line Reconfiguration (OLR), management of low power link states and cross-layer low-power mode control, and DSL Quality Management TR-197 [6]. These functions correspond roughly to the management functions that can be implemented by communicating over a DPU or DSLAM Northbound Interface, such as being proposed in TR-355 [10].
    - Inputting DSLAM or DPU state information; and correlating statistics across a cable or vector group that has faults, crosstalk, or noise which are correlated across multiple lines.
    - Configuration of vectoring groups / strategy.
- The remaining user plane and forwarding plane functions are expected to remain in the AN: The local nature of MAC learning makes it better suited to remain in the AN.
- Ethernet layer user plane OAM functions are inseparable from the forwarding path that they instrument and must remain in the AN. Other Ethernet layer OAM functions may be candidates for virtualization but they will need to be considered on a case by case basis.
- The strict requirements on latency associated with PHY level control of the ONUs make virtualization of these functions more difficult.
- Vector pre-coding of transmit signals in real-time in DSLAMs and DPUs would require a large amount of data capacity to virtualize and therefore this is unlikely to occur.
- The timing of real-time line updates such as OLR and Dynamic Timing Allocation (DTA) in DSLAMs and DPUs generally precludes their virtualization.

5.2.7 Decomposition of CPE

For wireline TR-178 [5] architectures, the CPE is often referred to as the Residential Gateway (for Residential Services) or Business Gateway (for L2/L3 Business Services). In some deployments, decomposition of individual network functions might be beneficial. In the latter case, it would be more efficient to fully decompose these devices into multiple functions; however, there is a desire to virtualize certain functions or set of functions. For Residential Services, virtualization and disaggregation options are described in TR-317 [8]. For Business Services, corresponding options are described in TR-328 [12]. The Virtualized functions can either stay at the customer premise (for TR-328 style services), which means that logically the NFVI portion of the CloudCO needs to be extended into the customer premise, or there might be a desire to move the virtualized functions into the CO in order to avoid moving the NFVI into the customer premise (for both TR-317 and TR-328 style services).

A CloudCO Domain needs to support all of the virtualization and disaggregation options offered in these Technical Reports (TR-317 [8] and TR-328 [12]).

5.3 Hybrid Architectures

CloudCO's flexibility in (1) allocating virtual functions on the NFVI, and (2) accommodating and supporting access technologies drivers makes it a very versatile platform to implement hybrid fixed and mobile networks and to facilitate convergence and/or seamless integration paths.

The Access Node and the BNG are fundamental elements of the CloudCO concept; likewise their counterparts in mobile networks are also expected to be deployed per the CloudCO's paradigm, either as legacy PNFs or as disaggregated and/or functionally virtualized implementations. For example, the Serving Gateway (S-GW)/ Packet Data Network Gateway (P-GW), as counterpart of the BNG, is a natural choice, providing IP connectivity in geographically dispersed locations. In the case of the evolved Node B (eNodeB), it may be expected to be deployed in locations even further out towards the customer. In a scenario where the eNodeB is supporting functionally disaggregation, e.g., in CloudRAN cases, some selected functionalities such as the baseband processing may be hosted in the CloudCO infrastructure, taking advantage of its widespread geographical presence over network sites. On the other hand, AAA and Policy Control functions such as Broadband Policy Control Function (BPCF) and Policy and Charging Rules Function (PCRF), and as well as certain user plane functions such as, e.g., also the Mobility Management Entity (MME) are expected to reside in more central CloudCO locations.

Some reference points which are in a box-centric model, defined as external reference points, may become CloudCO internal ones, implying that they might not be at all exposed to functions outside the CloudCO.

6 Architectural Framework

6.1 Introduction

A CloudCO specializes in management, control and orchestration of access and edge functionality. Figure 8 illustrates the boundaries of a CloudCO Domain and how it fits into an Operator's larger overall network.

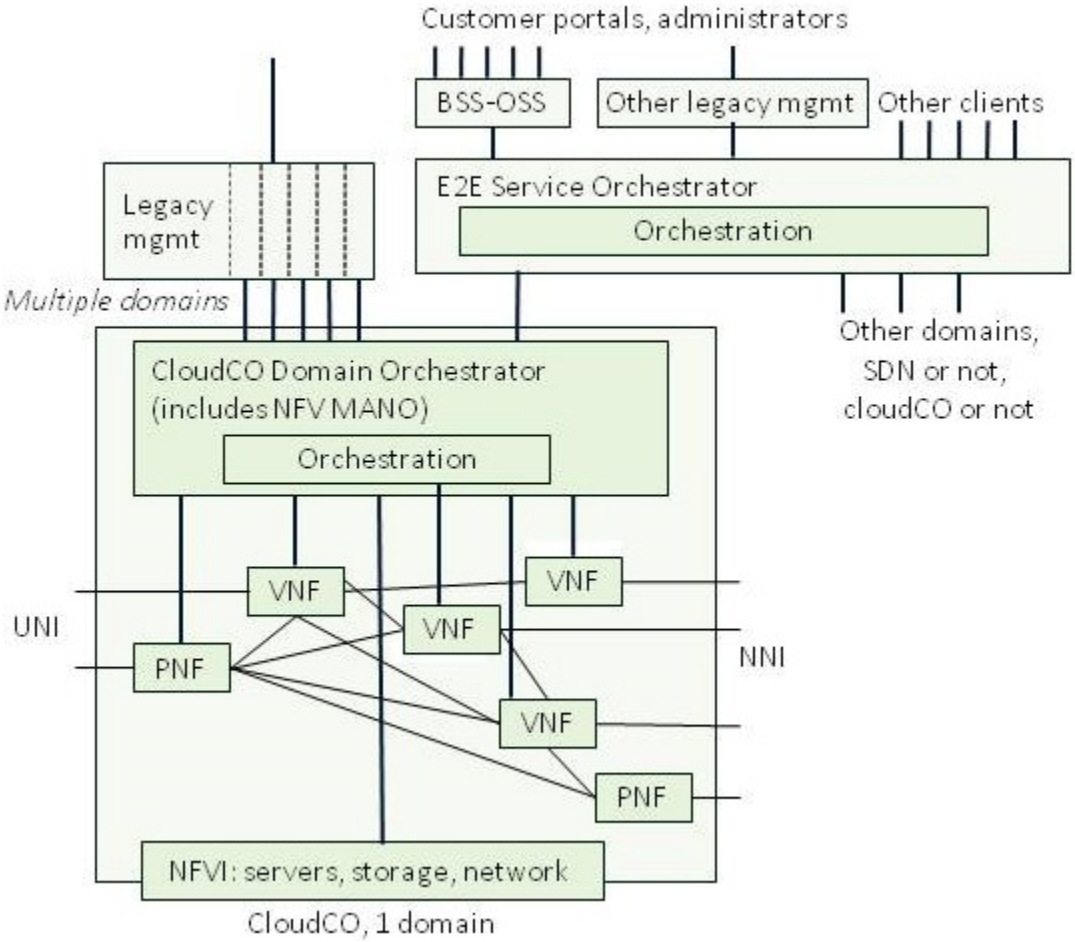

Figure 8: CloudCO in context of end to end service orchestration

It is important to recognize that customers, administrators and other possible third parties do not normally access a CloudCO Domain directly, but only through a superior entity, the End to End (E2E) Service Orchestrator, which has broader scope and visibility, and is in a position to resolve multiple client accesses. This broader scope spans the given CloudCO Domain, other possible CloudCO Domains within the same provider's network, non-SDN domains, and peer domains with partner carriers or providers. It is the responsibility of the End to End Service Orchestrator to coordinate all client interfaces with regard to their privileges and views, and to resolve any contentions that may arise.

Each CloudCO Domain is managed, controlled and orchestrated by a CloudCO Domain Orchestrator. Its NBI exists in the fully trusted provider domain. Essentially the NB API allows

onboarding of a service that is using PNFs, VNFs, Access I/O, Network I/O, and who are interconnected by a given Service Graph, and hides the complexity from the consumer of the API. Furthermore, in a CloudCO type of architecture the current specialization of access and edge network nodes becomes blurry in that the widespread distribution of NFVI throughout network sites allows flexible and dynamic deployment of edge functionalities flexibly in space and dynamically in time. This in turn translates into the ability to build up a service offering which is richer in content and easier to consume.

A CloudCO Domain could span multiple sites; potentially the whole access/edge network and a single instance of CloudCO Domain Orchestrator could orchestrate, manage and control the entire network.

Notwithstanding the above, there might be constraints/needs that warrants deploying multiple, unique CloudCO Domains, where the respective CloudCO Domain Orchestrators are orchestrated using an End-to-End Orchestrator, as shown in Figure 8.

To mention some of these constraints:

- Network-wide scalability of CloudCO Domain Orchestrators and CloudCO Domains over large networks.
- Segregated domains of CloudCO Domain Orchestrators that may serve the (interim) need for separation of segment/regional competencies and/ or other smooth migration purposes.

In these cases, the deployment of multiple CloudCO Domain Orchestrators would be needed. Each CloudCO Domain Orchestrator would supervise one CloudCO Domain spanning over a portion of the whole network. By allowing the NB API to span different locations, the hierarchy between the E2E Service Orchestrator and CloudCO Domain Orchestrators can be significantly reduced.

Figure 9 shows an example of multi CloudCO Domain Orchestrator deployment, recalling also the multi-site nature, intrinsic to a CloudCO Domain.

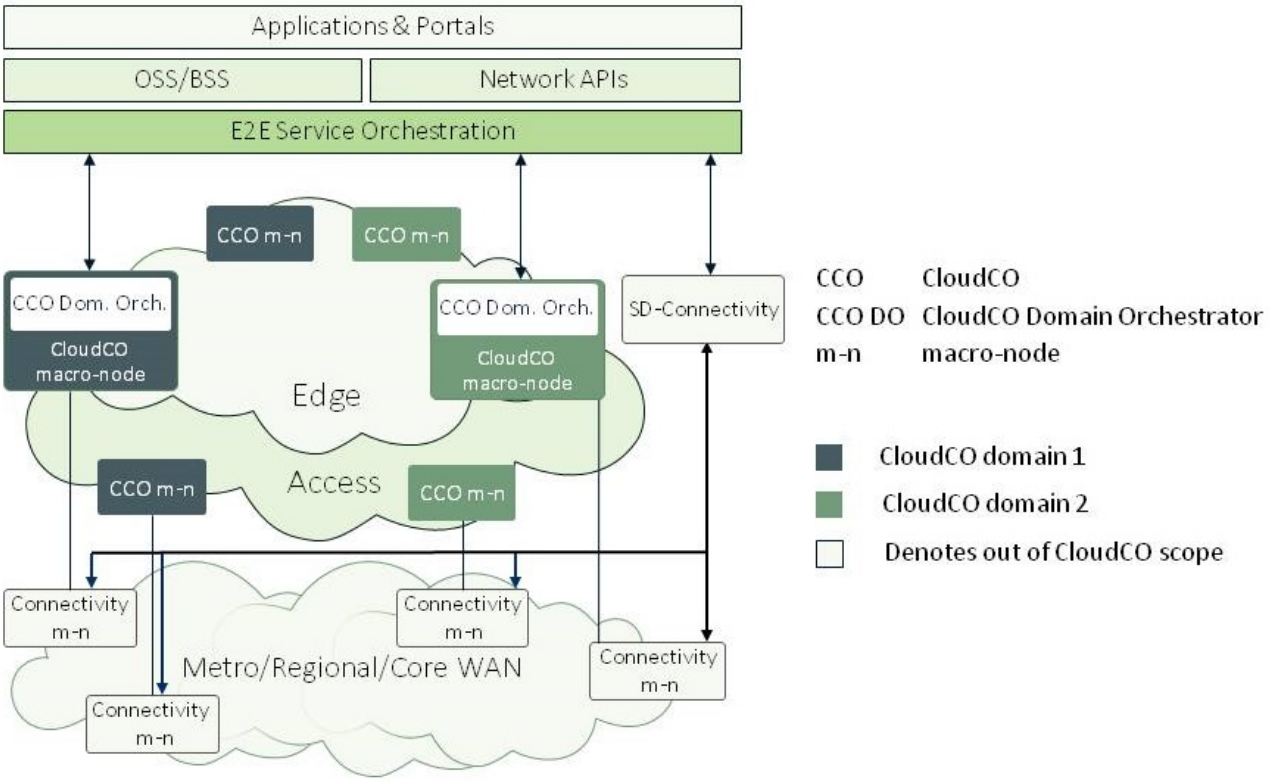

Figure 9: Example of deployment of multiple CloudCO Domain Orchestrators

Figure 9 synthesizes a number of the concepts described above and provides additional ones.

The objective of CloudCO is to enable a single Telco Cloud that consists of both Access and Edge functionality which is exposed seamlessly through open APIs to Service Layers.

CloudCO macro-nodes are assumed to support the same functionality (regardless of size, overall infrastructure, etc.) in that they implement Access as well as Edge functionalities enabled by the NFVI. Note that NFVI may be potentially deployed at each site. This enables lean and repeatable industrial processes for “Network Production” both on the suppliers’ side and on the Operators side.

The E2E Service Orchestrator coordinates multiple CloudCO Domain Orchestrators (CCO DO), deployed on specific nodes; the domain of each CloudCO Domain Orchestrator spans over portions of the whole access/edge network;

The end-to-end service chain requires geographic connectivity at metro, regional and core levels; The Wide Area Network (WAN) is depicted for completeness at the bottom of Figure 9 along with a Wide Area SDN Controller and Orchestrator, though all these elements are colored grey, since they do not belong to the CloudCO Domain.

This deployment example implies that the End to End (E2E) Service Orchestrator is able to coordinate multiple CloudCO Domain Orchestrators and the Wide Area SDN Controller and Orchestrator.

If the E2E Service Orchestrator is only expected to consume services via a network abstraction interface and not to implement network-related capabilities, some degree of cross-coordination among the underlying Orchestrators is needed.

An alternative option for a realization of the E2E Service Orchestrator with considerable network capabilities is by enabling cross-coordination among Orchestrators to be performed by a higher tier of orchestration. This higher layer of orchestration would act as a:

- Proxy for the service requests down to the appropriate network portion via the relevant Orchestrator (i.e., CloudCO and WAN).
- Reconciliation point for status, fault, PM inputs to coordinate service assurance action/reaction strategies.

Note that the two bullets above are not exhaustive and serve the scope to identify two orchestration flows E2E Service Orchestrator to higher network orchestration tier and lower to higher network orchestration tier.

Finally, it shall be assumed that the Network Orchestrators in Figure 9 perform a more sophisticated role in the sense that, as required and appropriate for their domain/hierarchic level, they have to support resource/status Orchestration, control of aggregate-traffic/granular-flow and ultimately support management.

Part of the E2E Service Orchestrator function is to coordinate user plane handoff when services cross domain boundaries, as most services will in fact do.

Because of the focus of CloudCO on the network access and edge, it is expected that the CloudCO Domain terminate at subscriber User to Network Interface (UNI) on the access side and at the Network to Network Interface (NNI) on the network side.

6.1.1 CloudCO Macro-node Description

While Section 6.1and Figure 8 illustrate a CloudCO in its wider context, this section and Figure 10 provide a first level view of the detail inside a CloudCO macro-node, i.e., how the CloudCO Domain is physically implemented at each of the multiple sites it spans across. The leaf-spine switch fabric is exemplary, not normative.

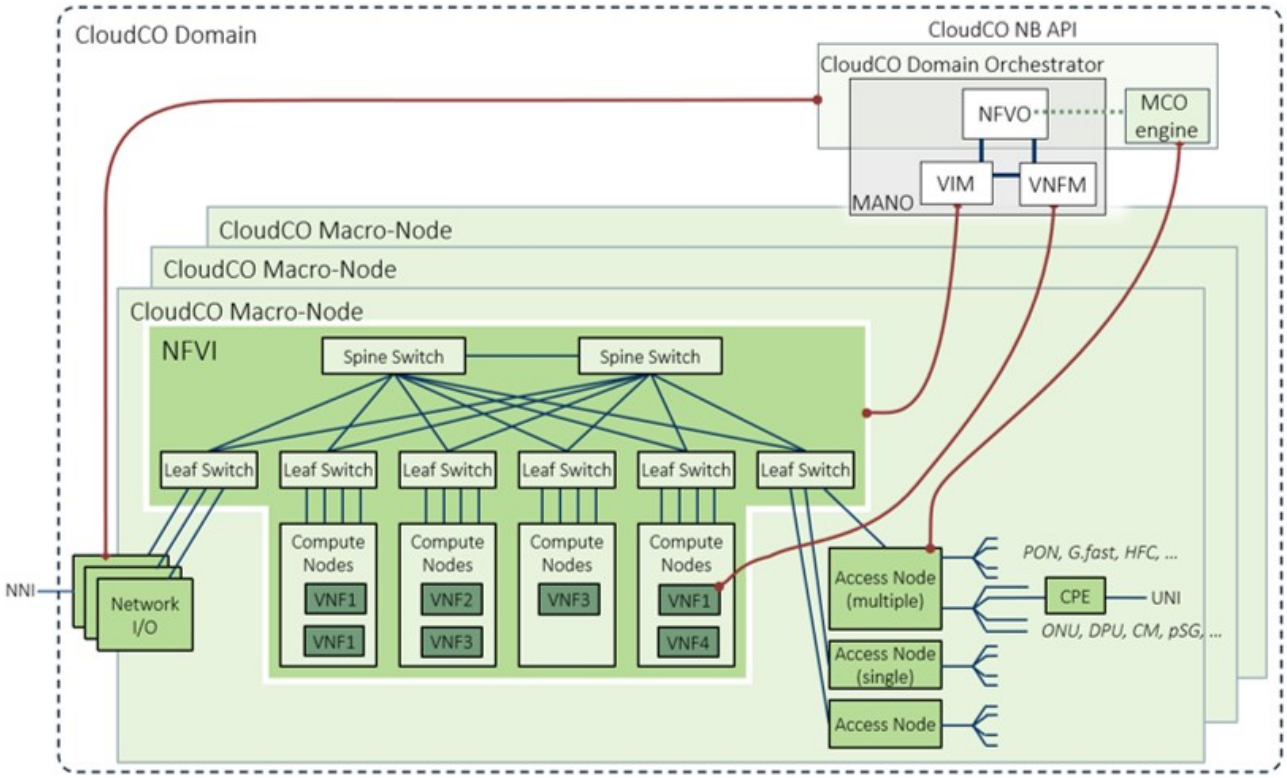

Figure 10: CloudCO macro-node structure

Major blocks in the CloudCO include the CloudCO Domain Orchestrator with built in NFV Orchestrator (NFVO), an NFVI, Network I/O (specialized nodes to interface the core network), Access Nodes (specialized nodes to interface the access network), and whatever physical CPE may be managed and controlled by the provider. Additional physical nodes may also be present. The interface between Access Node and CPE may use any physical media including wireless.

Note that for purposes of this architecture, the CloudCO Domain Orchestrator is a conceptual container for management, control and orchestration entities (some of which may be implemented as VNFs).

Moreover, note that the Network I/O may be implemented via the existing equipment and/or uplink modules integrated into the leaf switches and/or spine switches.

Dark red lines in Figure 10 exemplify management and control associations to the various components. For example, the VIM is shown with management and control responsibility for the entirety of NFVI. Likewise, VNF Manager (VNFM) provides management and control responsibility for the lifecycle and virtualization attributes of the VNFs, and the CloudCO Domain Orchestrator is responsible for all CloudCO functional components inside the NFVI (via its NFVO component) and outside the NFVI (via the Management Control Orchestration Engine described in Section 6.3).

Figure 13 shows a detailed functional architecture of the CloudCO, including block interactions highlighted via relevant reference points.

Physical Access Nodes can include different functions based on the types of access links supported, capacity, and other factors. An OLT supporting a PON-based AN can be modeled as either a single PON access port (in which case multiple ports are modeled as multiple OLTs) or as a node terminating multiple PON access ports. An AN may include a switch that forwards traffic between multiple access ports and one or more uplink ports, or it may pass traffic directly between access and uplink ports in a 1:1 configuration (or 1:2, to account for uplink port redundancy). An AN may or may not support multicast replication.

In any event, an AN that terminates contended upstream traffic from multiple sources (e.g., a PON) must mediate real-time upstream capacity allocation according to the aggregate of Service Level Agreements (SLAs) of the contending traffic sources.

The Network I/O side of the domain may likewise comprise nothing more than facility termination, or may comprise additional functionality such as layer mapping and multiplexing, Network I/O may be implemented via dedicated equipment and/or uplink modules integrated into the leaf switches and/or spine switches.

6.1.1.1 Multi-tenancy & Multi-service Support

Figure 8 and Figure 13 show a single CloudCO Domain Orchestrator and the related Northbound API interact with the E2E Service Layer. This implies the CloudCO Domain pertains to an Infrastructure Provider who owns the underlying hybrid Physical + NFV infrastructure. Ownership, here, means also accountability for resources availability, contracted SLAs, maintenance and so on and so forth.

The support of multiple tenants and/or, within the same organization, of multiple slices tailored for specific services (e.g., fixed vs. mobile, business vs. consumer vs. wholesale) occurs via the E2E Service Orchestrator (SO).

In the following paragraphs, “Tenant” is any “big or small consumer” of the CloudCO exposed services: i.e., different Service Providers, different Service Departments of the same organization, single residential or business subscribers.

Per each tenant, one tailored client interface is instantiated and can be envisaged with a degree of complexity and depth corresponding to the nature of the tenant. All these client interfaces interact with the same CloudCO Domain Orchestrator (or more than one, if this is the case) that overall controls the CloudCO facility.

The Client Interface via which each tenant accesses CloudCO “service” is designed in a way that allows exposure of the appropriate services and resource attributes, from simple subscriber services to more structured network services.

The client interfaces that exposes these latter network services SHALL be as structured and fine grained as needed and it may replicate the very level of detail of the underlying network, i.e.:

- of the M_{inf}, M_s, M_{fc} reference points which are used to expose the network functions (including fabric functions, if needed) and/or
- of the O_{cco}-$N_{f\text{-}sdn\text{-}xxx}$ Managers&Controllers.

Service Provider type of tenants MAY have their own orchestrators and/or managers and controllers but these elements SHALL not hook directly to the CloudCO resources (Network Functions (NFs), SDN Managers&Controllers, VNFM, VIM) but only to the CloudCO Domain Orchestrator via the Os-Ma-ccodo reference point from the E2E Service/OSS/BBS Layer.

This approach is as follows:

- Operationally sound and effective: it reconciles service requests via one single brain, the CloudCO Domain Orchestrator avoiding the need for transversal coordination of multiple orchestrators accessing to the same sections of the "orchestra" (i.e., SDN elements, MANO components or even deeper into the NFs basin). Tenants are exposed to network views with different levels of detail and depths in terms of management and control of the various resources via "avatar" versions of M_x and O_{cco}-$N_{f\text{-}sdn\text{-}xxx}$ reference points, be them M_x^T and O_{cco}-$N_{f\text{-}sdn\text{-}xxx}^T$. Instead, if multiple Managers&Controllers are hooked to the NFs or multiple orchestrators are hooked to the SDN Managers&Controllers, there is a high potential of issues related to handling conflicts, inconsistencies, compatibility and interworking constraints, concurrent access, priority booking, gate keeping, etc.

  Having a the CloudCO Domain Orchestrator as the entry point for all E2E Service Layers instances associated to each tenant allows to put together all "service" requests and constraints and translates them into appropriate actions towards the physical and the NFV legs of the infrastructure. Meanwhile this approach allow to exposes a look of the M_x^T and O_{cco}-$N_{f\text{-}sdn\text{-}xxx}^T$ reference points that is predefined in terms of structure and grants but which adapts in real time to actual resources availability and constraints.

  Having multiple orchestrators access directly to the very CloudCO Domain resources would require transversal coordination, definition of roles and priorities and again a potential for issues (beyond the technical ones mentioned above, possibly also of legal accountability and regulatory nature) that must be avoided by moving the complexity to the Management Control Orchestration (MCO) engine and the Os-Ma-ccodo reference point.

  The "avatar" M_x^T and O_{cco}-$N_{f\text{-}sdn\text{-}xxx}^T$ reference points are then the appropriate hook points of tenant controllers and NFVOs (where applicable) which are fully decoupled from those within the CloudCO. For example, the ability of a tenant NFVO to exercise requests over the CloudCO NFVI depends on what and how much of the ETSI NFV Os-Ma-nfvo reference point is exposed via the Os-Ma-ccodo reference point. This is about designing and architecting of this latter reference point which is addressed in another BBF Technical Report focusing on the definition of the CloudCO interfaces.

- Robust against hassles: it decouples network related operation (fully dealt with within the CloudCO Domain) from "service" related requests; a "service" can be defined the way the Infrastructure Provider and a tenant agree to. Depending on how the Os-Ma-ccodo reference point is designed it can expose a variety of services from a basic HSI access to full blown virtual networks "leased" to different Service Provider Tenants or network slices tailored to specific applications (e.g., mobile, Content Delivery Network [CDN], Vehicle to Everything

[V2X], etc.) serving a multi-service paradigm within the same Network Operator/SP organization.

- Forward looking: using network abstraction and APIs is a future proof approach to support multi-tenancy rather than using a legacy approach, i.e., attaching multiple EMS (in a legacy world) to the managed resources.

6.1.1.2 Multi-site support

As mentioned in section 6.1.1, Figure 10 shows that a CloudCO Domain spans a MANO/NFV and SDN complex under the control of a Northbound API, but it shall not be interpreted as if a CloudCO Domain spans a single location. However, in some cases it may be appropriate to consider that a CloudCO Domain spans one location only.

This is already clarified in Section 6.1 and this is further exemplified with some deployment cases where a CloudCO Domain spans multiple sites, even at the customer premises.
Where a CloudCO does not span multiple sites the full MANO framework may not be necessary.

6.1.1.2.1 VNFs hosted in multiple locations

In some cases, VNFs that are part of the service graph can be hosted in different locations. Figure 11 is an example of how this might look like.

Sometimes these locations have NFVIs managed by a unique VIM; sometimes these locations share a common VIM. When a common VIM is deployed, one has to make sure the delay constraints (e.g., round trip time constraints of the control plane connections between VIM and NFVI), are met between the VIM in one location, and the Compute Nodes that VIM is managing in a different location. Similarly, when multiple VIMs are deployed and in order to support real time requirements, care has to be taken in terms of the delay constraints between the CloudCO Orchestrator and the remote VIM.

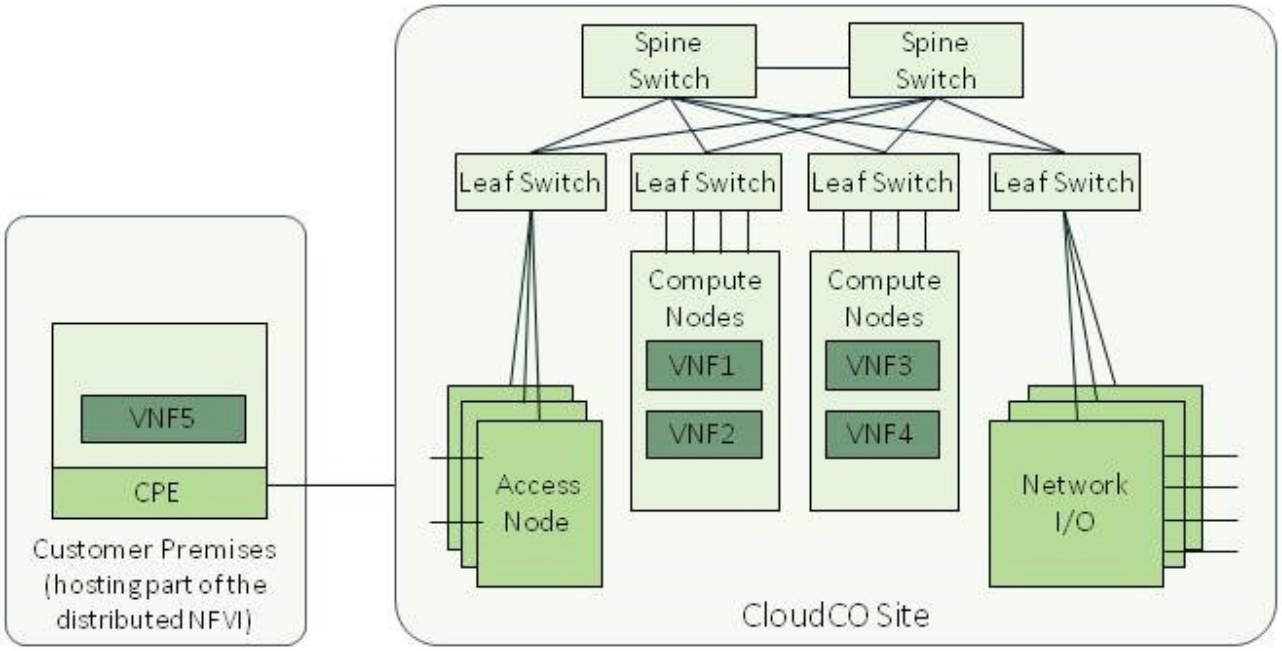

Figure 11: Example of deployment of VNF over multiple sites

6.1.1.2.2 Onboarding a Common Service for Geographically Dispersed Subscribers

In some cases, it makes sense to aggregate subscribers attached to Access I/O in different locations to a common NFVI, where the NFVI and its fabric could span multiple locations. Logically all these locations look like one CloudCO Domain. The diagram included in Figure 12 shows how this might look like ignoring the lack of the NFVI and CloudCO Domain Orchestrator. Note that the interconnection of CloudCO sites is over virtual network connections.

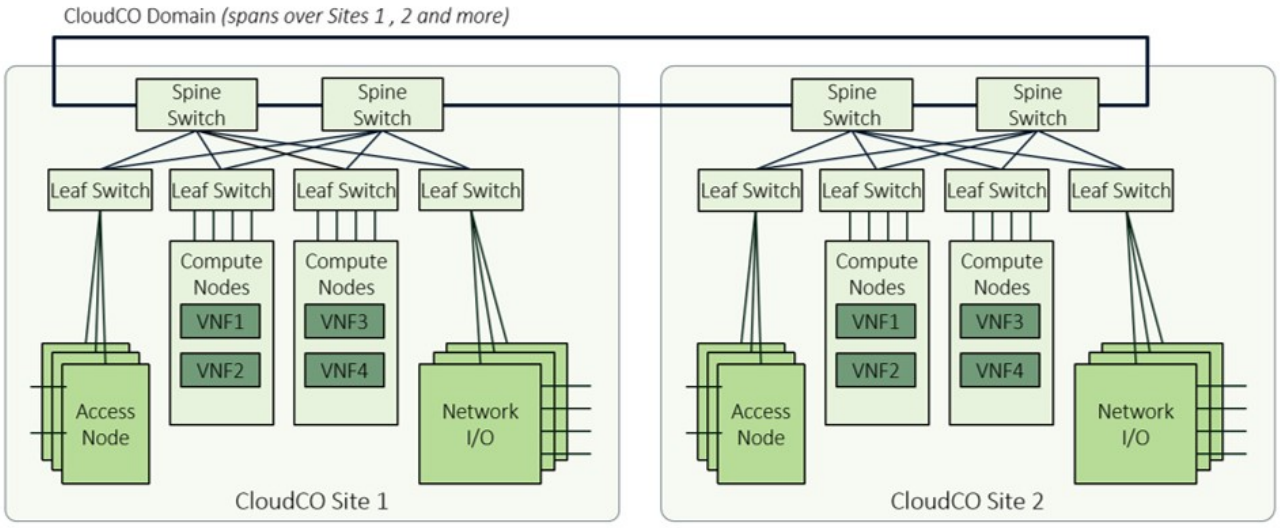

Figure 12: Example of deployment of a distributed switch fabric serving dispersed subscribers

6.1.1.2.3 Simplifying Orchestrator/SDN Hierarchy

By allowing the CloudCO NB API to span different locations, the Orchestrator/SDN hierarchy between the E2E Service Orchestrators and multiple CloudCO Domain Orchestrators can be significantly reduced. This case is depicted in Figure 9.

6.2 Functionality of the CloudCO

The concept of a CloudCO allows for arbitrary inclusion of Value Added Services, often onboarded onto the platform by a 3rd party, so there can be no complete or definitive list of the functions it contains, nor of the full set of functions that must be supported by the CloudCO NBI.

A CloudCO is at minimum an NFVI, a set of Access I/O and a set of Network I/O. The NFVI is sized according to need, and includes compute and storage (not shown) nodes, as well as a leaf-spine fabric. Section 5.2 describes additional physical considerations for access and network components, and allows for further physical devices that may be needed.

The functionality of a CloudCO is as follows.

The CloudCO must have its own life cycle management that includes the management of the components defined within the bounds of the NFV domain, as described by ETSI, but also the management of equipment, software, and facilities outside the NFV domain, including access topology and remotely located equipment. These may be provisioned or discovered, or any combination thereof, and information may be exchanged about actual and desired state with repositories that may exist beyond the CloudCO.

Such NFV and physical life cycle management functions may be provided via the NBI between E2E Service Orchestrator and the CloudCO Orchestrator, or via existing CO management interfaces, or any combination thereof.

Initially, it is expected that the CloudCO Orchestrator NB API will focus on service onboarding, instantiation, delivery and monitoring, rather than physical or virtual resource life cycle management. The NB API must therefore expose the set of services it is prepared to support. However, due to the fact that services are anchored, at least to endpoints, and probably also to other inventoried resources, the NB API must also expose the CloudCO resource inventory and topology.

Software upgrade, physical fault and performance monitoring, troubleshooting and diagnosis, are also required.

On the access side, the CloudCO must recognize the subscriber attachment and facilitate subscriber authentication, configure the subscriber equipment and its service path with subscriber-specific attributes. This may be facilitated by a DHCP agent, which may in turn consult local or remote subscriber information repositories and additional AAA or other functions (a logical TR-069 [1] Auto-Configuration Server (ACS) would be an example of a possible collaborating function). IP address assignment and Address Resolution Protocol (ARP) binding will be needed for most subscribers. CPE functionality may be partitioned between physical equipment (typically at the

customer premises) and virtual functions, typically hosted in the NFVI, which may be located in the CPE, in the Network, or both (see TR-317 [8] and TR-328 [12]).

6.3 CloudCO Domain Architecture in Detail

This section specifies the detailed logical and functional architecture of a CloudCO Domain. Figure 13 represents the logic and functional relationships among the blocks, via relevant reference points and related interfaces. The definitions of these reference points are provided in Section 6.3.1.

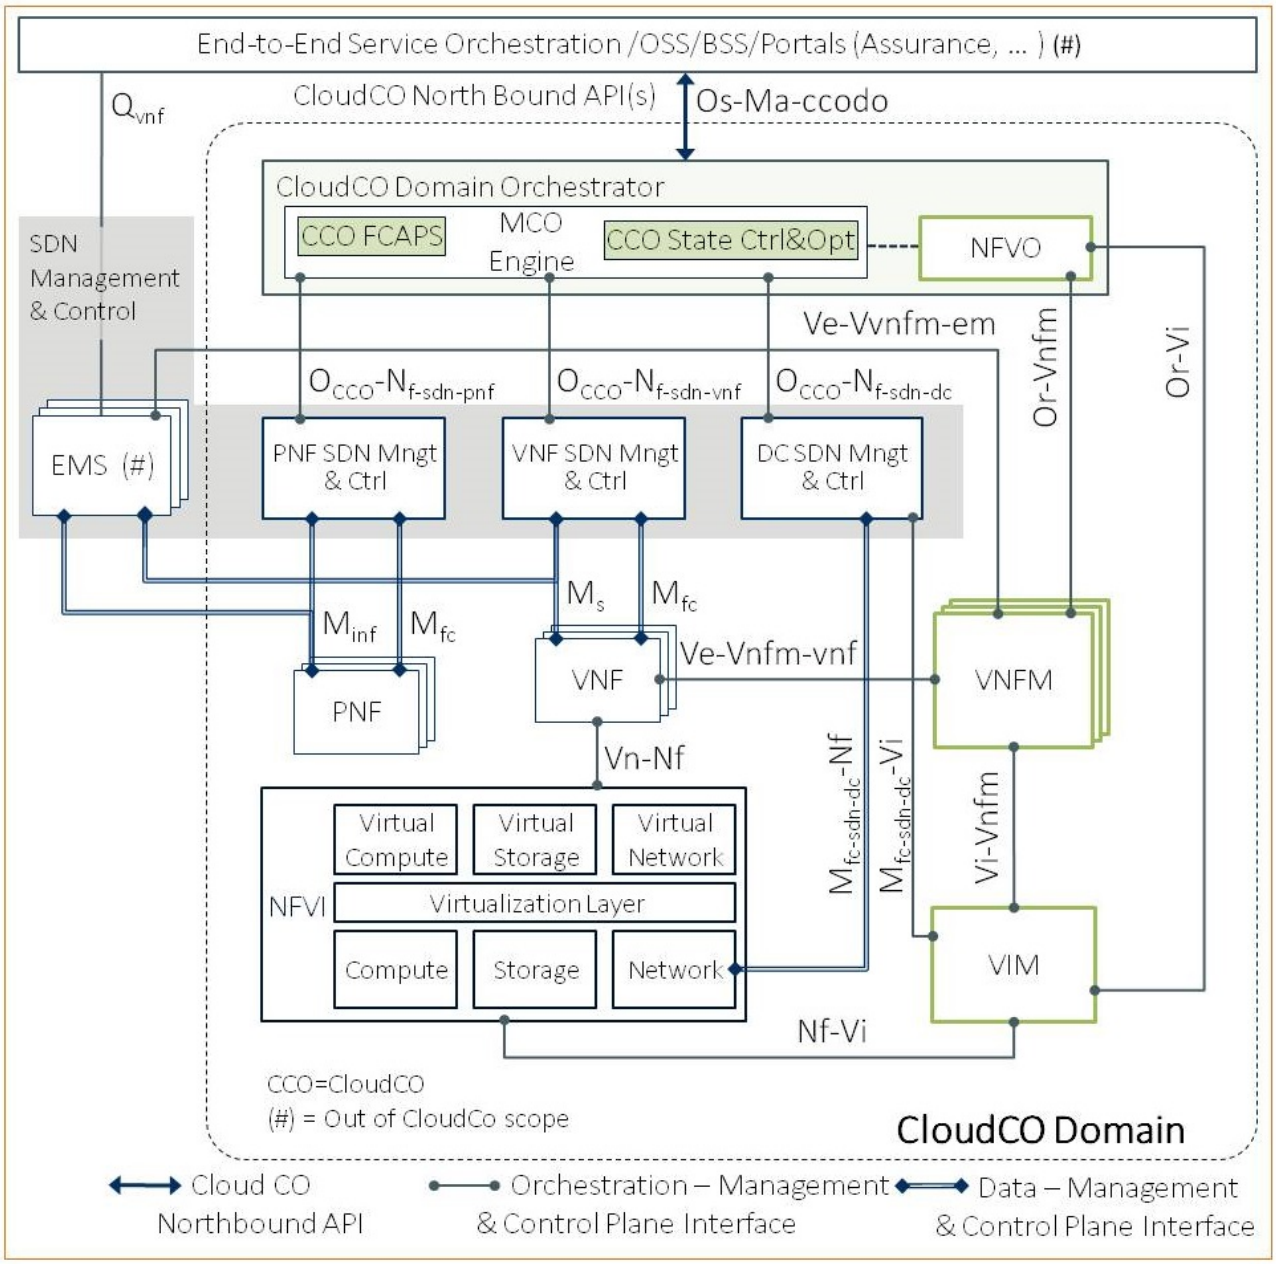

Figure 13: CloudCO reference architecture

- Green colored boxes and lines are related to NFV reference points as identified mainly in ETSI standards.
- Blue boxes and lines are related SDN reference points as identified in BBF and ONF standards.
- The CloudCO NB API is labelled as Os-Ma-ccodo and it can be considered an extension of the ETSI Os-Ma-nfvo reference point, to encompass the E2E Service layer to Network Orchestration layer relationship over a hybrid Access/Edge and NFVI domain.

As mentioned in Section 6.1, a CloudCO Domain could span multiple sites in which case it has to rely on Wide Area connectivity. This needs to be coordinated by a Wide Area SDN Controller & Orchestrator which is out of CloudCO scope and not shown in Figure 13.

The CloudCO Reference Architecture, see Figure 13, is a combination of an SDN and NFV architecture applied over a hybrid physical and NFV infrastructure. The NFV portion is used to orchestrate and manage the virtual functions and their supporting infrastructure, while the SDN portion is used to manage and control mostly user plane interactions among the PNFs, the VNFs and the switch fabric, as well as redirecting of certain in-band control packets to the relevant SDN controller applications. In order to build service segments, including forwarding across the switch fabric, as well as configuring and setting up connectivity to PNFs and VNFs (onboarded on the NFVI), all these resources have to be orchestrated together. This is the role of the CloudCO Domain Orchestrator function.

The CloudCO Domain Orchestrator is the central function in the architecture and it is also home to the CloudCO NB API, and as such it delivers the necessary Service Abstraction Layer, hiding the internal operations of the CloudCO from the NB API.

The CloudCO Domain Orchestrator operates over a distributed SDN access/edge domain relying on the resources and capabilities of a NFV Infrastructure ideally largely available in this whole multi-site domain. Under this standpoint the CloudCO Domain Orchestrator shall be envisaged as a “brain” with two coordinated “hemispheres” (i.e., The MCO Engine and the NFVO).

The MCO Engine is a specialized module within the CloudCO Domain Orchestrator that oversees all tasks related to exploiting the CloudCO Domain as a network asset. MCO Engine expresses a continuum of tasks: (1) resource/status Orchestration, (2) aggregate and individual traffic flow Control and Management and (3) CloudCO state transitions and supervision tasks.

The MCO Engine also implements the essential logic needed to interact with the NFVO via a cross-orchestration reference point internal to the CloudCO Domain Orchestrator. This tight integration of the two CloudCO Domain Orchestrator “hemispheres” is key to guarantee effective operation of the SDN elements over the management and control plane, while the MANO components ensure that the VNFs over the user plane evolve in space and time as required by the CloudCO state and configuration transitions.

The MCO Engine operates its tasks via a proper set of SBIs towards the PNF, VNF and DC SDN Managers&Controllers as directly instructed via the CloudCO NB API.

As shown in Figure 13 these Southbound Interfaces (SBIs) are in principle separated as they point to specialized SDN elements. Indeed, the nature of the interaction exercised by the MCO over them is essentially the same in that it governs the centralized Management & Control Plane of the CloudCO Domain. This is done via 'commands' to the appropriate SDN element, which manages and controls a part of the distributed user plane made of PNFs, VNFs and NFs.

As shown in Figure 13, the MCO Engine contains the following functionalities:

- CCO State Control & Optimization, which includes a closed-loop control and optimization of the state of CloudCO resources, on a time-continuous basis.
- An aggregated high-level orchestration of similar functionalities exercised in turn by the specific SDN Controllers and relevant MANO components, e.g., key Fault, Configuration, Accounting, Performance and Security capabilities. This encompasses the derived Service Assurance (SA) logics and algorithms governed centrally by the MCO Engine. Then in a waterfall fashion, SDN and MANO elements potentially embed SA logics as appropriate for their specific domain as related to their domain of action.

Generally speaking, the CloudCO NB API shall simultaneously support an arbitrary number of clients, such as OSS, applications, signaling inputs, other managers & controllers, accessing to an abstract view of the network per a Network as a Service (NaaS) paradigm. It is anticipated that each client is configurable independently, with at least security attributes, tailored information model including SLA, gatekeeper policy to protect the server from client misbehavior, and other attributes necessary or appropriate.

The CloudCO Domain Orchestrator is expected to provide a default administrative client interface, with effectively and satisfactorily available capability and visibility into the CloudCO.

As can be seen in Figure 13, a CloudCO Domain encompasses all the orchestrated functionalities triggered via a single, uniquely addressable instance of the CloudCO NB API. This means that the SDN Management&Controllers and the VNFM can operate over PNFs and VNFs, which are placed across different physical locations.

The NFV part closely follows the ETSI Reference model (NFVI and MANO, i.e., VIM/VNFM/NFVO; see the functions in green on the diagram included in Figure 13).

The ETSI architecture allows one or multiple VIMs per NFVO. Various VIMs can be deployed as part of the solution. In case of a TR-328 vBG, where virtual functions are running at the customer site, a dedicated Access VIM may be connected to the NFVO controlling the NFVI at the customer site together with a CloudCO VIM connected to the NFVO controlling the CloudCO NFVI.

Various VNFM's can be deployed as part of the solution, and they can be onboarded using the NB API. If a single VNFM is deployed, then the generic CloudCO VNFM could be made solely responsible for the lifecycle management of all VNFs. An implementation could choose to realize the generic VNFM as part of the overall CloudCO Domain Orchestrator function thereby hiding the interfaces between the generic VNFM and the NFVO (in such case the VNFM would not appear as an independent block). The CloudCO Domain Orchestrator could in that case handle the service specific VNF configuration and other FCAPS activities. The advantage of this approach is that

there is no need for a dedicated EMS functionality, while the per-VNF configuration does not need to be exposed at the NB API.

There is also an option to onboard additional VNFMs that could manage a dedicated set of VNFs. For that set of VNFs, the configuration of the service specific configuration and other FCAPS related matters could be realized by these additional VNFMs or realized via the EMSs, while the EMSs interact directly with the OSS/BSS, rather than via the CloudCO Domain Orchestrator. The EMSs themselves, and the OSS/BSS integration are out of scope of the CloudCO, but the interface to attach them to the architecture is in scope.

The SDN portion (the blocks in blue in Figure 13) controls the physical infrastructure (PNFs and switch fabric) within the CloudCO, as well as any associated external SDN controlled VNF in the user plane that might be needed.

The PNF SDN Manager&Controller interfaces with the Access Node and Network I/O devices, as well as any other devices that have PNFs that are deployed inside the CloudCO. An implementation could use a common SDN controller to interface with the Access and Network I/O devices.

The VNF SDN Manager&Controller interfaces with the VNF(s) to handle related operation dynamics of the VNF instantiated by VNFM within the CloudCO, e.g., a virtual Router controlled via an OpenFlow interface.

The DC SDN Manager&Controller directly accesses the NFVI networking resources to implement functions (e.g., L3 routes in the switch fabric) that the VIM is not supposed to do.

Figure 13 shows that the CloudCO NFs (considering also the Networking NFs [NNFs] of the switch fabric) could be accessed via different "concurrent elements" and reference points. For example, PNFs could be accessed via an EMS along with the related SDN Manager&Controller element. The same applies to VNFs which are also accessed via a dedicated reference point from the VNFM. Finally, both the VIM and the DC SDN Manager&Controller may access the network resources of the NFVI.

This multiplicity of options implies that:

- Whenever different blocks access to NFs any kind of conflict or inconsistency shall be avoided by means of proper orchestration resolved at the appropriate reconciliation point depending on the involved "concurrent elements" (i.e., the CloudCO Domain Orchestrator or the E2E Service Orchestrator). Alternatively this reconciliation could be exercised via a direct interface between the involved blocks; this is the case of the VIM and the DC SDN Manager&Controller which, as shown in Figure 13 can access the NFVI networking resources via the [Nf-Vi/]N and the $M_{fc\text{-}sdn\text{-}dc}$-Nf reference points respectively; to avoid mismanagement of these networking resources the $M_{fc\text{-}sdn\text{-}dc}$-Vi reference point allows coordination of the two blocks allowing also a client-server relationship that is dependent on the actual features of the SDN Manager&Controller and the VIM, e.g., the DC SDN Manager&Controller may access the NFVI networking resources only through the VIM via the $M_{fc\text{-}sdn\text{-}dc}$-Vi reference point and the $M_{fc\text{-}sdn\text{-}dc}$-Nf reference point is simply not

implemented or reversely the VIM may access the NFVI networking resources via the DC_SDN controller and the [Nf-Vi/]N interface that is only partially implemented.

- Generally speaking the scope of management and control exercised by "concurrent elements" over the same NFs shall be complementary; this is particularly meaningful for the pairs VNF SDN vs. VNFM and DC SDN vs. VIM,
- The "concurrent elements" over a set of NFs in Figure 13 can also represent two distinct points in time of a migration path from legacy managers towards evolved Managers&Controllers with enhanced functionalities and innovative interfaces.

The management and control interface to pre-existing physical functions need not be revised. Adaptation software can exist in VNF form, resident in the NFVI, supporting the necessary protocols and data models (e.g., OMCI [22], TR-069 [1], proprietary protocols, etc.).
An example of such mediation software is the Broadband Access Abstraction layer described in Section 6.3.2.7.

6.3.1 CloudCO Reference Points & Catalogs

This sub-section defines relevant CloudCO reference points based upon the architecture in Figure 13 and relevant catalog in the following tables.
Table 2 describes ETSI-NFV reference points defined in [14], [15], [16], [17], [18], [19] and [20] and that are relevant for the CloudCO architecture.

Reference Point	Location	Description
Os-Ma-nfvo	Between the CFS/RFS Orchestration and Assurance and the NFVO	The Os-Ma-Nfvo reference point provides management of Network Service Descriptors and VNF packages; lifecycle management of Network Services and VNFs; Policy management and/or enforcement of Network Services, VNFs and NFVI resources. [20] describes the information exchanges across this reference point. In the CloudCO context this interface may be (partly) exposed via the Os-Ma-ccodo reference point.
Or-Vnfm	Between the NFVO and VNFM	The Or-Vnfm reference point provides management of NFVI resources for a VNF including information needed for authorization, validation, reservation, allocation and release of NFVI resources. In addition, lifecycle management of VNFs is provided. [18] describes the information exchanges across this reference point.
Ve-Vnfm-em	Between the EMS and the VNFM	The Ve-Vnfm-em reference point provides life cycle management of VNFs managed by the EMS and VNFM. [19] describes the information exchanges across this reference point.
Or-Vi	Between the NFVO and the VIM	The Or-Vi reference point provides management of NFVI resources including information needed for allocation and release of NFVI resources. [16] describes the information exchanges across this reference point.

Vi-Vnfm	Between the VIM and the VNFM	The Vi-Vnfm reference point provides management of NFVI resources including information needed for validation, reservation, allocation, update and release of NFVI resources. In addition, software image management of VNFs is provided. [17] describes the information exchanges across this reference point.
Ve-Vnfm-vnf	Between the VNF and the VNFM	The Ve-Vnfm-vnf reference point provides lifecycle management of VNFs managed by the VNFM. [19] describes the information exchanges across this reference point.
[Nf-Vi/]N	Between the VIM and NFVI Layer	The NF-Vi reference point is comprised of sub-reference points for interfaces associated with requesting infrastructure connectivity services ([Nf-Vi]/N), hypervisor services ([Nf-Vi]/H) and compute services ([Nf-Vi]/C). The [Nf-Vi/N] reference point is described in clause 5.2 of [15] and Clause 5.7.4 of [14] describes the information exchanges across this reference point.
[Vn-Nf]/N	Between the NFVI Layer and the Network Service Layer	The [Vn-Nf]/N reference point provides transparent network services to VNFs as described in clause 5.1 of [15].

Table 2: ETSI NFV Reference Points

Table 3 specifies additional reference points required for the interactions among the blocks defined in the CloudCO architecture.

Reference Point	**Location**	**Description**
M_{inf}	Between the (Resource Facing Service) RFS Orchestration and Assurance, ACS, EMS or PNF SDN Manager&Controller and the network elements within the physical infrastructure	This is the reference point for FCAPS on infrastructure NEs in the MSBN.
M_s	Between the RFS Orchestration and Assurance, ACS, EMS or VNF SDN Manager&Controller and the VNFs.	This is the reference point for FCAPS on user facing VNFs in the MSBN.
M_{fc}	between SDN Manager&Controller and the Network Functions.	This is the reference point for Flow Control of NFs (e.g., service parameters and forward table configuration).
Q_{vnf}	between the RFS Orchestration and Assurance and EMS	This is the reference point to manage the lifecycle aspects of the VNF in the case where the EMS utilizes the Ve-Vnfm-Em reference point.
O_{cco}-$N_{f\text{-}sdn\text{-}pnf}$ O_{cco}-$N_{f\text{-}sdn\text{-}vnf}$	Between the CloudCO Domain Orchestrator and the SDN	This is the reference point for the CloudCO Domain Orchestrator to interact with the

O_{cco}-$N_{f\text{-}sdn\text{-}dc}$	Managers and Controllers	SDN Managers and Controllers.
$M_{fc\text{-}sdn\text{-}dc}$-Nf	Between the DC SDN Controller and the NFVI physical networking resources	This is the reference point for the DC SDN Controller to access to the NFVI physical networking resources, mainly to exercise connectivity at L2 and L3.
$M_{fc\text{-}sdn\text{-}dc}$-Vi	Between the DC SDN Controller and the VIM	This is the reference point for the DC SDN Controller and VIM coordination to manage NFVI networking resources.
Os-Ma-ccodo	Between the E2E Service Orchestrator and the CloudCO Domain Orchestrator	The Os-Ma-ccodo reference point allows to expose an abstracted view of the CloudCO Domain resources to the Northbound Service Orchestration (and demand) entities and to consume such resources per a Network-/Subscriber-as-a-Service paradigm. Given the CloudCO Domain Orchestrator contains a NFVO, this reference point may expose directly to the E2E Service Orchestrator the functions associated to Os-Ma-Nfvo reference point: management of Network Service Descriptors and VNF packages; lifecycle management of Network Services and VNFs; policy management and/or enforcement of those resources and the NFVI in general. Alternatively, the Os-Ma-Nfvo reference point may be embedded CloudCO Domain Orchestrator, between the NFVO and the MCO Engine and this latter would be then the only block to interface with the E2E Service Orchestrator.

Table 3: BBF Reference Points

Table 4 and Table 5 describe the ETSI NFV and additional BBF repositories that support the NFVO function operating in the hybrid SDN-NFV context of the CloudCO.

In an architectural document like this Technical Report, these tables represent examples that shall be refined at the stage of actual software implementation.

The catalog of orchestration, control and management packages in Table 5 assumes that these functions are implemented as VNFs. That would not be applicable in the case where they are prebuilt applications already running at the time of creation of the CloudCO Domain.

Catalog	**Primary Entities**	**Description**
NS Catalog	NFVO	Represents the specifications of all of the on-boarded Network Services, VNFs and NFVI Resources. The NS Catalog is further

		described in clause 5.4.4 of [14].
VNF Catalog	NFVO, VNFM	Represents the specifications of all of the on-boarded VNF Packages, supporting the creation and management of the VNF Package (VNF Descriptor (VNFD), software images, manifest files, etc.). The VNF Catalog is further described in clause 5.4.5 of [14].
VNF Instance Repository	NFVO	The NFV Instances repository holds information of all VNF instances and Network Service instances. Each VNF instance is represented by a VNF record, and each NS instance is represented by an NS record. Those records are updated during the lifecycle of the respective instances, reflecting changes resulting from execution of NS lifecycle management operations and/or VNF lifecycle management operations. The VNF Instance Repository is further described in clause 5.4.6 of [14].
NFVI Resource Repository	NFVO	Represents information about available/reserved/allocated NFVI resources as abstracted by the VIM across Operator's Infrastructure Domains. The NFVI Resource Repository is further described in clause 5.4.7 of [14].

Table 4: ETSI NFV Data Repositories

Catalog	**Primary Entities**	**Description**
CloudCO Orchestrators, VNF Catalog of Controllers & Managers, Mediation Layers	NFVO, VNFM	Represents the specifications of all of the on-boarded VNF "management" Packages specialized in orchestration, control and management functions. This catalog encompasses, for example, the CloudCO Domain Orchestrator, the PNF-VNF and DC SDN Manager&Controller VNFs. An example of a Mediation Layer is the Broadband Access Abstraction layer.

Table 5: BBF Data Repositories

6.3.2 CloudCO blocks role and functions

This section provides a definition of roles and functions of management, control and orchestration elements in the CloudCO architecture.

6.3.2.1 NFV Orchestrator (NFVO)

In the CloudCO context the NFVO is a component of the CloudCO Domain Orchestrator.

It has two main responsibilities:

- The orchestration of NFVI resources across multiple VIMs.
- The lifecycle management of Network Services.

For a complete list of NFVO capabilities refer to section 5.4 of [14].

6.3.2.2 VNF manager (VNFM)

The VNF Manager is responsible for the lifecycle management of VNF instances.

For a complete list of VNFM capabilities refer to section 5.4 of [14].

6.3.2.3 Virtualized Infrastructure Manager (VIM)

The VIM is responsible for controlling and managing the NFVI compute, storage and network resources, usually within one Operator's infrastructure domain.

For a complete list of VIM capabilities refer to section 5.4 of [14].

6.3.2.4 Management Control Orchestration (MCO) Engine

The MCO Engine is a component of the CloudCO Domain Orchestrator.
It expresses a continuum of MCO tasks as well as CloudCO state transitions and supervision tasks.

The following list expresses the set of functions performed by the MCO Engine:

- Acting as the interaction module between the CloudCO Domain and service level systems (E2E Service Orchestrator, OSS/BSS, etc.) via the CloudCO North Bound API.
- Exposing via the CloudCO North Bound API an abstract view of the hybrid network topology and related status of physical and virtual resources therein instantiated. Examples of exposed information related to the CloudCO Domain status are the full configuration and tenant ownership of the instantiated resources or their current availability for consumption from the E2E Service Orchestrator, as well as any notification of resource alarms status or resource out of order maintenance conditions.
- Acting as the upper level of a two-level SDN hierarchy enabling real time dynamics and advanced network automation of the CloudCO Domain edge and access resources, i.e., the PNFs, VNFs and Networking functions which collectively realize the CloudCO Domain user plane.
- Orchestrating the tasks of specialized Manager&Controller blocks related to PNFs, VNFs and DC Networking I/O resources.
- Providing internally an advanced Data Collection Function (DCF) related to monitoring and surveillance of network performance parameters and other data related to the CloudCO Domain instance lifecycle diagnostics.

In order to achieve the above, the MCO Engine must also be provided with:

- A Service Building Framework, where service developers can use ‘building-block’ services to create their own aggregate services, along with an API to consume them. Aggregate

services can be used to create yet other aggregated services. Building-block services can be single-tenant, or can be an instance of a multi-tenant application.
- The essential interaction with the NFVO via an orchestration reference point internal to CloudCO Domain Orchestrator
- An internal FCAPS function, depicted as a single block in Figure 13, which is responsible for the global FCAPS management framework of the CloudCO Domain.
- An internal CloudCO State Controller and Optimizer function, depicted as a single block in Figure 13, featuring closed loop tracking and optimization of the CloudCO resources. Given the current internal status of the CloudCO Domain and the pending requests from the Service level systems, such as for addition/release/reconfiguration of available resources, this function computes and guides the transition to a new stable and resilient state of the CloudCO Domain. This new state is then reached via actuation of a proper set of MCO Engine commands at the SBI towards the SDN Manager&Controllers, potentially preceded by an NFVI scale in/out action, requested by the MCO Engine to the NFVO.

6.3.2.5 PNF and VNF SDN Managers&Controllers

The PNF and VNF SDN Managers&Controllers are responsible for FCAPS and Flow Control management functionalities respectively for PNFs and VNFs. They implement the Management & Control Plane that governs the overall user plane represented by the service graphs that chain PNF, VNF and networking resources together within the CloudCO Domain.

The following list expresses the non-exhaustive set of functions performed by the each of these SDN Managers&Controllers over their target NFs. These functionalities are exposed by means of interfaces and consumed by the CloudCO Domain Orchestrator:

- Configuration for the NFs provided by the PNFs and VNFs.
- Fault management for the NFs provided by the PNFs and VNFs.
- Accounting for the usage of PNFs and VNFs.
- Collecting performance measurements for the functionalities provided by the PNFs and VNFs.
- Security management for the PNFs and VNFs.
- Overall flows control across the space-time distributed user plane represented by the PNFs and VNFs as interconnected via the networking resources within the CloudCO Infrastructure domain and complemented by wide area connectivity services.

6.3.2.6 DC SDN Manager & Controller

The DC SDN Manager & Controller directly accesses the NFVI networking resources to implement functions (e.g., L3 routes in the switch fabric) outside the scope of VIM control.

6.3.2.7 Broadband Access Abstraction (BAA) layer

The functions identified in Sections 5.2.6 and 5.2.7 as candidates for disaggregation can be implemented within NFVI in the CloudCO. Some of these functions, such as those associated with subscriber management, can be coupled with associated functions in a controller or other virtualized elements. Other functions such as configuration, reporting and alarming remain associated with access network devices even once they are virtualized. In the CloudCO architecture, these functions are located logically in between one or more control and management elements which interact with access network devices themselves.

This calls for the introduction of an abstraction layer to expose, via a standardized NB API, a simplified functional view of access devices which is vendor independent and in some cases also access technology independent.

The concept of access abstraction is a powerful paradigm not only under a network operations perspective but also as an enabler of lean industrial processes for network production (e.g., for design, update and upgrade cycles) both on the suppliers' side and on the Operators side.

The functional diagram for such an abstraction layer is shown in Figure 14.

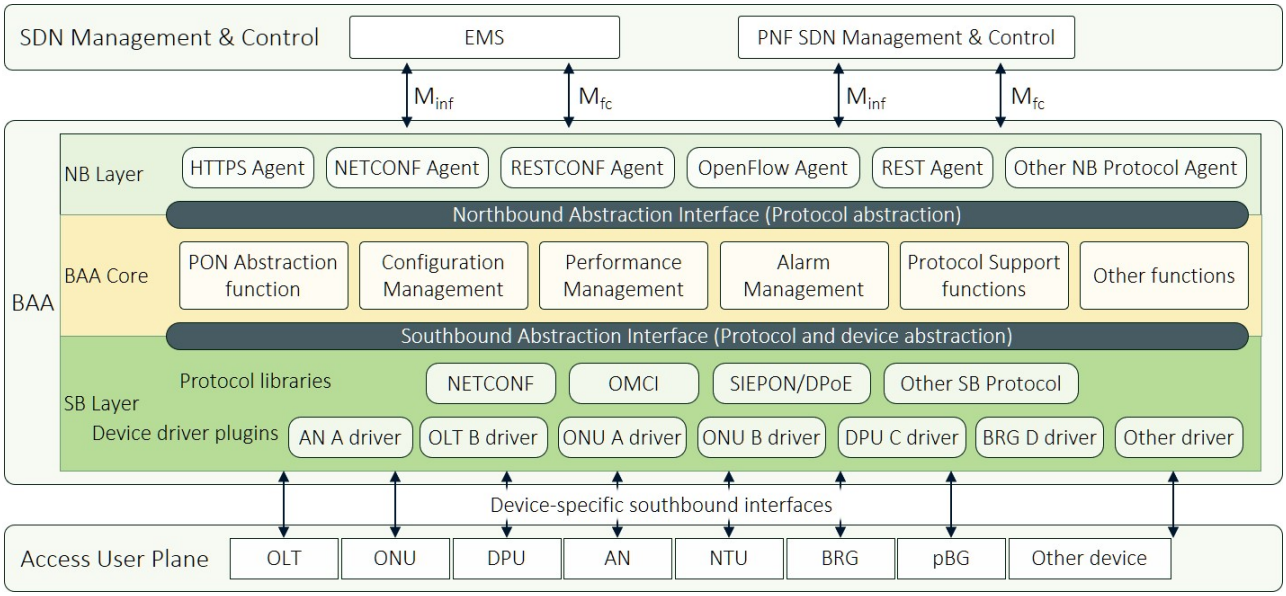

Figure 14: Broadband Access Abstraction layer

The Broadband Access Abstraction (BAA) layer fits within the overall CloudCO architecture (refer to Figure 13) as shown in **Figure 15**. As seen in **Figure 15**, not all network functions need or use a BAA layer. VNFs and Non-Access PNFs can interface directly with their SDN Management and Control function. Access PNFs may be accessed directly or via a BAA layer. For those Access PNFs that are managed and controlled via a BAA layer, the interface is modified as shown on the left hand side of **Figure 15**. Access and Non-Access PNFs that interface directly to Management and Control functions are shown in the middle of **Figure 15**.

The northbound interface from the BAA layer to the PNF SDN Management and Control functions is functionally the same as for other PNFs (i.e., directly managed Access and Non-Access PNFs), consisting of the M_{inf} and M_{fc} interfaces. However, the BAA enables a clean separation of device management in M_{inf} and flow control in M_{fc} that may not be possible in directly managed PNFs. More importantly, it allows a set of PNFs (such as an OLT and multiple ONTs within a PON) to be controlled as a single network function over M_{fc}. In this case the BAA Core takes care of translating L2/L3 "requests" from the PNF SDN Management and Control functions into technology and device specific settings.

The southbound interface from the BAA layer to Access PNFs may be device specific as discussed below.

The elements in the SDN Management and Control Plane in **Figure 15** are the same as their corresponding elements in Figure 13. As in Figure 13, the EMS elements are out of scope for the CloudCO framework, but they use the M_{inf} or M_s interfaces defined in CloudCO as shown.

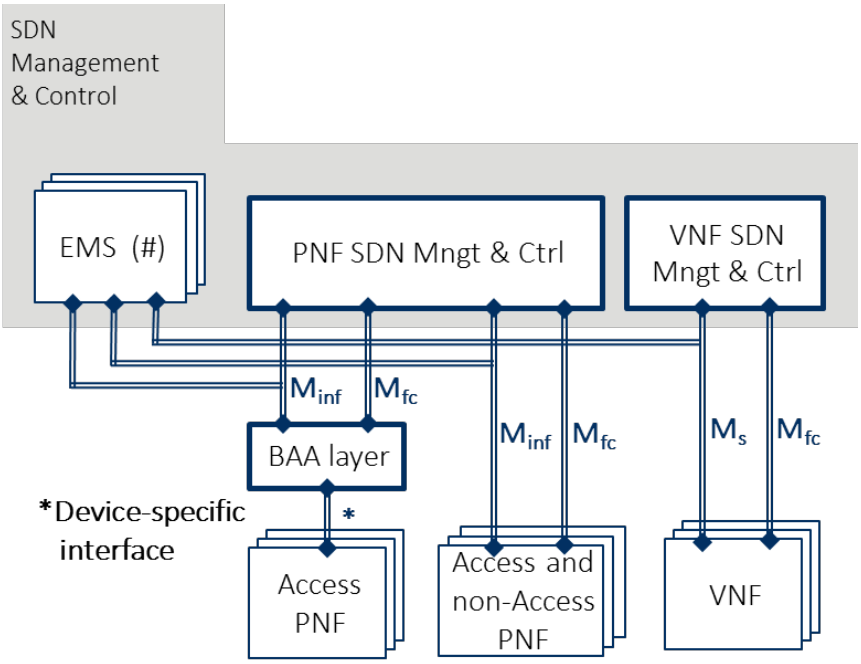

Figure 15: BAA within the CloudCO architecture

The disaggregated functions contained within the Core of a Broadband Access Abstraction (BAA) layer can be defined independent of the protocols used to communicate with the network elements to the south or with the control and management elements to the north. This independence is achieved by using software abstraction interfaces at the northbound and southbound edges of the BAA Core (NBI, SBI, respectively) combined with protocol-specific (and device-specific in the case of SBI) plugins which convert generic messages from the northbound control and management elements into protocol-specific messages used to communicate with external network elements and vice versa.

Via the BAA layer, for example, a service for a Fiber to the Home subscriber can be configured, without needing to be aware of the vendors of OLT and ONU, or the specific PON technology deployed (apart from knowing what rates can be offered). Additionally, the BAA layer will take care of coordinating ONU and OLT configuration and the specific protocol used between them. The same applies to other FTTx deployments and related access technologies (e.g., FTTdp/G.fast, Fiber To The Antenna (FTTA)/ Gigabit Ethernet (GbE), etc.).

When the BAA layer virtualizes the configuration and management of a PON, the OLT and ONUs comprising the PON are exposed northbound as a single AN. A PON abstraction function in the BAA core performs the conversion between this abstract view at the NBI and the management, configuration, reporting and alarming functions for each of the physical devices comprising the PON at the SBI, hiding the device level details.

The BAA layer's SBI contains device driver plugins that support communication with the access devices in the network. A device driver plugin may be device-specific in that it provides a low-level interface to device hardware or accommodates vendor variations between devices. Alternatively, it may be a generic device driver designed to interface with devices from multiple vendors. In either case, the device driver must comply with the Southbound Abstraction Interface API which is the standards-based interface between the BAA Core and the access devices. This interface is specified via the required data models and procedures, which in turn govern the interworking between systems from different sources. Since communication with specific devices relies on device driver plugins, the interface between the plugin and the device is device-specific and is not considered a point of interworking. Device driver plugins may use southbound protocol libraries provided as common resources, or they can embed their own protocols as needed.

At northbound the BAA layer communicates with one or more control and management elements which may include access network managers, SDN controllers, and orchestrators. These elements may use different protocols to communicate with functions in the BAA Core. By applying protocol plugins at the BAA layer's NBI, the elements to the north can be redefined and interfaces can be updated to use a different protocol (e.g., RESTCONF or NETCONF) with minimum redesign.

The Northbound Abstraction Interface is specified via the required data models and procedures, which in turn govern the interworking between the BAA Core and the control and management elements to the north. Since the NBI relies on common protocol plugins, there is no analogy to the device-specific interfaces at the SBI, and the behavior of the protocol plugins should be defined by standards per each of the applicable protocols. However, the data carried by the protocols is specified at the abstraction interface.

A broad range of access devices may be managed via functions virtualized in the BAA layer. These devices include:

- OLTs and ONUs TR-156 [3], TR-167 [4].
- DPUs TR-301 [7]. Both the PMAs for a population of DPUs and the PMA Aggregator may be implemented via BAA.
- CPE equipment TR-069 [1], including Bridged RGs TR-317 [8] and physical Business Gateways (pBGs) TR-328 [12].

- Other ANs TR-101 [2], TR-178 [5].

The BAA layer can be implemented as a virtual function hosted on the CloudCO NFVI. It can also be embedded within a physical AN, potentially enabling upgrade and migration paths for legacy access devices. Wherever it is hosted, it must meet requirements associated with CloudCO dynamics, including those for virtual function onboarding and lifecycle orchestration. For example, a BAA layer hosted within an OLT and providing virtualized management for the devices on a PON must be able to support device-plugins for all devices on the PON such as ONUs and DPUs.

6.3.2.7.1 Example of BAA applied to a PON AN

The abstraction of a PON AN is illustrated in **Figure 16** to show how the BAA layer presents a northbound abstraction of a PON AN in the forwarding plane as a logical switch hiding its PON-based implementation, while high level commands are translated to OLT and ONU detailed configurations through the BAA layer.

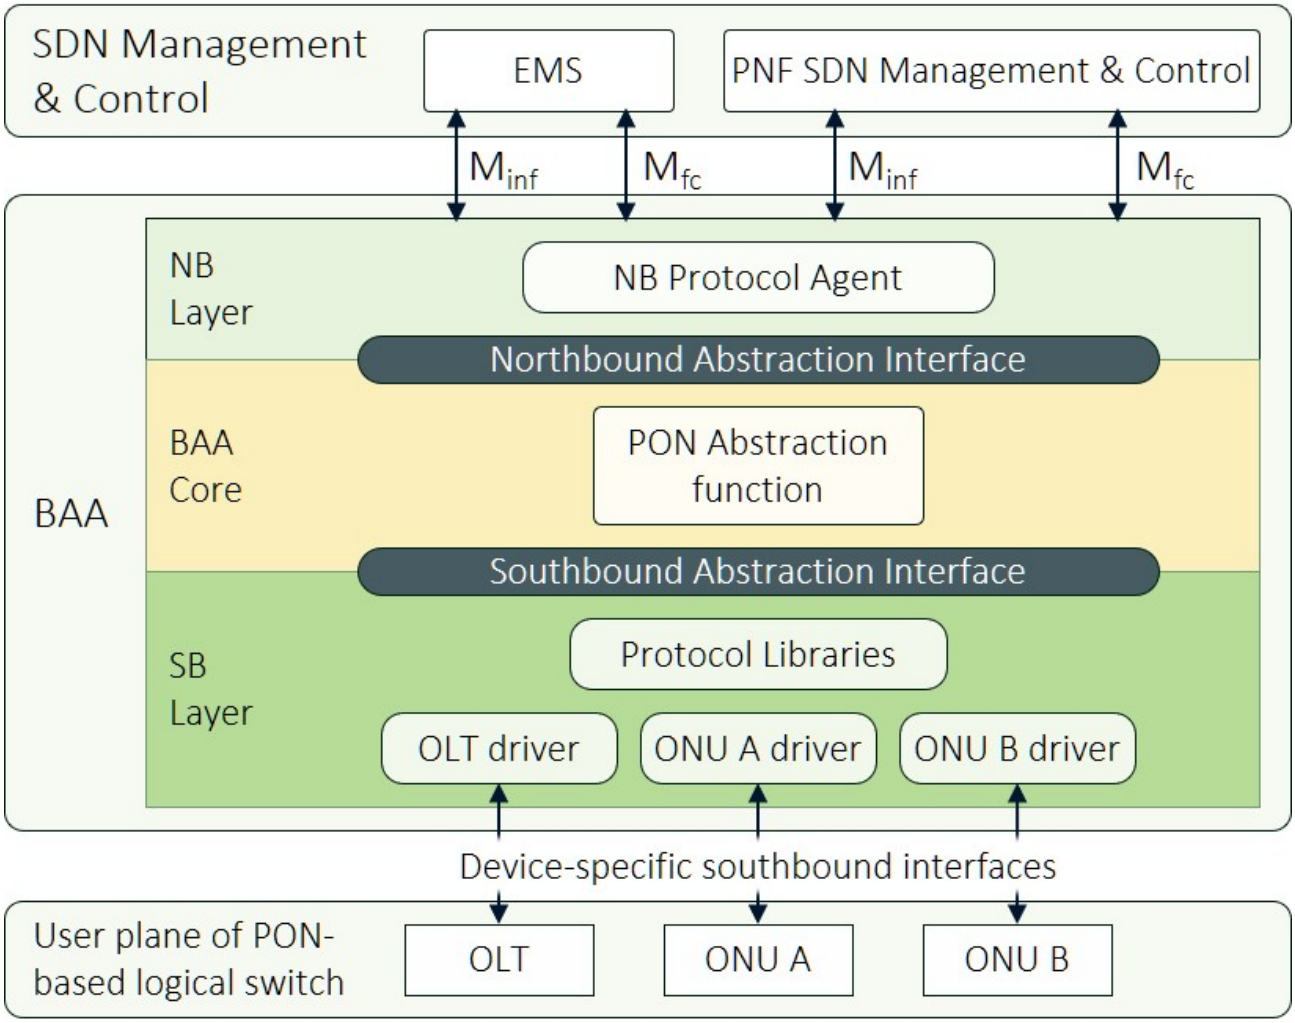

Figure 16: BAA applied to a PON AN

The CloudCO control and management elements configure the forwarding entries in the logical switch needed for the processing of packet classification, VLAN ID (VID) action, forwarding, or other services. Based on some decomposition policies, the PON Abstraction function within the BAA Core translates the forwarding entry commands into the forwarding entries for the OLTs and ONUs, to map the ingress/egress ports of the logical switch to the ingress/egress ports on the OLT and ONUs.

6.4 Network QoS Framework

6.4.1 Overview

A CloudCO is expected to support a number of business relationships. By definition, these are going to be constructed from pairwise agreements between seller and purchaser. Further, it is not hard to envision complex recursive hierarchies whereby numerous classes of resources are aggregated via such pairwise relationships, additional value added, and sum then resold to a further purchaser higher in the food chain.

For the purpose of discussing network QoS, a simplification is necessary and made possible by the fact the network offers types of resources, such as bandwidth, latency, priority and availability, anchored at a particular point in the value chain. Therefore, this framework considers three actors in a recursive arrangement of buyer/seller pairwise relationships; a customer who is the end consumer of a service offered by a retailer, a retailer who purchases networking resources from one or more wholesalers as components of the service they offer to connect their service to their customer, and the wholesaler who has an inventory of network resources to sell.

In traditional networking a user to network interface (UNI) is a pairwise interface between a service end point and a service access point and reflects a business relationship between a purchaser and a seller, and the pairwise nature of the interface implies a specific 1:1 cardinality between resource consumption on the part of the purchaser and policing of resource consumption on the part of the seller, functions located at the opposite ends of the UNI.

It is necessary to revisit this in the context of CloudCO because an architecture that uses cloud technologies provides the potential of a shared UNI such that a 1:1 cardinality between consumption and policing does not necessarily exist, and it is desirable to leverage this property in how virtualized services are architected.

The following architectural principles and requirements guide this framework:

1. A CloudCO is expected to be a convergence vehicle such that both the access and the WAN will be shared by a number of users with heterogeneous needs. Therefore a QoS framework is necessary for fair and proper sharing of network resources. A CloudCO architecture should be able to support equivalent QoS capabilities as defined in TR-101 [2] and TR-178 [5] .
2. A CloudCO has deconstructed the 1:1 relationship between cloud hosted network functions and either access or WAN network resources
3. Irrespective of the implementation, a UNI is a business interface and therefore needs to offer clear separation of interests, roles and responsibilities. The buyer will send network traffic to

the seller with shaped and scheduled characteristics in accordance with an agreed profile, and the seller will police the received traffic to that profile. Profile compliant traffic will be delivered by the wholesaler to the far end in accordance with agreed transfer characteristics.

4. The network internal to the cloud is assumed to be an any-to-any non-blocking network designed to minimize congestion, hence latency and jitter, but is clear that it is a simplification and may depend on its topology. This permits the function that polices access or WAN resources to be placed at the periphery of the cloud network, and allows maximum flexibility in architecting network functions, while allowing minimal complexity in the fabric implementation. Non-blocking also ensures that the transfer characteristics of the fabric have minimal impact on contracted transfer characteristics.
5. A cloud network may offer transit to services that have been shaped and policed elsewhere in the network. This would be in a situation where the only resources local to the CloudCO consumed by a retailer were network resources, or that the retailer had architected their service offering as a combination of functions deployed both local to and remote from the CloudCO.
6. Network resources in the access are sold as a P2P service between a customer physical facility and either retailer end point(s) the CloudCO or some other end point in the wholesaler's network.
7. More than one P2P service may transit the customer physical facility. In which case the service is delineated by VLAN marking.
8. Network resources in the WAN are sold as either transit; a P2P service between a retailer or customer and some remote network end-point, or peering; a P2C service in which a policed amount of traffic is handed off possibly to a third-party wholesaler.

Note that the following topic is out of scope of this framework:

- Access and WAN network resources that are sold independently of the network resources internal to the cloud. Cloud fabric resources may be separately policed at the point of retailer attachment to the fabric. This framework does not preclude this, but it does not address this point as the current state of the art. Note that if the cloud fabric resources are separately policed at the point of retailer attachment to the fabric then it will not be possible for "access to access", "access to WAN", and "access to cloud internals" to be separately policed.

Figure 17 illustrates a useful subset of the possible connectivity configurations to illustrate some key cases to be addressed.
In each of the connectivity instances, the portion of a path between a shaping instance and a policing instance is a UNI:

- Connectivity instance 1 is a Point to Point (P2P) access service that connects a UNI on a physical facility connecting a customer site to a retailer UNI within the CloudCO. The customer and retailer ends of the UNIs have traffic shaped to a profile agreed with the wholesaler, which is policed by the wholesaler on the network side of the UNIs.
- Connectivity instance 2 is (from the point of view of the CloudCO) a P2P transit service that connects a UNI on a physical facility on a customer site to a remote endpoint outside the CloudCO. The traffic will transit the CloudCO but is assumed to have been both shaped and policed by external agents to seamlessly fit into the CloudCO, access and WAN traffic matrices.

- Connectivity instance 3 is a P2P WAN transit service connecting a UNI to retailer to an external point remote on the WAN.
- Connectivity instance 4 is a Point to Cloud (P2C) WAN service connecting a UNI to a typically Multi-Point to Multi-Point (MP2MP) network. The retailer shapes the outbound traffic and the wholesaler polices it according to an agreed contract. There may not be a corresponding far end UNI that polices traffic, as this may be Internet access via local peering instead of a P2C service with an SLA. Examples of a P2C service would include Layer 2 VPN (L2VPN), Layer 3 VPN (L3VPN).
- Connectivity instance 5 is a P2P access services that connect a UNI on a physical facility connecting a customer site to a retailer within the CloudCO, the Cloud fabric being the vehicle whereby multiple VNF user side service endpoints are able to connect to a single network side service access point. The customer and retailer ends of the UNIs have traffic shaped to a profile agreed with the wholesaler, which is policed by the wholesaler on the network side of the UNIs. Note the same class of construct could also be used for communication to the Network I/O.

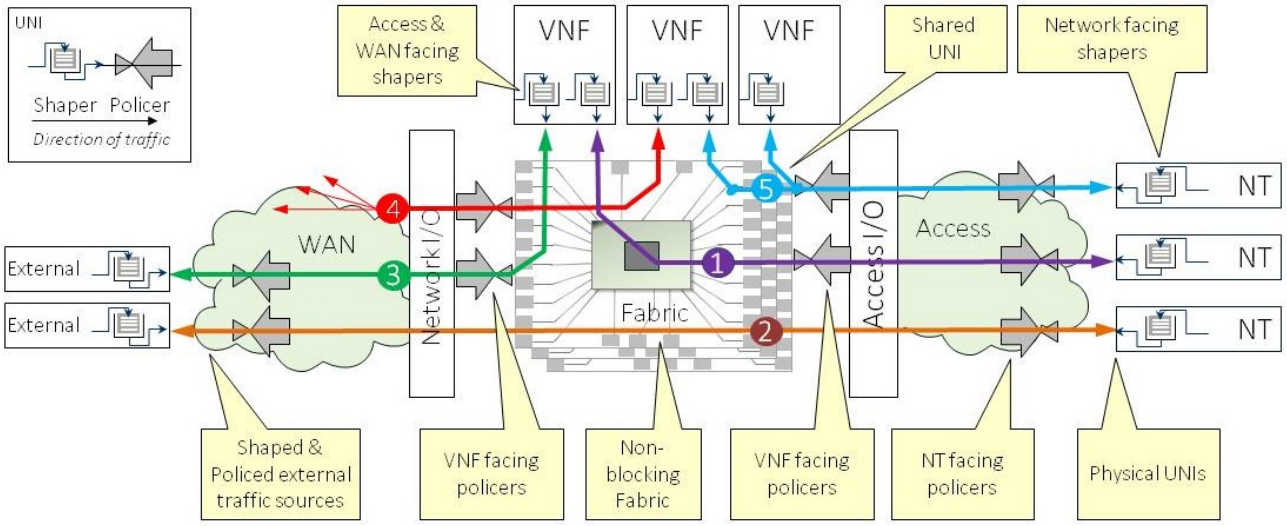

Figure 17: Possible connectivity configurations

6.4.2 Connectivity Constructs

This section provides a brief description of possible connectivity configuration constructs. The figures shown in this section contain several new BBF reference points that are listed in Table 6.

Interface	Description
W_A	The WA reference point identifies the service end point at the retailer end of a shared UNI (bracketed by WAA and WA) between the wholesaler and retailer domains in an access facing service. This subclasses the ETSI NFVI [Vn-Nf]/N interface.
W_{AA}	The WAA reference point identifies the service access point at the wholesaler end of a shared UNI (bracketed by WAA and WA) between the wholesaler and retailer domains in an access facing service. The cardinality of WA to WAA is not necessarily 1:1, nor are the protocol stacks at WA and WAA necessarily the same.

W_I	The W_I reference point identifies the interface internal to a VNF system.
W_N	The WN reference point identifies the service end point at the retailer end of a shared UNI (bracketed by WNN and WN) between the wholesaler and retailer domains in a network facing service. This subclasses the ETSI NFVI [Vn-Nf]/N interface.
W_{NN}	The WNN reference point identifies the service access point at the wholesaler end of a shared UNI (bracketed by WNN and WN) between the wholesaler and retailer domains in a network facing service. The cardinality of WN to WNN is not necessarily 1:1.
X_N	The X_N reference point identifies the boundary between the Network I/O and the NFVI for transit traffic.
X_A	The X_A reference point identifies the boundary between the Access I/O and the NFVI for transit traffic.

Table 6: New BBF reference points

6.4.2.1 The 1:1 Access Model

The 1:1 model sees the wholesaler offer a P2P service (in Metro Ethernet Forum (MEF) terms referred to as Ethernet – LINE (E-LINE)) between the retailer and a specific customer network termination, see Figure 18. This is identified at the retailer side termination as a LAN instance identified by either native TR-101 [2] tagging (single S-tag, S-tag/C-tag depending on the TR-101 tagging mode employed by the wholesaler) or includes a virtual network identifier native to the NFVI as defined by IETF NVO3. If the NFVI fabric is not native Ethernet, then the NFVI will interwork between the technology used within the NFVI fabric and the TR-101 [2] tagging in the access I/O. This is identified at the customer side termination as either a tagged or untagged UNI. The customer side is unchanged by this framework.
For downstream traffic the retailer may shape the traffic according to the contract agreed with the wholesaler. For upstream traffic the customer's CPE may need to perform the corresponding shaping function. Alternately the facility rate could be configured to correspond to the service rate eliminating the requirement for shaping by the Network Terminal (NT).

Figure 18 shows the policing of downstream traffic located in the Access I/O. Alternatively, policing could be performed in a VNF or PNF similar to the location shown in Figure 19b.

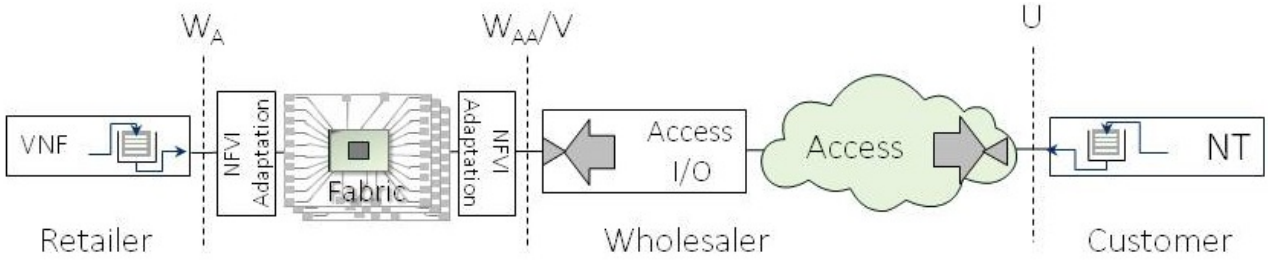

Figure 18: 1:1 access model with a locally hosted VNF

6.4.2.2 The n:1 Access Model

The n:1 model sees the wholesaler offer an n:1 service between the retailer and a specified set of customer service end points that are served by a common OLT or access node, see Figure 19. In MEF terms this is generally referred to as an Ethernet- TREE (E-TREE) service, although in this case it is topologically constrained. The n:1 model is considered to have several benefits, primarily relating to scale in the access I/O implementation, but also in possibly allowing some implementation “slack” in software-based shapers, and in permitting the retailer to manage oversubscription themselves.

Packets transmitted as part of an n:1 service are identified at the retailer side service end point by a TR-101 [2] native single S-tag or a virtual network identifier as defined by IETF NVO3. Note that the actual frames at the V-interface may be single or double tagged as per TR-101 [2], in the double tagged case the C-tag will be exposed to the VNF, while the S-tag might need interworking to the technology used inside the NFVI fabric. If the NFVI fabric is not native Ethernet, then the NFVI will interwork between the technology used within the NFVI fabric and the TR-101 [2] tagging in the access I/O. The n:1 service is identified at the customer side service end point as either a tagged or untagged UNI. Such a construct embodies a split horizon forwarding paradigm such that customers cannot communicate directly, and the traffic matrix is modelled as a bundle of individual P2P contracts.

For downstream traffic the retailer may first shape the traffic according to the contract agreed with the customer, then shape the aggregate of customers served by a common AN or OLT according to the contract agreed with the wholesaler prior to presentation to the service access point. For upstream traffic the customer’s CPE may perform the corresponding shaping function. The wholesaler polices the downstream aggregate at the service access point and the individual upstream tributaries at the customer facing service access points.

Figure 19 shows two locations where the wholesaler can police downstream traffic. In Figure 19a, the policer is shown within the Access I/O. This alternative provides for a policer within a defined PNF, but it constrains the topology of the E-TREE service since all customer service endpoints must be served by a common OLT or access node. Alternatively, Figure 19b shows the wholesaler’s policer implemented in a VNF or PNF. This configuration requires an extra network element, but it allows customer service endpoints to be served by any combination of OLTs or access nodes.

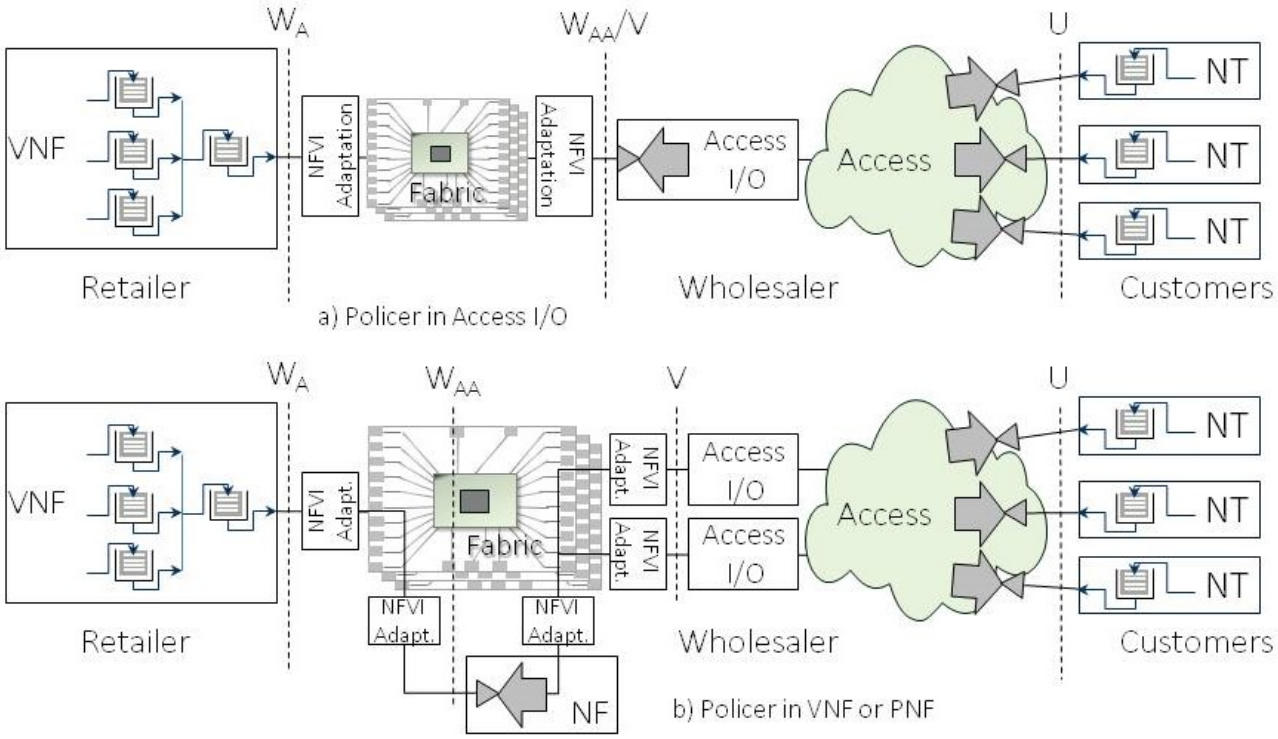

Figure 19: n:1 access model with a locally hosted VNF

6.4.2.3 WAN P2P & P2C

From a connectivity point of view, both WAN P2P and P2C models manifest themselves as a UNI. A private P2C model will interconnect a number of other P2C network service access points, see Figure 20. A P2C model may also involve control plane peering for routing exchange such that the customer prefixes reachable via the retailer VNF are advertised into the wider network. It is also possible to envision a scenario where the wholesaler service access point of the P2C UNI is implemented on other than the WAN I/O (e.g., a VNF).

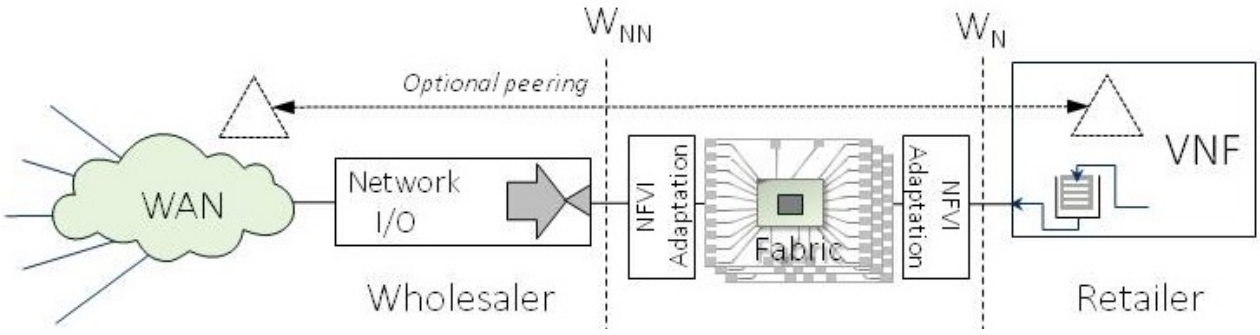

Figure 20: P2C WAN model

The WAN P2P model is simply a private connection where the endpoint is at some other point in the WAN, see Figure 21.

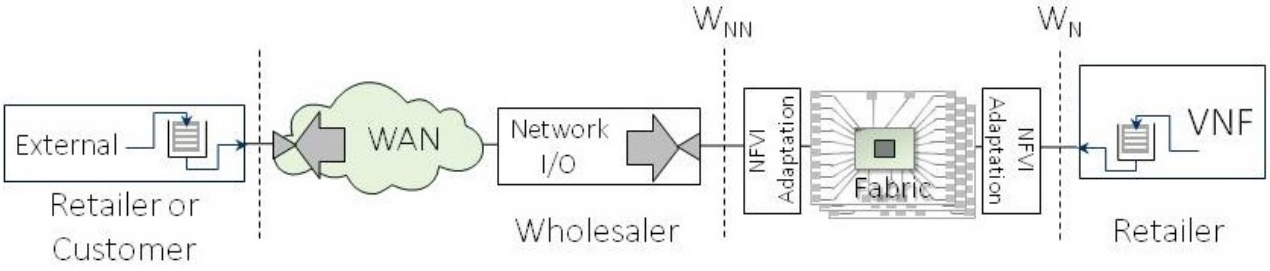

Figure 21: P2P WAN model

6.4.2.4 E2E transit

End to end transit involves traffic that passes through the fabric, but the UNIs are external to the fabric. The implication is that the resource reservation has been reconciled with the overall traffic matrix by the wholesaler to ensure all traffic contracts can be honored. It would be possible to do both 1:1 or n:1 models where the network side service access point is located external to the CloudCO; see Figure 22.

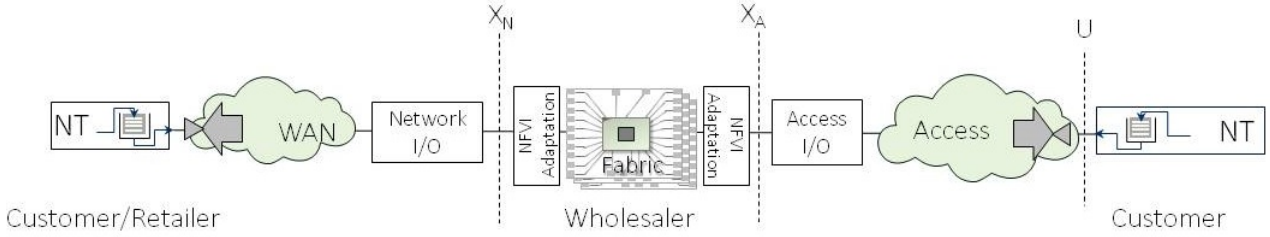

Figure 22: E2E transit

6.4.3 Implications of a Shared UNI

The presence of a fabric as a UNI on the retailer side of either access or WAN connectivity raises the possibility of having multiple service end points feed a single wholesaler service access point. This has implications both for user plane OAM and traffic management.
A potential consequence of this would be excessive packet drop as it would be difficult to coordinate a set of traffic shaping instances such that 100% of a traffic contract could be consumed. This can be avoided by how the retailer VNF system is architected. A non-exhaustive set of strategies would include:

1. A primary shaper, and all other retailer sources offer only minimal load.
2. Dynamic coordination of shaping entities at longer time scales.

It is not the intention of this framework to make recommendations as to how a shared UNI could work.
From an OAM point of view, what is normally a point to point service would now have a number of end points, while the tools for instrumentation of the path were primarily pairwise in nature. This has a couple of implications:

1. Instantaneous measurements of packet loss cannot be made. Only long-term correlation with a defined margin of error of the counts of a set of sources and single sink, or a single source with a set of sinks can be made.

2. To achieve fault detection in a defined time, it requires probes from all in the set of sources to the sink at a common frequency sufficient to detect failure amplifying the load on the sink, and a multicast probe from a single source to a set of sinks at the defined frequency.

6.4.4 Combining in and out of Profile Traffic

It is possible to consider modes of operation that contain a mix of traffic that is: (a) within a contracted profile and (b) purely best effort and not shaped to a profile. The best effort traffic is considered to be untrusted and simply marked to have a lower precedence than the "in-profile" traffic. In this scenario, the best effort traffic is assumed to employ adaptive flow management (e.g., Transmission Control Protocol (TCP) or Adaptive Bit Rate (ABR) as examples) to achieve a degree of fairness within the best effort traffic class but the QoS framework cannot exclusively depend on this for either fairness or suitable precedence. Therefore, this approach would still require policing of best effort traffic sources:

- Within the envelope of a defined service rate of an access contract where the traffic matrix may be the Σ of both shaped and untrusted best effort traffic.
- To have appropriate marking to avoid theft of service.

Such an approach provides the infrastructure owner with a suitable toolkit for supporting applications with a variety of network requirements.

6.5 Physical and Virtual Network Function Reconstruction

6.5.1 Distributed Routing Options

As noted in Section 5.2.5, BNG user plane L3 forwarding related network functions can be disaggregated and distributed across the CloudCO physical infrastructure. There are several ways of accomplishing:

1. The switch fabric is leveraged to create a CloudCO wide virtual distributed router (from a user plane perspective). This can be done by injecting the proper L3/ Layer 4 (L4) flows into the forwarding tables at run-time by the responsible SDN controller depending on the necessary service graph. PNFs attached to the switch fabric can also be leveraged in the same way. This needs to be done in a multi-tenant way, i.e., the ability to create many distributed router instances.
2. All the compute nodes are involved in creating a virtual distributed router (from a user plane perspective) inside the compute hypervisor virtual switch, again creating a CloudCO wide disaggregated BNG, offering the same default gateway to all the virtual workloads that run inside the NFVI, and offering routing across the fabric for network traffic. Packets that need to be sent across the switch fabric are encapsulated using e.g., network overlays (IETF NVO3), making the switch fabric configuration very static. VLAN-based access and VLAN-based hand-off to WAN circuits can be achieved by bridging the overlay into a VLAN, either on the hypervisor, or on the network switch, leveraging SDN Control.

6.5.2 User plane Programming as a Result of Disaggregation

As TR-178 compliant network nodes may be disaggregated, and their control plane functions are being virtualized, their user plane functions need to be programmable to allow certain packet types to be redirected to these control plane functions. The same applies to the switch fabric elements, as they are also in the forwarding path across the potential service chain.

6.5.3 Virtualization Options

Most of the control plane and management plane functions in the disaggregated BNG and AN are moved into the NFVI for virtualization. For a subset of subscribers, the option exists to completely virtualize a BNG (including user plane).

Some of the Residential Gateway (RG)/BNG functionality, including user plane, can be moved into the NFVI, as per TR-317 [8] and TR-328 [12].

6.5.4 Service Chaining

An example of service chaining is provided in Figure 4 of TR-345 [9], where the user plane forwarding model is shown. Each function in this forwarding model can be composed by two main functions:

- User plane;
- Control plane.

Service chaining can be achieved using IETF NVO3 techniques, i.e., network virtualization overlays. In this approach some network functions run as VMs or Containers, while network overlays perform basic L2/L3 steering of packets in-between and to certain NF's. The overlay functionality allows flexibility in creating the dynamic service chain as well as decoupling from the underlying switching hardware. The underlying hardware only needs to be capacity planned according to the throughput requirements between NF's and are not directly involved with service chaining themselves, which is a beneficial simplification.

6.5.5 Offering Functionality 'As-a-service'

CloudCO allows TR-317 [8]/TR-328 [12] style virtual gateways inside the NFVI, and these can be offered 'as-a-service', denoted as "subscriber-as-a-service", to other services, without the service consumer having to worry about e.g., scaling in/out, lifecycle management.

Multiple instances of these 'subscriber-as-a-service' services can be aggregated and combined with a multi-tenant Virtual Distributed Router Service (see Section 6.5.1), and that combination can be offered again 'as-a-service', denoted as IaaS, to a 3rd party consumer, making wholesale and 'as-a-service' offerings.

VAS can be onboarded, service chained and wrapped together with these ‘as-as-service’ building blocks to offer ‘VAS-as-a-service’.

6.5.6 Cloud Central Office Northbound API Description

The Northbound API data model on top of a CloudCO Domain has many different actors that engage with it, either through the API natively, or through a user interface that sits logically on top of the API. Note that the term ‘service’ used throughout this text can be a simple service like ‘allow-subscriber-onto-network’, or a composite service, i.e., a composition of different services to form a single service, e.g., internet access (combining the ‘allow-subscriber-onto-network’ with additional services that provide e.g., home-LAN DHCP/ Network Address Translation (NAT) services and a routable public IP address.). Every time a new service is created, a new API data model is built to allow the different actors access to their appropriate attributes of service functionalities. There needs to be an easy way for the actors to determine which objects they can read or change.

The different actors that can leverage the API data model are:

1. Service Consumer, i.e., the actors of the system that receive a service from the system. This can be to change service parameters or to monitor the service performance or monitor service parameters. Furthermore, the user interface generated for them will be specific to the objects they can manage or read.
2. Service Developers, i.e., the actors of the system who creates the service. Next to changing or reading service characteristics, the API data model needs to be able to accommodate scaling up/down the service, as well as onboarding (i.e., placement, Operating System details, etc.) the service, with the system hiding how scaling up/down is achieved.
3. Service Providers. The API data model can be leveraged to onboard a service on top of the CloudCO Domain, monitor it, change settings, create connectivity between services, etc. Further, the user interface generated for them will be specific to the objects they can manage or read.

Note that a service can be one or several instances of functionality running on a compute/storage resource, connected into a network or set of networks. These networks can be of different types and instantiated on physical or overlay user planes. Another point to consider is that the API data model should not assume the resource and/or network that a service is leveraging is running inside the CloudCO, in order allow the service to make potential use of IaaS offerings of 3rd parties. On the other hand, a service can also be composed of other already defined services. This introduces a notion of hierarchy, whereby a parent service can be combined from several building block child services. This means that this new parent service should be able to be consumed by its own dedicated and auto-created API data model. Furthermore, many different (from a state point of view), but identical (from a functional point of view) child services will need to be able to make use of a common parent service, in order to make this notion of service composition scale. In other words, isolation needs to occur between these child services to make sure that state from the different service users do not get mixed up, without the need to make this explicit via the API data model.

Furthermore, it is to be noted that in some cases the Operator running the CloudCO can be both a Service Developer (creating 'core' building-block services) as well as Service Provider (onboarding 3rd party, parent services and hooking them into their own child services).

6.5.7 Cloud Central Office Northbound API: Capabilities

The CloudCO NB API needs to support the following characteristics:

- The CloudCO Domain will expose a user interface to access certain API attributes.
- The CloudCO NB API will need to support different actors e.g., Service Consumer, Service Provider, Service Developer. Each actor will have different privileges in terms of access privileges and what parts of the API they can access.
- The CloudCO NB API needs to allow actors specific access to statistics and service attributes.
- The CloudCO NB API needs to automatically generate actor specific user interfaces.
- The CloudCO NB API will need to support role based access control.
- The CloudCO NB API needs to allow retrieving the necessary data model schemas for a given operation, depending on the actors, and the role within the actor identity, this is known as a declarative API and makes the API self-documenting.
- The CloudCO NB API needs to support 3rd party VNF onboarding and support general lifecycle management of a service, or components of a service.
    - Note that:
        - The term 'service' can mean an atomic service, or a composition of different atomic services. Note that a service can be one or several instances of functionality, i.e., VNFCs running on a compute/storage resource, connected into a (set of) network(s).
        - VNF onboarding describes the process by which a VNF is made available to a network functions virtualization (NFV) platform, allowing its life cycle operations, such as deployment, scaling, healing, software upgrade and termination, to be automated.
- The CloudCO NB API must be extensible to allow new services. As a result of developing and onboarding a new aggregated service, the appropriate API extensions have to be developed by the Service Developer, in order to allow lifecycle management of the service.
- The CloudCO NB API needs to be able to support scaling the service up or down, while the system hides how this scaling is achieved.
- The CloudCO NB API needs to be able to support composing services, out of various 'building-block services', where the actual 'building-block' service is running outside the CloudCO premise (e.g., 3rd party IaaS functionality).
- Several child services can make use of a common parent service, such that the child services and the instance of the parent service together compose a new service. In this case the CloudCO Domain needs be able to automatically create the API data models to allow the different actors to interact with the new composite service.
- The CloudCO NB API needs to support different child services derived from a common parent that are independent, i.e., the different child services are not cognizant of each other.

7 Dynamic Behavior of a CloudCO

This section describes the relevant CloudCO dynamic behaviors throughout its lifecycle which has typically a much longer horizon than the lifecycle of a VNF or other virtualized components.
Under this stand point, the CloudCO dynamic behaviors touch upon key phases for a CloudCO, such as (1) its bootstrap (somewhat comparable to the network creation phase in traditional networks), (2) the onboarding of VNFs (with due distinctions and caveats this bears some resemblances to service creation, though VNFs are a sub-part of a service chain) and (3) the exercise of performance monitoring, diagnostics and optimization (service assurance).

7.1 CloudCO Domain Bootstrap

This sub-section identifies a baseline reference for the CloudCO Domain dynamics from the beginning of its lifecycle with the creation of the instance itself to user plane VNFs instantiation and up to service bootstrap. For simplicity, this lifecycle does not consider legacy constraints.

The described CloudCO bootstrap dynamics are also helpful to identify the necessary CloudCO reference architecture building blocks and the basic functionalities required at the key system interfaces in order to support specific scenarios as those identified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios.
Furthermore, this baseline will be used as guidance for considerations about software and hardware reference implementations in future CloudCO sub-projects.

It is recognized that these baseline CloudCO bootstrap dynamics may need to be adapted or enhanced due to further technical investigation and use cases analysis.
This may apply, for example, in the two diametrically opposite cases below:

- Legacy: support of legacy elements in the user and/or management and control planes not easily integrated into the functional CloudCO architecture depicted in Figure 13.
- Evolutionary: support of CloudCO enhancements via autonomic networking or other forms of self-learning-based network and service operation.

Figure 23 describes a baseline CloudCO bootstrap dynamic based on the main characteristics of the CloudCO architecture.

- The CloudCO Domain Orchestrator consists of a SDN orchestration and control capability represented by the MCO Engine which cooperates with the included NFVO.
- The above mentioned SDN capability is expressed as a continuum of MCO tasks as well as CloudCO state transitions and state supervision.
- The MCO Engine, directly instructed via the CloudCO NB API, is responsible for the management and control continuum of the CloudCO Domain, by means of multiple application controller SBIs to PNF/VNF/DC controllers.
- The ensemble of CloudCO NFVI and MANO components hosts the execution and manages the lifecycle of the SDN controllers, therein implemented as VNFs.

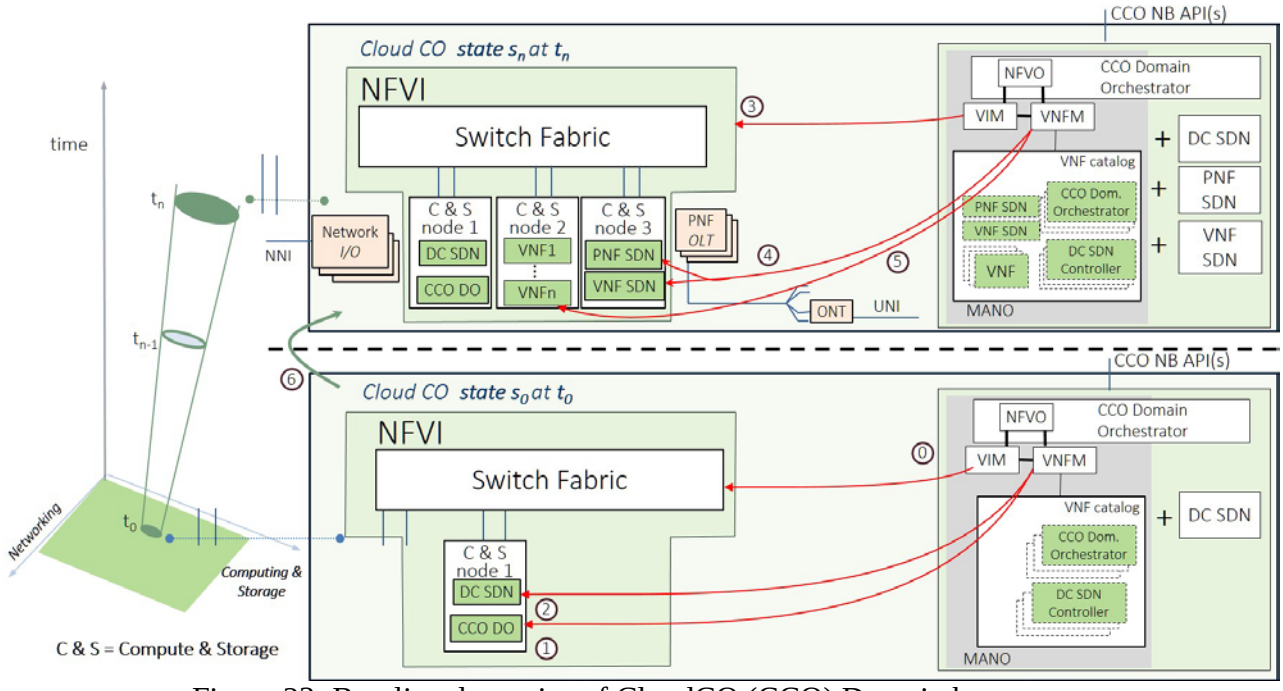

Figure 23: Baseline dynamics of CloudCO (CCO) Domain bootstrap

A pre-requisite for realizing the baseline dynamics of a CloudCO Domain bootstrap as depicted in Figure 23 is that, just before time t_0, an essential data center infrastructure, for instance consisting of a switch fabric and a minimal server farm, is up and running under the control of a MANO system.

An example of a CloudCO Domain bootstrap could be as follows:

- Step 0: bootstrap of the NFV Infrastructure is supervised by the VIM. It is assumed that NFVI, VIM and VNFM have been instantiated.
- Step 1: the VNFM, that comes along with a catalog of VNFs that can be instantiated as needed throughout the CloudCO lifecycle, instantiates the CloudCO Domain Orchestrator VNF.
- Step 2: the VNFM instantiates the DC SDN Manager&Controller VNF. This step is optional if the VIM is functionally adequate to handle all NFVI networking needs.
  Note that as indicated above the DC SDN and the CloudCO Domain Orchestrator are implemented as VNFs. This is not the only implementation option.
  At time t_0, the initial CloudCO Domain operating state (S_0) is reached and the described entities are steadily running on the NFVI server farm execution space:
- Step 3: Additional storage & Compute Nodes (2 and 3 in Figure 23) are added under VIM control.
- Step 4: the VNFM instantiates the PNF and VNF SDN Managers&Controllers VNFs that take care of management and control respectively of the Access Nodes and Network I/Os and of the VNFs that realize the CloudCO user plane.
  Note that as indicated above the PNF and VNF SDN controllers are implemented as VNFs. This is not the only implementation option.

- Step 5: the VNFM instantiates VNFs (e.g., virtual BNG (vBNG), virtual Evolved Packet Core (vEPC), etc.) as needed to realize the CloudCO user plane per the initial network status needed to accept service requests.
  Note that this step may go somewhat beyond the basic bootstrap of the CloudCO Domain, depending on the instantiated user plane VNFs, and may extend to the use cases described in the following sections.
  Note that the VNFM function to bootstrap the CloudCO Management & Control elements could be implemented as a script.
  At time t_n, state S_n is reached: the CloudCO Domain operates as a specialized hybrid physical and NFV Infrastructure domain, where NFVI/MANO and Access/Edge SDN capabilities are exploited in a coordinated fashion as guaranteed by the interaction between the MCO Engine and the NFVO within the CloudCO Domain Orchestrator.
  From time t_n on, requests for Network-as-a-Service and end-user services can be issued and consumed through the CloudCO NB API(s).
- Step 6: (overarching step before/from t_0 to t_n): the CloudCO Controller & Optimizer computes and updates the current CloudCO state based on the overall resource status. Based on the inputs from the E2E Service Orchestrator, the MCO Engine, in coordination with the CloudCO Controller & Optimizer, decides on CloudCO state transitions and converts them into specific actions implemented via the CloudCO Domain Orchestrator interfaces.

7.2 Dynamic Behavior of CloudCO Interfaces

The following chapter describes a set of features that are expected to be supported across the CloudCO interfaces, based on the use cases described in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios.

This set is not meant to be complete but will help in setting expectations for future work. VNFM interfaces (OR-vnfm, Ve-Vvnfm-em, Vi-Vnfm, Ve-Vnfm-vnf) are not discussed in this document as they depend on the applications/services deployed on the infrastructure and are typically not involved with service chaining.

7.2.1 O_s-$M_{a\text{-}ccodo}$ (CloudCO Northbound API)

This interface is already described in detail in Sections 6.5.6 and 6.5.7. The attributes are referred to specific to different actors. One is the Service Developer, who is also responsible for authoring the relevant API's for a given developed Service. There is the Service Provider, who can use the API to Create, Read, Update and Delete (CRUD) the service. And finally there is the Service User/Consumer who has often only the ability to read or update the service parameters.
Based on the use cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable:

- A Service Provider to exercise CRUD operations on a new instance of a pre-created service type. Note that not all of the state that can be read/created needs to reside on the CloudCO DO, but could equally be kept into lower level functional blocks such as a Service-specific VNFM.

- A hierarchy of Service Providers, i.e., wholesale versus retail, whereby the retail Provider acts as a consumer towards the wholesale provider, while acting like a provider to a Service User.
- A Service Provider to assign a Service User to service instances.
- The creation of multiple instances of multiple services concurrently, protected by the appropriate role-based access control (RBAC), and the necessary creation of multi-tenant services.
- The reading of certain diagnostics and performance data of the service.
- The Receiving abstract optimization triggers to dynamically adapt the system.
- Instantiation of monitoring and diagnostic facilities.
- A Service Developer to on-board new service components, and wrap them into an Aggregate Service using existing services and/or these new service components. Note that there is an assumption of a set of initial building-block services which need to enable:
    - Instantiation of SDN Applications.
    - Redirection of certain traffic types (such as 802.1x, DHCP) towards a given SDN Application.
    - Keeping track of an access-line AAA state as a result of received traffic on the access-line.
    - The creation of an IP forwarding Service across the CloudCO Domain.
    - Combination of several access-lines with one IP forwarding Service, taking into account the current access-line state in the SDN Applications.
    - To Service Chain other VNFs into any service.

7.2.2 Or-Vi

This is the interface between the Orchestrator and the NFVI VIM, as such it will be leveraged to set up and service chain VNFs. Based on the use cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable, but is not limited to:

- Creation of Virtual Networks across all the Compute Nodes in the NFVI.
- Creation of Virtual Distributed Routers across all the Compute Nodes in the NFVI.
- Creation of Virtual Routers to route between VLANs on Compute Node Uplinks and Virtual Networks.
- Creation of Bridges between VLANs on compute Node Uplinks and Virtual Networks.
- Association of subnets to Virtual Networks.
- Attaching Virtual Networks to Virtual Distributed Router Logical Interfaces.
- Offering IP address management to workloads on these Virtual Networks.
- Configuration of Security attributes for these networks (Distributed Firewalling).
- Configuration of Security attributes on Virtual Routers (Firewalling and NAT).
- Configuration of Load Balancing attributes on Virtual Routers.
- Attaching workloads to Virtual Networks.
- Extending Virtual Networks to switches in the switch fabric (so called Hardware Virtual eXtensible LAN Tunnel End Point [VTEP]),
- Doing a variety of CRUD operations for workloads, such as :
    - Create workload.

    - Set metadata and other variables such as network addresses, GuestOS, etc.
    - Move/Migrate workloads.
    - Delete workloads.
    - Start, Stop, Backup/Snapshot, Reboot, pause, resize workloads.
- Configuration and enforcement of Resource Management policies for the workloads, in a multitenant fashion.
- Configuration of QoS policies and enforcement for the virtual networks.
- Offering multitenant access to virtual networking and virtual compute across a shared NFVI.

7.2.3 O_{cco}-$Nf_{-sdn-pnf}$

This interface is between the CloudCO Domain Orchestrator and the PNF controller(s). Based on the Use Cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable:

- The CloudCO Domain Orchestrator to install and configure SDN Applications logically on the PNF SDN Controller. Note that the controller itself could be a virtual appliance itself.
- Signaling the appropriate information originating from the SDN applications upstream towards the CloudCO Domain Orchestrator.
- Doing FCAPS for any PNF.

7.2.4 O_{cco}-$Nf_{-sdn-vnf}$

This is the interface between the CloudCO Domain Orchestrator and the VNF controller(s). Based on the Use Cases specified in the BBF Technical Report focusing on the CloudCO use cases and Scenarios, the interface needs to enable:

- The CloudCO Domain Orchestrator to install and configure SDN Applications logically on the VNF SDN Controller. Note that the controller itself could be a virtual appliance itself.
- Signaling the appropriate information originating from the SDN applications upstream towards the CloudCO Domain Orchestrator.
- Doing FCAPs for any VNF.

7.2.5 $M_{fc-sdn-dc}$-N_f

This is the interface between the CloudCO Domain Orchestrator and the physical fabric of the NFVI. Based on the Use Cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable:

- Signaling the appropriate information originating from the SDN applications upstream towards the CloudCO Domain Orchestrator.
- Configuration and/or manipulation of flow tables inside the Switch Fabric (e.g., configure VLANs or install L2/L3 Flows).

7.2.6 M_{inf}/M_s

These are the configuration interfaces between the SDN Controllers and the PNFs or VNFs respectively. Based on the Use Cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable:

- Configuration and reading of PNFs and VNFs attributes (e.g., configure VLANs).

7.2.7 Mfc

This is the Flow Control interface between the SDN Controllers and the PNFs or VNFs. Based on the use cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios, the interface needs to enable:

- Configuration and/or manipulation of flow tables inside the PNFs and VNFs.
- Reception of redirected user plane packets from PNFs upstream to the SDN Applications and configuration of flow tables to appropriately route those packets through the user plane.
- Receptions and sending of Flow Control Messages and ability to instantiate Flow Control Rules.

8 Next Steps

8.1 CloudCO Application Notes

Using the Architectural Framework described in this document, a Service Developer/Provider can now create 'CloudCO Application Notes'. These CCO Application Notes describe:

- A Use Case for service creation on the CloudCO infrastructure.
- As such it describes the actors of the system.
- It describes what application components need to be onboarded on the CloudCO SDN and NFVI infrastructure to make the Application Note work.
- It describes the resulting Northbound API expectations in enough detail such that all service consumption can occur by the various actors.
- It describes the necessary interface level interactions between the various functional elements to make the use case work, in detail.

In essence, the Application Notes are a more detailed version of the use cases specified in the BBF Technical Report focusing on the CloudCO Use Cases and Scenarios. In particular, these use cases can be used as a source of inspiration.

A template will be provided by the Broadband Forum.

8.2 CloudCO Interface descriptions

An Application Note will lead to a detailed description of how interfaces between CloudCO functional elements need to function and what attributes need to be exchanged to make the Application Note work. From this future work will need to choose:

- Which transport protocol will be chosen to deliver those attributes.
- What information elements will be used to carry the attribute elements.

The interface descriptions and protocol choice only needs to fulfill the requirements for one given Application Note. As such the interface descriptions can grow over time (other information elements can be added, as the amount of application notes grows.

Interface descriptions as an output of Application Notes will be documented in future work.

8.3 CloudCO Test Cases

Using the Application Notes and Interface descriptions as an input, test cases will be derived. These test cases will describe how to create the services explained in the Application Notes, and also what functionality needs to be loaded onto the CloudCO infrastructure. It will therefore also test the attribute exchanges across the CloudCO interfaces.

Broadband Forum will not test the test cases, but will outsource this work to Open Broadband Labs. It is crucial that the Application Notes are detailed enough such that Open Broadband Labs can onboard the necessary pieces of software to create a working end to end system.

End of Broadband Forum Technical Report TR-384