TECHNICAL REPORT

TR-408

Cloud CO Migration and Coexistence

Issue: 1

Issue Date: September 2020

Notice

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

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT AND ANY IMPLIED WARRANTIES ARE EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE USER’S OR IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY USER, IMPLEMENTER, OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT, INCLUDING BUT NOT LIMITED TO, ANY CONSEQUENTIAL, SPECIAL, PUNITIVE, INCIDENTAL, AND INDIRECT DAMAGES.

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.

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	22 September 2020	22 September 2020	Ding Hai ChinaUnicom	Original

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

Editor: Ding Hai, China Unicom
Tim Carey, Nokia

Work Area Director(s): George Dobrowski, Morris Creek Consulting
Bruno Cornaglia, Vodafone

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

Table of Contents

List of Figures

List of Tables

Executive Summary

This Cloud Central Office (Cloud CO) document provides guidance to Service Providers as they transform their existing Multi Service Broadband Networks (MSBN) toward a Cloud CO based network that is more adaptable, agile, scalable and dynamic than the existing MSBN, while reducing costs by retaining investments that the Service Provider has made in their MSBN. In this context, this document identifies considerations that Service Providers need to consider when determining the deployment strategy of either migrating the existing MSBN or deploying a Cloud CO based network that coexists with the existing MSBN. The document also provides transformation assistance by identifying approaches to migrating their existing MSBN and discusses when and how to introduce the Cloud CO infrastructure or leverage their existing IT infrastructure. Finally, the document provides examples of how functionality in the management, control and user planes can be migrated toward Cloud CO.

1 Purpose and Scope

1.1 Purpose

TR-384 [3] defines a reference architecture for a Cloud CO (CCO), including the decomposition and virtualization of functions which were previously hosted in network elements deployed within the Service Provider's network. It also describes the behavior and interactions of functional modules as well as a new management architecture.". The target Cloud CO architecture is fundamentally different in many ways compared to existing broadband networks defined by TR-101 [1] or TR-178 [2]. There is a difference in the functional partitioning as CCO doesn't require a specific nodal placement of functions, as well as their distribution caused by nodal decomposition and disaggregation. A byproduct of functional partitioning is that the Cloud CO interfaces may also be different, including newly exposed interfaces for the decomposed functions as well as significant changes in the management infrastructure and interfaces. As a result, The Cloud CO architecture may require planning that includes standardization activities, which is the purpose of the work in BBFs' Cloud CO project stream.

One possibility for Service Providers is to introduce Cloud CO in new deployments, AKA "Greenfield". However, for most Service Providers, Greenfield deployments covers only a small fraction of their market, and it may be harder to justify the transition to a new networking and operation paradigm only for Greenfield deployments. Service Providers globally have made huge investments in networks providing varying broadband services. In fact, some of networks that are termed "legacy" or "existing", were deployed very recently. It may therefore be important to identify methods for Cloud CO transformation of existing MSBN while still allowing for the continued realization of some of the investments made in the existing MSBN.

The purpose of this Technical Report is to suggest methods by which an existing MSBN can gradually migrate to a Cloud CO one, and / or coexistence of the existing MSBN and Cloud CO architecture within segments of the same network.

1.2 Scope

The choice of how to transition to a Cloud CO architecture may differ. The main factors for deciding on which method may include the objectives of the transformation, the existing network situation, the speed and scale of the desired transition and more. This Technical Report identifies a few transformation use cases that may be typical and might cover the majority of situations Service Providers can encounter during the transition of their networks toward the Cloud CO architecture. It then specifies the steps required to realize each of those use cases.

The details of the Cloud CO architecture, such as its components, interfaces and behavior are defined by other specifications within the Broadband Forum's Cloud CO family of specifications. While it is not the purpose of this Technical Report to specify such details, it may influence other Cloud CO specifications.

1.2.1 Collaboration between Open Source and Standardization

Since the initiation of the Cloud CO Project, the Broadband Forum recognized the need to deal with this highly transformational architecture with the development of a richer ensemble of deliverables and a methodology aiming at blending Open Source practices with some key fundamentals of Standards bodies.

For this reason the Cloud CO Project encompasses different types of deliverables:

- Reference Architecture TR-384
- Interfaces Specifications
- Software Reference Implementations

- Coexistence and Migration (this Technical Report)
- Exemplary Implementations and Testing

Early in the development process of the TR-384 Cloud CO Reference Architecture it was realized that a new approach to the development of use cases would also be necessary. This new process places an emphasis on the creation of Cloud CO Application Notes (AppNotes) to capture specific use cases of the Cloud CO architecture. In common practice, a Cloud CO Application Note covers a well-focused use case, providing additional details, such as the interface between specific components, or the inclusion of various open source implementations in the use case's design or implementation.

The Cloud CO Application Note development process provides a review and commenting framework, prior to the publication of the application note on the BBF website as an informative resource as compared to the normative specifications of the interfaces and information models for the Cloud CO Reference Architecture, such as TR-413 "SDN Management and Control Interfaces for CloudCO Network Functions" [5].
The development of Application Notes provides guides and rapid development on the innovation of interface and information model definitions, which then follow the conventional review and approval process.

Figure 1 below illustrates the Cloud CO related deliverables.

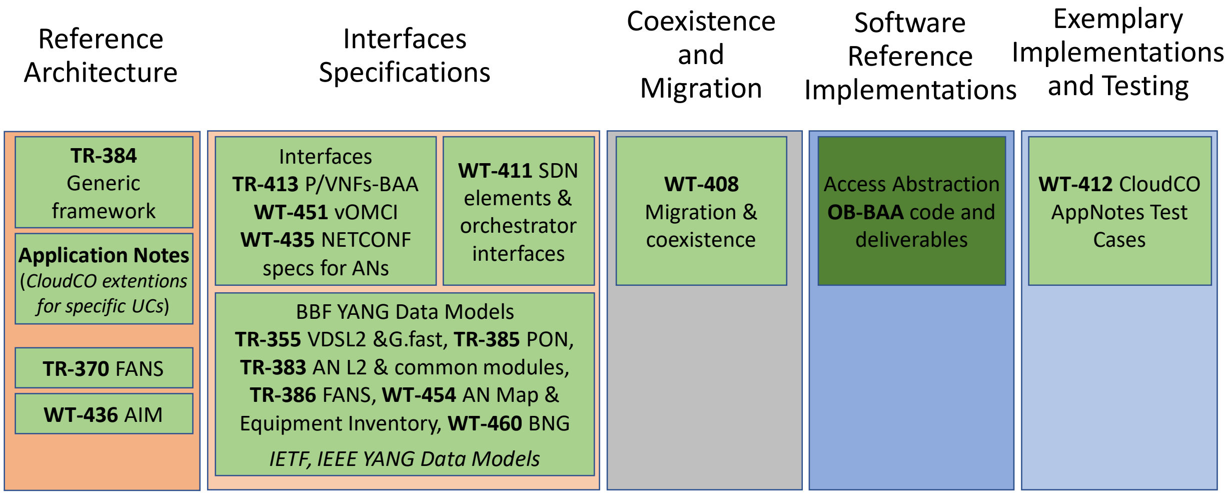
Reference Architecture
Interfaces Specifications
Coexistence and Migration
Software Reference Implementations
Exemplary Implementations and Testing
TR-384 Generic framework
Application Notes (CloudCO extentions for specific UCs)
TR-370 FANS
WT-436 AIM
Interfaces TR-413 P/VNFs-BAA WT-451 vOMCI WT-435 NETCONF specs for ANs
WT-411 SDN elements & orchestrator interfaces
BBF YANG Data Models TR-355 VDSL2 &G.fast, TR-385 PON, TR-383 AN L2 & common modules, TR-386 FANS, WT-454 AN Map & Equipment Inventory, WT-460 BNG
IETF, IEEE YANG Data Models
WT-408 Migration & coexistence
Access Abstraction OB-BAA code and deliverables
WT-412 CloudCO AppNotes Test Cases

Figure 1: Cloud CO Project deliverables

To enable the rapid process and prototyping of application notes, the Broadband Forum's Open Broadband efforts provide open source reference implementations for some of the components used within the Application Notes.

These Open Source projects are organized and overseen by the Broadband Forum and provide direct feedback into both the Application Note and Specifications processes. At the same time the reference implementations of these Open Source projects are the natural transition point with the lab and testing activities of Open Broadband as shown in Figure 2.

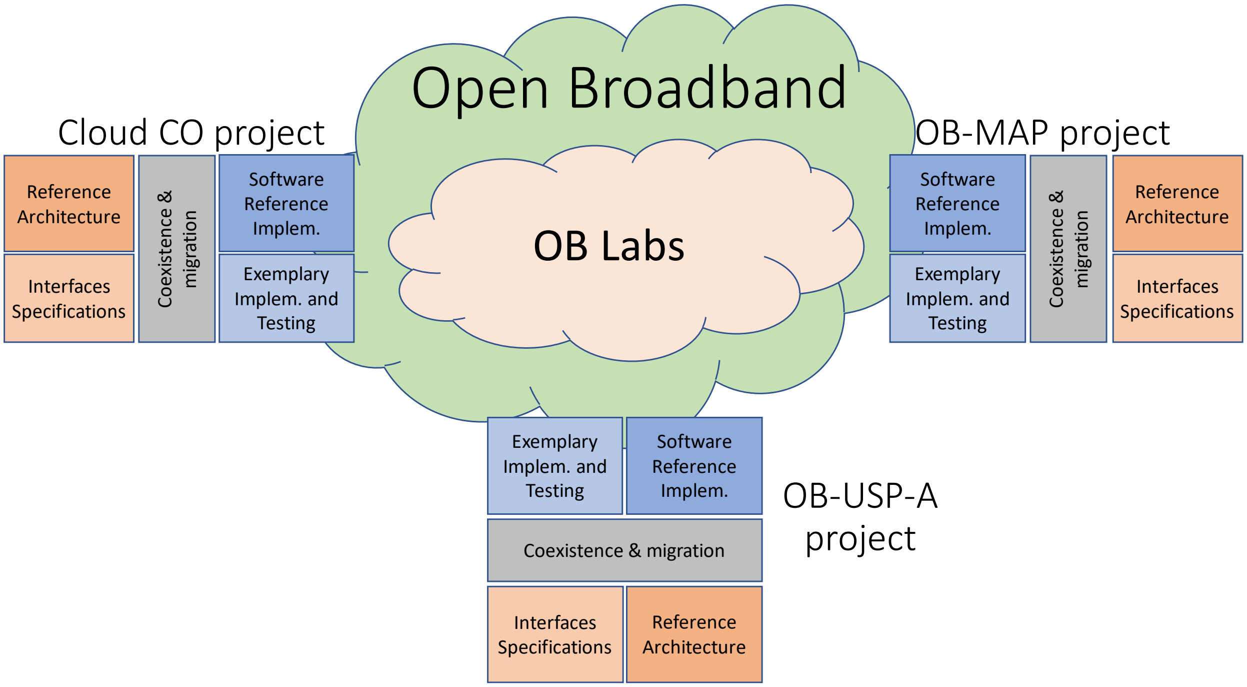
Open Broadband
OB Labs
Cloud CO project
Reference Architecture
Interfaces Specifications
Coexistence & migration
Software Reference Implem.
Exemplary Implem. and Testing
OB-MAP project
Software Reference Implem.
Exemplary Implem. and Testing
Coexistence & migration
Reference Architecture
Interfaces Specifications
Exemplary Implem. and Testing
Software Reference Implem.
OB-USP-A project
Coexistence & migration
Interfaces Specifications
Reference Architecture

Figure 2: Open Broadband and BBF Open Source projects

Open Broadband Labs work closely with the Open Broadband projects and BBF members to develop and deploy application notes, creating a specific instance of an Application Note, with additional details (e.g. hardware, software, and deployment choices), test plans, and testing results. These test results, when possible are provided back to the open broadband projects and the Broadband Forum to complete the development / deploy / test cycle.

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119.

MUST This word, or the term "REQUIRED", means that the definition is an absolute requirement of the specification.

MUST NOT This phrase means that the definition is an absolute prohibition of the specification.

SHOULD This word, or the term "RECOMMENDED", means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.

SHOULD NOT This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.

MAY This word, or the term "OPTIONAL", means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

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]	TR-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[2]	TR-178 Issue 2	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2017
[3]	TR-384	*Cloud Central Office (CloudCO) Reference Architectural Framework*	BBF	2018
[4]	TR-317	*Network Enhanced Residential Gateway*	BBF	2016
[5]	TR-413	*SDN Management and Control Interfaces for CloudCO Network Functions*	BBF	2018

2.3 Definitions

Cloud CO Cloud Central Office

Coexistence A deployment model where Cloud CO (CCO) infrastructure is deployed alongside the existing MSBN infrastructure inside the CO. Subscribers can be attached to one or both of the infrastructure types.

Migration The process by which an existing network architecture is converted to a CCO architecture. Each migration step may involve conversion or replacement of existing hardware, software, or introduction of new services, infrastructure and functions. Various migration types can be considered.

Service Enhancement The offering of a new service or enhancement of existing service by CCO. Service enhancement may happen either in coexistence or migration scenario.

SDN M&C Software Defined Networking Management and Control

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AAA	Authentication, Authorization and Accounting
AN	Access Node
AppNote	Application Note
AR	Augmented Reality
BAA	Broadband Access Abstraction
BBF	Broadband Forum
BNG	Broadband Network Gateway
BRAS	Broadband Remote Access Server
BRG	Bridged Residential Gateway
BUM	Broadcast, Unknown Unicast, and Multicast
CCO	Cloud CO
CCDO	CCO Domain Orchestrator
CO	Central Office
CDN	Content Delivery Network
CPE	Customer Premise Equipment
DC	Data Center
DHCP	Dynamic Host Configuration Protocol
EM	Element Management
EMS	Element Management System
FW	Firewall
IETF	Internet Engineering Task Force
I/O	Input/Output
IP	Internet Protocol
IT	Information Technology
LAN	Local Area Network
LSL	Logical Subscriber Link

L2	Layer 2
MANO	Management and Network Orchestration
MEC	Mobile Edge Compute
MSBN	Multi Service Broadband Network
NAT	Network Address Translation
NERG	Network Enhanced Residential Gateway
NETCONF	Network Configuration (management protocol)
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NMS	Network Management System
OLT	Optical Line Termination
OSS	Operational Support System
PNF	Physical Network Function
RG	Residential Gateway
SDN	Software Defined Network
TR	Technical Report
UAM	Universal Access Manager
VAS	Value Added Service
VLAN	Virtual Local Area Network
vCDN	Virtual CDN
vCPE	Virtual CPE
vDHCP	Virtual DHCP
vFW	Virtual FW
vG	Virtual Gateway
VIM	Virtualized Infrastructure Manager
vNAT	Virtual NAT
VNF	Virtualized Network Function
vOLT	Virtual OLT
vUP	Virtual user plane
VR	Virtual Reality
WA	Work Area
YANG	Yet Another Next Generation (data modeling language)

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, HVAC and fire protection between traditional Central Offices and datacenters is out-of-scope for this document.

3.2 Security

Security provides "a form of protection where a separation is created between the assets and the threat." This Technical Report enables the sharing of a common infrastructure and interfaces between various use cases that may be operated by different departments (e.g., wireline and mobile) or different tenants (other Service Providers). This Technical Report also provides an increased opportunity for Service Providers to dynamically control the network service behavior, with the use of API's.

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

3.3 Privacy

A multi-tenant Cloud CO hosts functionality for a set of actors with potentially competing interests that the Cloud CO 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 Cloud CO is a system of sufficient complexity that it will expose new attack vectors to malicious parties that have access to the Cloud CO 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 in transit and at rest, so that it cannot be understood even if it is intercepted.

This document does not define any specific mechanisms.

4 Overview of Cloud CO Migration and Coexistence

The transformation of today's MSBN toward a cloud based Central Office goes beyond the infrastructure and network components and represents a multi-faceted process across a Network Operator/Service Provider organization.

Focusing on the sole network transformation activities two major strategies, applicable to the Management, Control and User (Data) Planes, can be identified:

- Migration, which consists in the replacement of existing network components with new Cloud CO (CCO) components or their upgrade (if possible) to make them compliant with the CCO deployment.
- Coexistence of the existing network with deployed CCO components either in overlay or via interworking.

The above strategies need to consider and are conditioned by CCO's enabling technologies, e.g., virtualization, SDN automation, distributed network capabilities and service applications.

Migration activities can be classified as the following types:

- **Function migration**: replacing/upgrading a given network function of existing nodes with an equivalent function hosted on the CCO infrastructure.
- **Node migration**: replacing/upgrading an MSBN node, including all of its functions, with corresponding functionality hosted on the CCO infrastructure.
- **Subscriber migration**: the process of moving one or more subscribers from today's MSBN to the CCO infrastructure. Subscriber Migration would typically occur in a coexistence scenario.
- **Service migration**: the process of replacing a service offering from a today's MSBN with equivalent functionality hosted on the CCO infrastructure, including all the subscribers of that service. Service Migration would typically occur in a coexistence scenario.

Migration activities, such as a Function and Node migration, are performed by disaggregating the functionality that was provided by MSBN network elements and moving one or more of these functions to be hosted on the CCO infrastructure. The considerations for disaggregating network elements are highlighted in section 4.1 of this Technical Report.

In addition, the following sections of this Technical Report additionally describe the dimensions, the types of elements that are affected during the transformation activity (e.g., subscribers, functions).
These sections are informative as migration strategies and associated coexistence timeframes do not occur in isolation of other activities, are conditioned by multiple non-technological factors and vary across Service Providers.

The Cloud CO's flexibility in allocating network functions on the NFV Infrastructure and supporting technologies needed to adapt the existing MSBN network functions makes the Cloud CO a very versatile platform that builds on a hybrid physical/virtual infrastructure, a major paradigm shift with respect to the existing MSBN.

An integral function at the heart of the Cloud CO is the Domain Orchestrator that operates over distributed SDN domains (e.g., Access, Edge) relying on the resources and capabilities of an NFV Infrastructure. Under this standpoint the Cloud CO Domain Orchestrator is envisaged as a "brain" with two coordinated "hemispheres" (i.e., MCO Engine, NFVO). The tight integration of these "hemispheres" guarantees effective operation of the SDN Management and Control Plane elements while the NFV components ensure that the User Plane network functions evolve in space and time as required by the Cloud CO state and configuration transitions.

Effective SDN Management and Control functionality that is coordinated with NFVI capabilities can be applied to micro-transitions related to Cloud CO's daily operations and is a powerful method for implementing and tuning macro-transitions like CAPEX/OPEX driven migrations (e.g., Node or Function migration), revenue driven migrations (e.g., Service or Subscriber migrations) or a mix of both.

4.1 Considerations on Migration / Coexistence

The Service Providers' existing MSBN is expected to gradually migrate to a network deployment that is represented by the CCO architecture, and/or today's MSBN and CCO deployments can coexist for a significant time period. For today's MSBN to either migrate to the CCO architecture or ultimately coexist with CCO deployments, several migration phases can be required.

The following factors should be considered when migrating the today's MSBN toward the CCO architecture or when today's MSBN coexists with the CCO:

1. Existing network situation: Service Providers globally have made huge investments their existing MSBN, it is important to identify methods for CCO transformation of today's MSBN while protecting some of those investments.
2. Migration for increased service agility: turning the existing MSBN into an agile and automated one capable of introducing new services/features rapidly.
3. Migration and coexistence based on reduced total cost of operations: The decision to migrate functionality within the existing MSBN needs to consider the total cost of operations that includes aspects such as increased system integration costs to systems that interact (e.g., OSS, NMS) with the management and control elements of the CCO architecture; additional capital expenditures; software licensing internal training and organizational re-structuring.

In addition to the factors that a Service Provider would take into consideration upon selecting a CCO deployment strategy, Table 1 identifies advantages and disadvantages of choosing a migration or coexistence strategy.

	Migration		Coexistence
+	Can migrate functions gradually as their virtualized version become available	-	Requires full readiness of virtualized functions to get started
+	Some independence of management migration may be possible (e.g. via the CCO BAA layer)	-	Requires readiness of CCO management plane
+	Existing network elements can be phased out in stages	-	Requires all 'existing' equipment types until all subscribers are migrated
-	Require upgrades to 'existing' equipment to direct selected traffic to the NFVI fabric	+	'existing' equipment remains unmodified through subscriber migration
-	May require temporary hooks to abstract / hide 'existing'/CCO differences	+	CCO remains “Greenfield Clean” – No such hooks needed
+	CCO NFVI scales by functionality	+	CCO NFVI scales by capacity
-	Higher Risk: one function changes for all subscribers (probably)	+	Lower Risk: down to one subscriber at a time

Table 1: Comparison of Migration/Coexistence Strategies

4.2 Transforming the Management Plane

In many strategies, the management plane is the first plane to be transformed within the existing MSBN. Management Plane transformation is necessary in order to offset increases in operational costs Service Providers have seen when they introduce virtualization into the MSBN.

When the management plane is transformed from network management elements with the vertical view of resources and custom integration toward domain controllers and associated orchestration functionality that is based on standards, the result is an SDN Management and Control layer that enables automated, programmable and always-on network functionality needed for SDN automation.

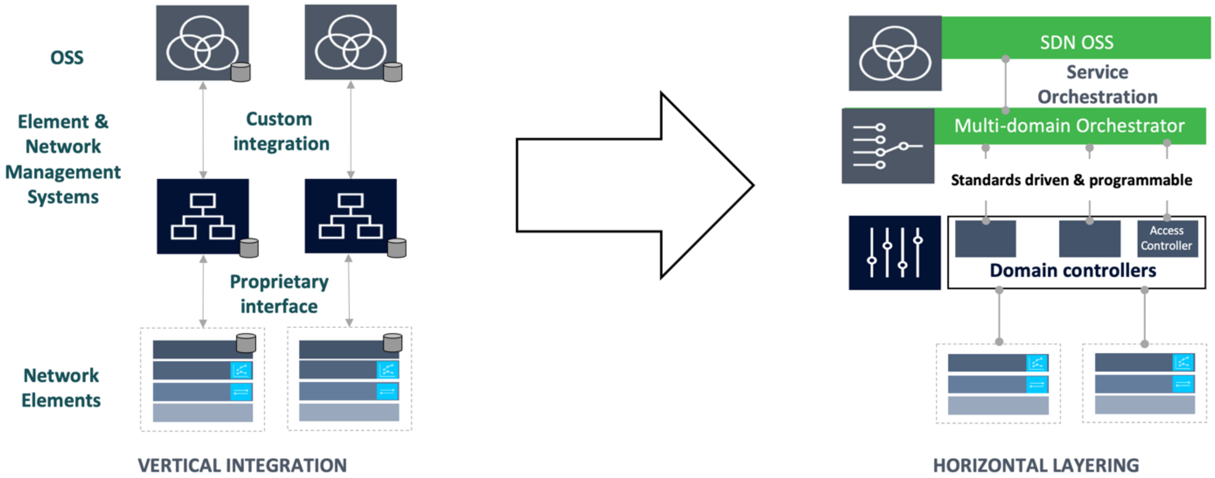

Figure 3: Network Management to SDN Domain Controller Transformation

4.2.1 Migration of the Management Plane

The migration of the management plane suggests a path for existing element management system functionality to be replaced by or integrated into SDN Manager and Controllers. In this section of the Technical Report, a possible path for migrating the management plane is described.

In the existing network, element management (EM) functions manage the respective domain’s (e.g., Edge, Access, Premises) network elements (e.g., BNG, OLT, RG) as depicted in the Figure 4. The High layer OSS and NMS systems interact with the element management systems to provide cross-domain and in some cases cross-vendor management functionality.

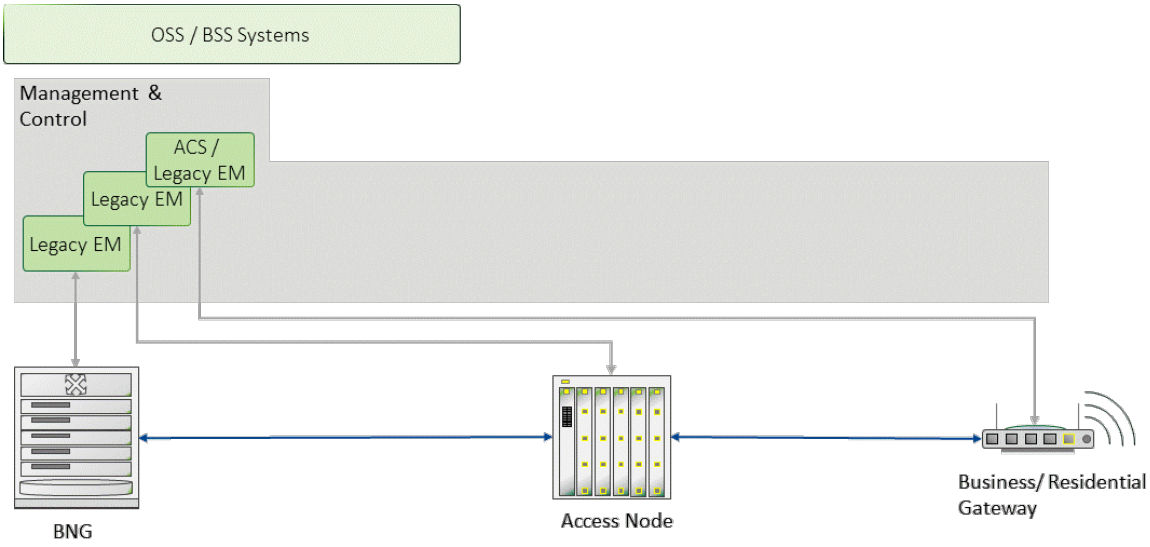

Figure 4: Management plane transformation with existing management functionality

When transforming the existing network management infrastructure toward the SDN based management and control elements, SDN management and control elements can be introduced and work in conjunction with the existing network management infrastructure. Depending on the Service Provider's preference the SDN management and control element can manage and control multiple domains of the broadband network as depicted in Figure 5 or there can be SDN management and control elements that are specialized to provide control of specific domains (e.g., Edge, Access) as depicted in Figure 6.

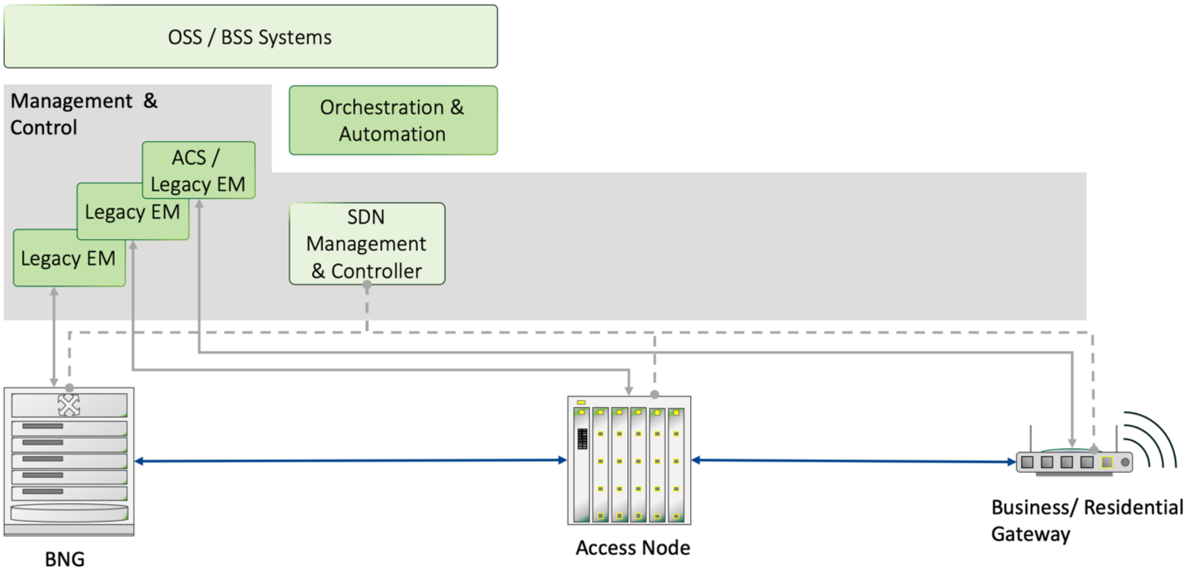

Figure 5: Management plane transformation with existing and SDN Management functionality (control of multiple domains)

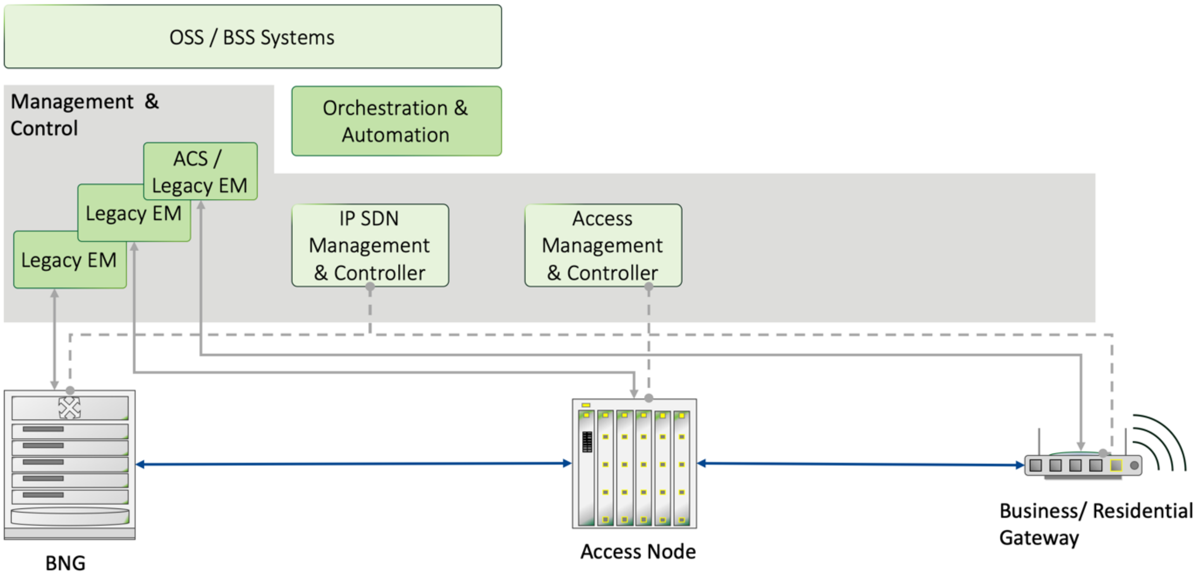
OSS / BSS Systems
Management & Control
Orchestration & Automation
ACS / Legacy EM
Legacy EM
Legacy EM
IP SDN Management & Controller
Access Management & Controller
BNG
Access Node
Business/ Residential Gateway

Figure 6: Management plane transformation with existing and SDN Management functionality (control of specific domains)

Once the SDN management and control elements are in place the SDN management and control element to maintain the functionality and the associated service function chain for any virtualized MSBN control or user plane functions is realized as depicted in Figure 7. At this point the existing element management functions might no longer be necessary in the solution.

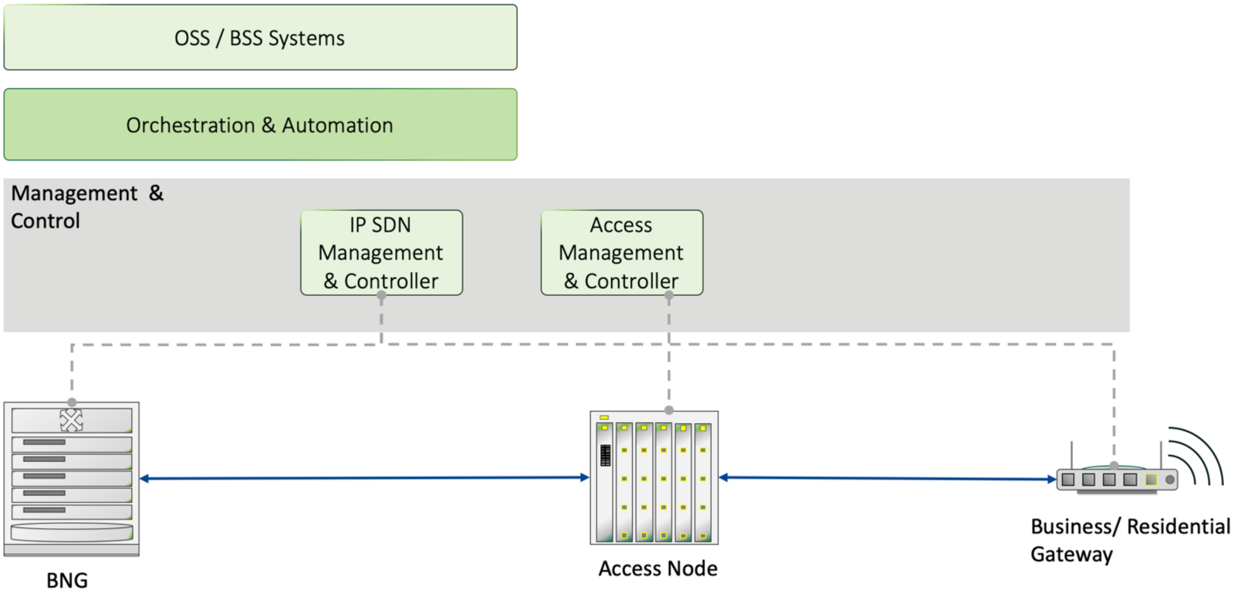
OSS / BSS Systems
Orchestration & Automation
Management & Control
IP SDN Management & Controller
Access Management & Controller
BNG
Access Node
Business/ Residential Gateway

Figure 7: Network Management transformation with SDN Management and Control

4.2.1.1 Geographical Distribution of the Management Plane

A key decision when migrating the functions of the management plane is the location of the functions. In many cases the existing functionality of the management plane is located either in the Service Provider's central or regional DCs. However, in some cases management functions of a specific technology can be in the Service Provider's CO as depicted in Figure 8.

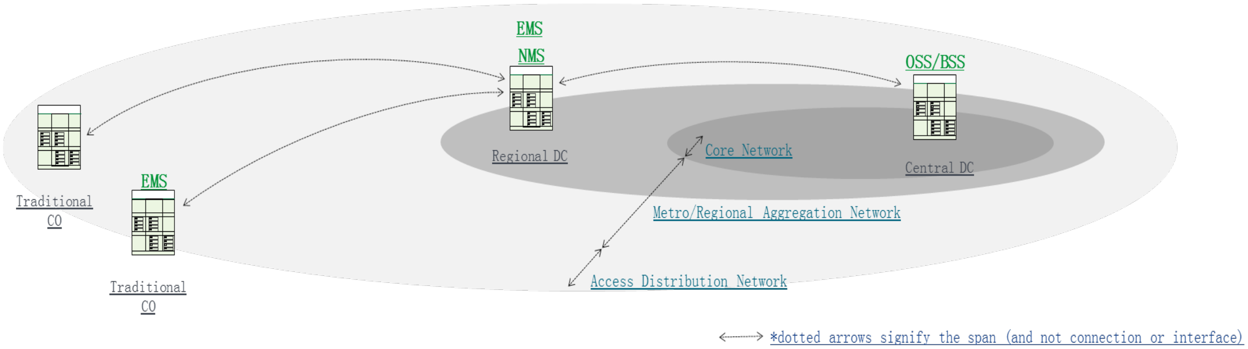
EMS
NMS
OSS/BSS
Regional DC
Core Network
Central DC
EMS
Traditional CO
Traditional CO
Metro/Regional Aggregation Network
Access Distribution Network
*dotted arrows signify the span (and not connection or interface)

Figure 8: Locations of the existing management plane

In these existing networks, the management plane is traditionally been developed using server-based Information Technologies (IT). As the IT capabilities for the Service Provider has been transformed to use Web-based virtualized technologies, many of the existing management plane deployments have also been transformed. This transformation has resulted in virtualization of the management plane functions in the Central and Regional DCs as depicted in Figure 9.

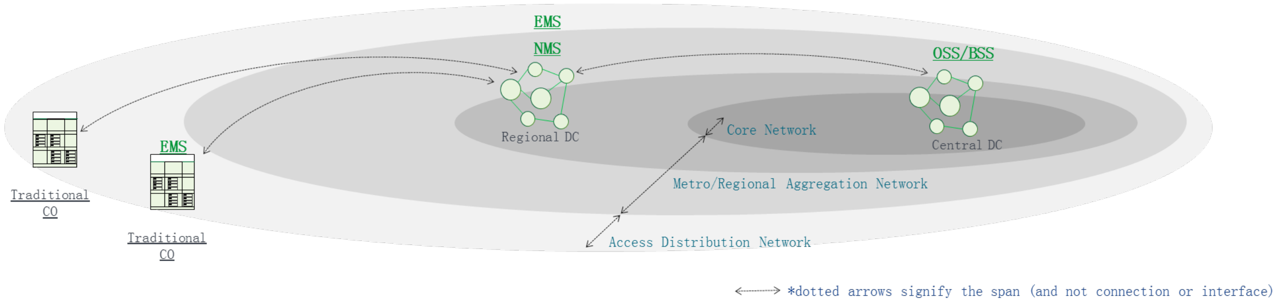
EMS
NMS
OSS/BSS
Regional DC
Core Network
Central DC
EMS
Traditional CO
Traditional CO
Metro/Regional Aggregation Network
Access Distribution Network
*dotted arrows signify the span (and not connection or interface)

Figure 9: Virtualization of the existing management plane

As the management plane of the CCO is expected to be virtualized, the migration of the management plane toward Domain Controllers and Orchestrators can take advantage of the existing IT infrastructure and ecosystem, at least taking advantage of the existing Rackspace, power and cooling, and in some cases existing compute facilities. In Figure 10, the existing management plane is migrated toward the CCO concepts of Domain Controllers (SDN M&C) and Orchestration (CCDO) functions in the location(s) where there is virtualized IT infrastructure. In the case of the Access network in the Central Offices, the BAA layer could provide the distributed management and control of existing Access network elements.

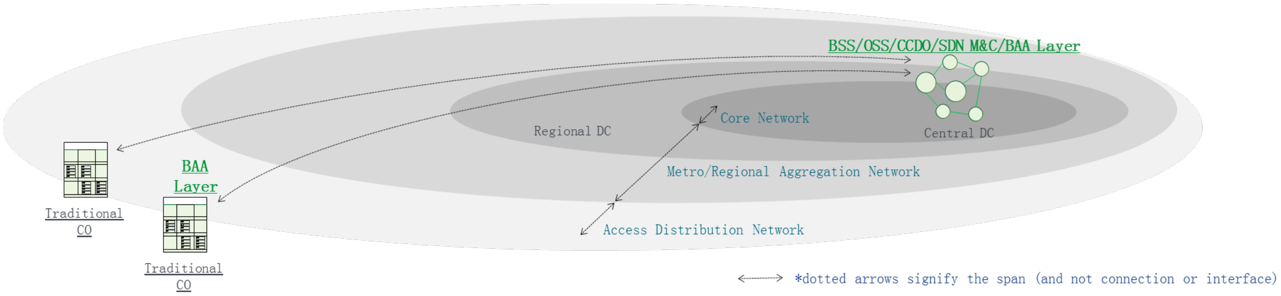
BSS/OSS/CCDO/SDN M&C/BAA Layer
Core Network
Regional DC
Central DC
BAA Layer
Traditional CO
Metro/Regional Aggregation Network
Traditional CO
Access Distribution Network
*dotted arrows signify the span (and not connection or interface)

Figure 10: Migration of the existing management plane toward CCO

As more and more control and user plane functions are virtualized and the associated virtualized infrastructure is located on the Aggregated or Far-Edge locations, the Domain Controller and Orchestration functions can be further distributed as depicted in Figure 11.

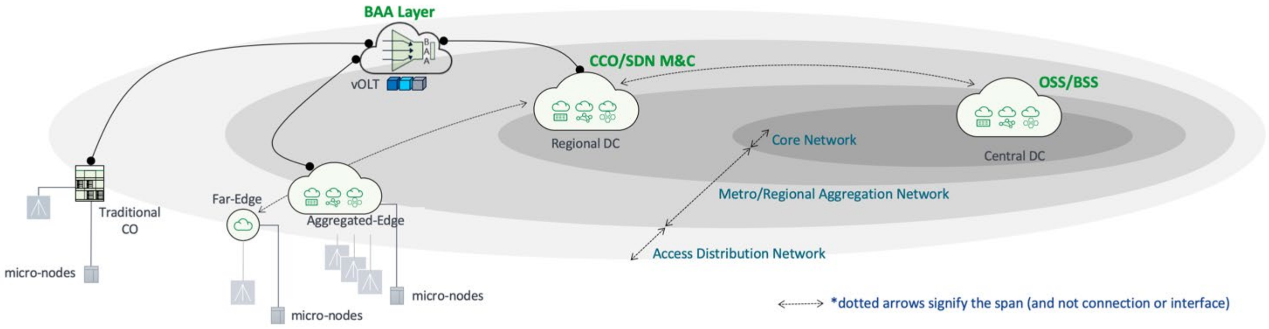

Figure 11: Distributed Cloud CO management plane

Figure 11 like the other figures in this section of the Technical Report is illustrative only of the principles that a transformation of the management plane can follow:

- Take advantage of existing IT infrastructure, at least when it comes to power, cooling and rack-space.
- Centralize management and control functions as much as reasonable, distributing the functions when there is a technological or business need.

4.3 Transforming the Control and User Plane

4.3.1 Considerations for Network Function Virtualization Infrastructures

While it is conceivable that the (network) management function migration described in section 4.2 could leverage existing IT infrastructures, a different picture emerges when considering network control plane and user plane functions. Migrating of some control plane functions and all user plane functions could involve newly implemented NFVI's at locations where these functions reside e.g. at BNG and Access Node locations, as well as more central locations (regional DC). We refer to these NFVI's as 'CCO Macro-Nodes' through-out this document, as per TR-384 [3]. These NFVI's do not need to host the Management & Orchestration parts of the Cloud CO architecture, only the VNFs needed in scope of the migration. The same consideration applies to the VIM layer as well as any other components that are needed to manage and operate the virtualization infrastructure itself, and which are typically abstracted by the VIM layer. We can refer to these components as 'NFV Infrastructure Management'.

SDN, MANO and NFV Infrastructure management components can therefore be more centrally deployed, and this would require an NFVI. We will refer to this 'management' NFVI as the 'CCO MGMT-Node'. Various CCO MGMT-Nodes could be deployed, typically across at least 2 availability zones, with one CCO MGMT-Node taking care of several regional CCO Macro-Nodes. Separating the NFVIs into 'Management' (CCO MGMT-Nodes) and 'Resource' (CCO Macro-Nodes) also follows the standard best practices for Cloud Deployment, as both NFVIs will have different requirements in terms of scale and performance. This allows the service provider to centralize the deployments of a lot of supporting components of the Cloud CO architecture and allows the Service Provider to deploy these components in facilities that already are equipped for IT infrastructure, such as Central and Regional Data Centers.

CCO Macro-Nodes can also be deployed together with CCO MGMT-Nodes at the same facility (e.g., Central or Regional Data Center). This allows the operator to migrate the Network management and some control plane functions without having to deploy CCO Macro-Node NFVIs at the edge locations.

CCO MGMT-Nodes might not be able to leverage the existing IT infrastructure in some cases for a range of reasons, such as availability of facilities and equipment (e.g., rooms, racks, servers), infrastructure management technologies used by the IT infrastructure, and operator organizational structure. However, they will be deployed using existing IT Best Practices, and quite likely at the same locations as the existing IT infrastructure, using the same power/cooling and potentially rooms. It is to be noted that sharing the same virtualization infrastructure for both IT and NFV workloads can have its advantages from an operational model point of view. However, this would require organizational alignment within the Service Provider. It is likely that, after deploying the CCO MGMT-Nodes, that all already-migrated management plane functions, would be re-migrated to them, in case they had been running on the existing IT Infrastructure.

Figure 12 gives an overview of the high-level architecture with the various regions that are supported within the Cloud CO architecture. In the figure the CCO-MGMT Nodes are centrally deployed, supporting various CCO Macro-Nodes at various Edge Locations for every region.

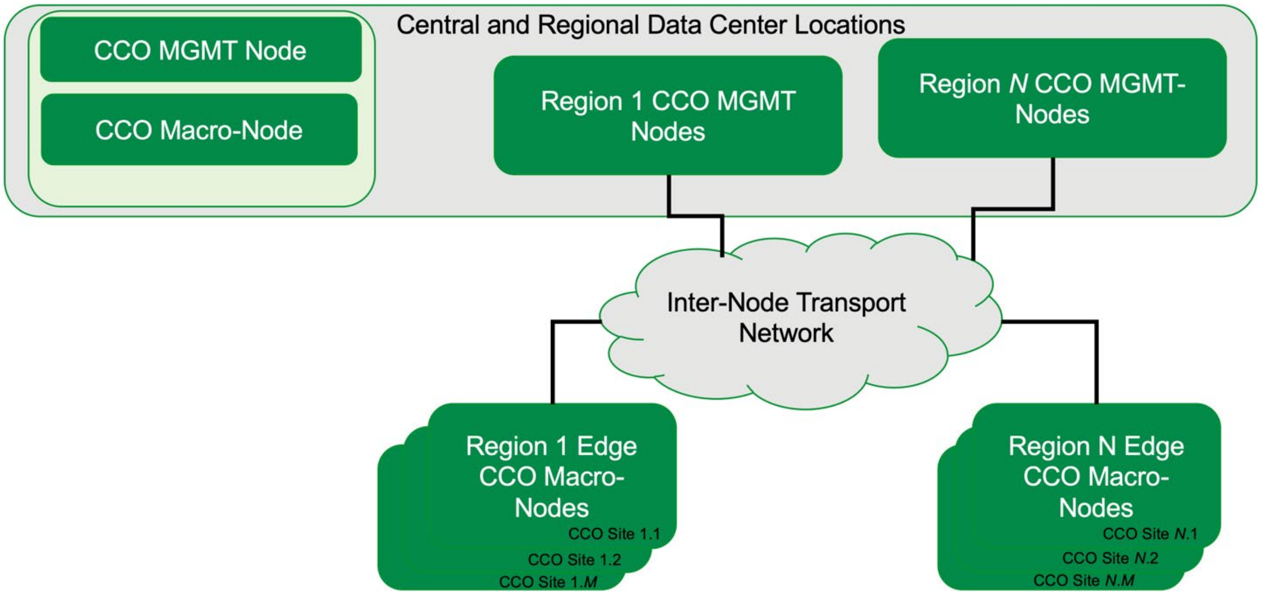
Central and Regional Data Center Locations
CCO MGMT Node
CCO Macro-Node
Region 1 CCO MGMT Nodes
Region *N* CCO MGMT-Nodes
Inter-Node Transport Network
Region 1 Edge CCO Macro-Nodes
CCO Site 1.1
CCO Site 1.2
CCO Site 1.*M*
Region N Edge CCO Macro-Nodes
CCO Site *N*.1
CCO Site *N*.2
CCO Site *N*.*M*

Figure 12: CCO Macro and MGMT Nodes

Next to the Regional/Central Data Center, we can foresee additional locations to be used to deploy network functions. Disaggregation allows us to distribute management plane, control plane and user plane network functions across those locations, even for one given service. The various locations include:

- Customer Premise: where managed services such as SD-WAN can be placed, some of them virtualised.
- Far Edge Office: ultra-low latency apps such as AR/VR can be placed here and are typically constrained by cooling/power/space.
- Near Edge Central Office: User Plane functions that require high user density, low latency and high throughput will be placed here, as well as, distributed control plane functions. Examples include. CDNs, NERG, and the Mobile User Plane.
- Regional Data Center: All Management Plane functions and Centralised control plane functions (if the user plane to control plane protocols allows it). An example is the Mobile Control Plane VNFs.

The following table summarizes where various components of the TR-384 architecture resides in Figure 12.

	NFV Infrastructure Management Component	Network Management Plane Function	Network Control Plane Function	Network User Plane Function
Central/Regional Data Center	X	X	X	X
Near Edge Central Office			X	X
Far Edge Office				X
Customer Premise				X

Table 2: Locations for TR-384 Architectural Elements

The design of a CCO Macro-node depends on what functionality is required, what scale is needed in terms of deployed workloads, latency and load, and where the CCO Macro-Node is deployed. If the CCO Macro Node is deployed at the Central Office ('Near Edge'), it will typically involve a leaf-spine fabric with various

PNFs and compute hosts attached, as this will provide the capability to scale out the CCO Macro-node efficiently and without incurring down time. Network user plane migration will definitely be a trigger to deploy these CCO Macro Nodes at those locations. However, when the CCO Macro-Node is deployed at a more remote location ('Far Edge'), or even at Customer Premise locations, a set of switches with two or more compute hosts might be sufficient.

For a deeper discussion on deployment options, see Annex A.

The following sections expand more on how and when to deploy CCO Macro-Nodes to allow migration and transformation of control and user plane functions.

4.3.2 Considerations for Control and User Plane Transformation

The technology, network and service migration processes for introducing a CCO architecture, in full or in parts, into existing deployments and the associated coexistence phases are quite complex.
This section identifies several dimensions to be considered for planning and executing such a migration. Moving across these dimensions results in a multitude of combinations that describe potential migration scenarios.

Each Service Provider will define the graduality, depth and extent of the transformation of their existing MSBN over each dimension that ultimately defines the realization of their own migration plan.

The following dimensions are discussed below:

- CCO Macro-Node transformation
- Geographical penetration
- Service migration and enhancements

These dimensions are only partly orthogonal as NFVI evolution, service and management migration as enabled by VNF onboarding, investments and time are all variables of the same decision and planning equation.

4.3.2.1 CCO Macro-Node Transformation

This dimension captures the transformation path of the Central Office's existing MSBN towards using an NFVI that supports a CCO Macro-Node as described in section 6.1.1 of TR-384 [3]. Note that this CCO Macro-Node does not need to be co-located with the actual MSBN equipment. The choice of a co-located CCO Macro Node is dependent on factors such as required latency, user termination density, and throughput requirements, as pointed out in section 4.3.1

In this dimension, one or more network functions provided by the Central Office's existing MSBN is migrated to the CCO Macro-Node instance by virtualizing the network function on the CCO's NFVI.

For example, selected control plane functions provided by access and edge nodes within legacy Central Offices can be virtualized and deployed on the CCO Macro-Node's NFVI. The migration can be performed in stages for selected functionality. As part of those stages it can also be determined whether it makes sense to also deploy a CCO Macro-node inside the Central Office itself at a certain point in time.

For example, the following sequence of figures depicts the migration of a BNG and OLT access node by virtualizing the functions of the respective nodes and migrating the functions on the CCO Macro-Node NFVI's. The example depicts that not all network user plane or control plane functions require migration or that the migration occurs all at once. Additionally, the example depicts that the CCO Macro-Node's NFVI could either be centralized or distributed based on the needs of the network functions that have been deployed.

In Figure 13 the assumption is that the service provider transforms a few virtualized functions that might not require co-location of those functions with the existing MSBN equipment. For example, the AAA and DHCP capabilities of the BNG and access node could be virtualized along with a newer value-added service (VAS) and hosted on the centralized CCO Macro-Node inside the Regional Data Center. The figure also does not take into account any Management Plane functions which would have been migrated already in a previous phase.

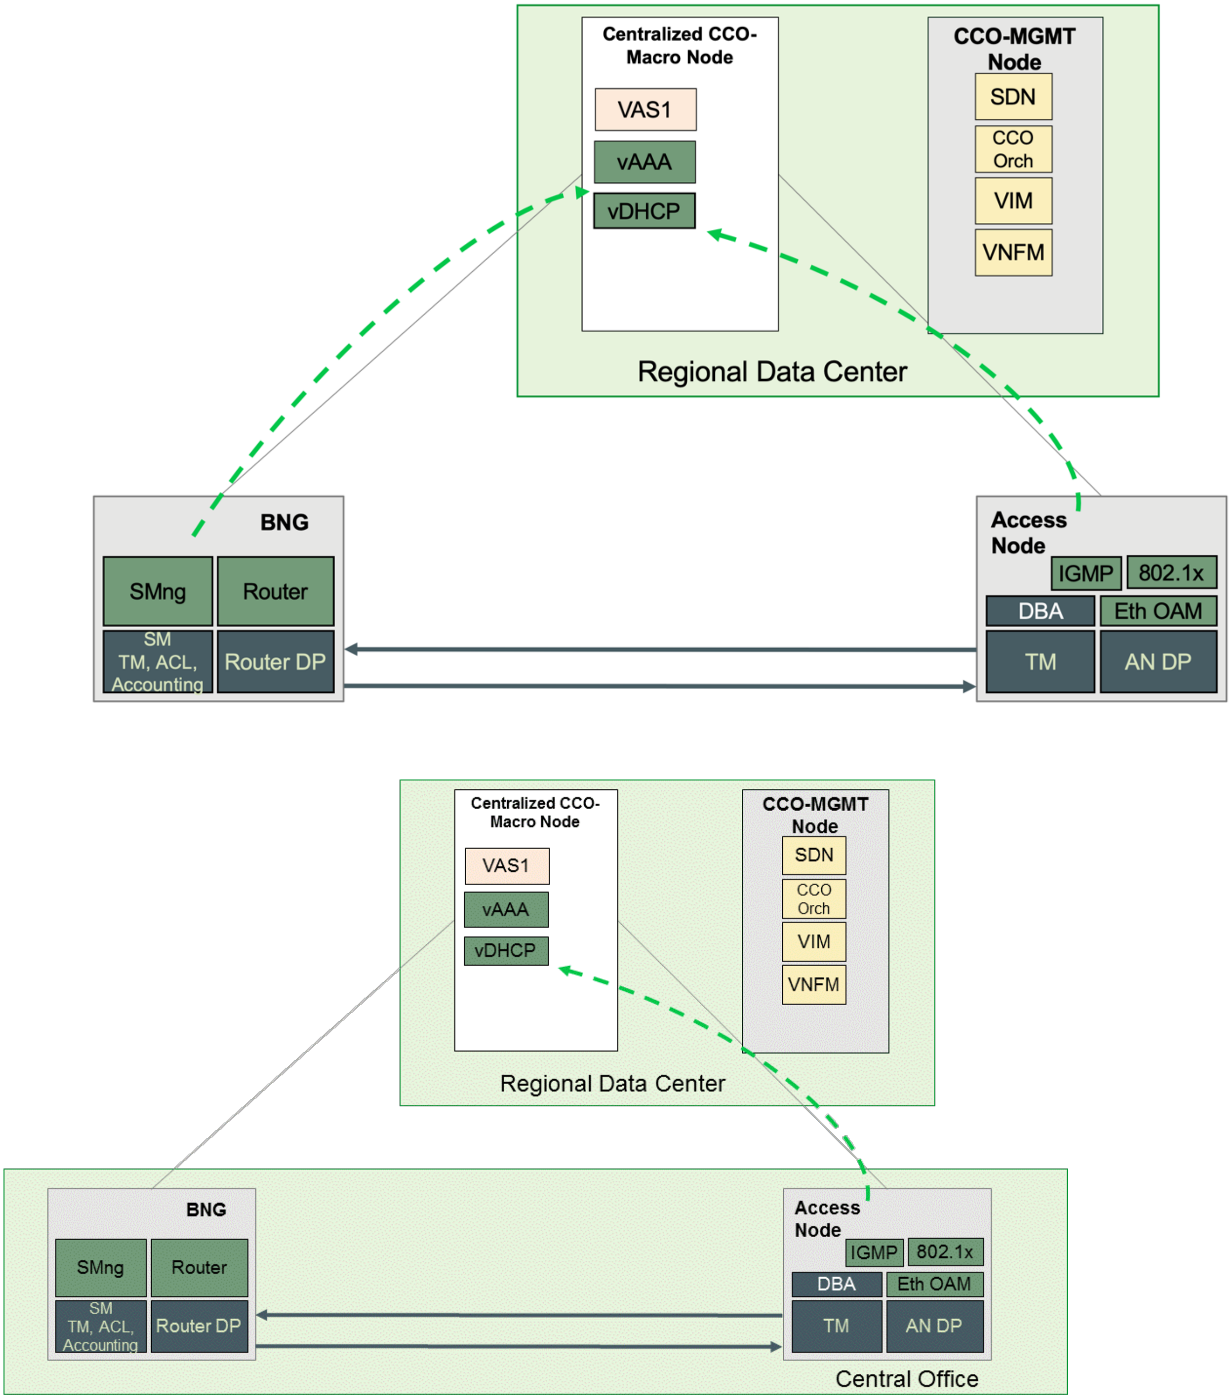

Figure 13: CCO Macro-Node transformation without a fabric

As the service provider decides to migrate additional functionality that require co-located NFVI to connect the functions, the service provider can introduce a full CCO Macro-Node inside the Central Office that permits the needed connectivity as depicted in Figure 14.

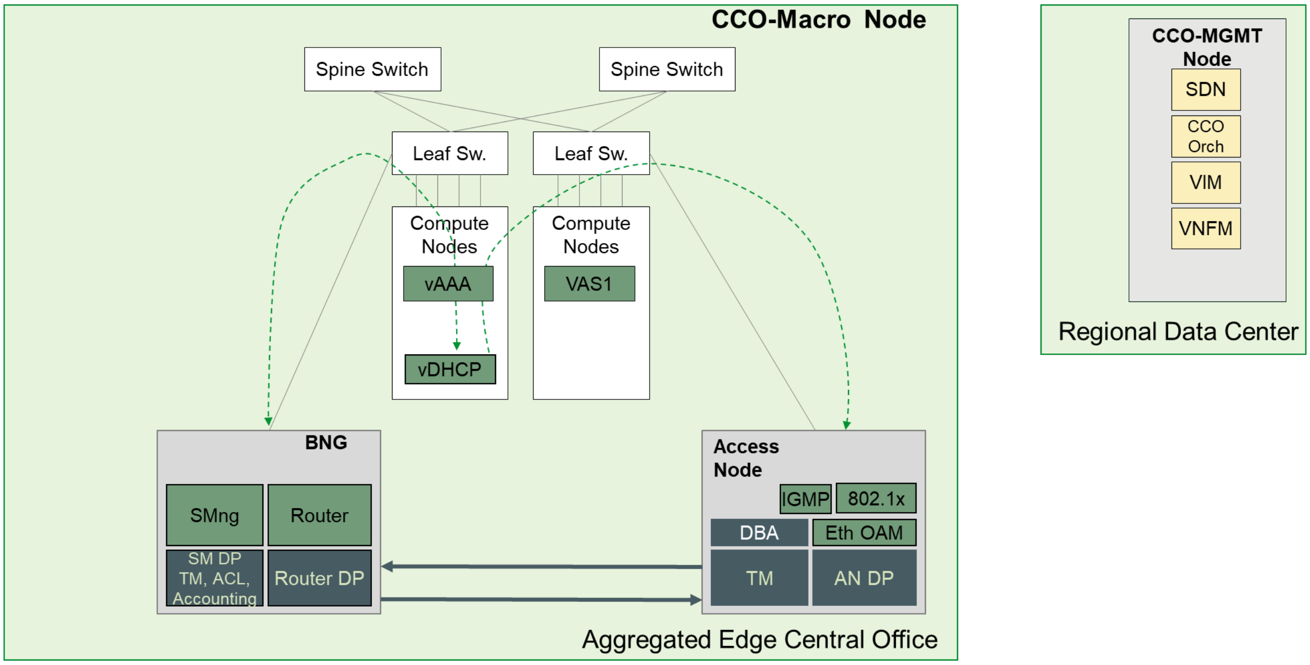

Figure 14: CCO Macro-Node transformation with a minimal fabric

With the CCO Macro-Node now deployed in the Central Office, the service provider can determine which functions make the most business sense to migrate. In this example, the service provider has decided to migrate the control functions of the BNG and OLT access node to the co-located CCO Macro-Node, migrate the Management Functions to the centralized CCO Macro-Node, but keep the latency sensitive user plane functions within the nodes as depicted in Figure 15. The fabric is able to scale out as needed by adding extra leaf switches if required without incurring downtime. For more information on fabric scaling see Annex A.

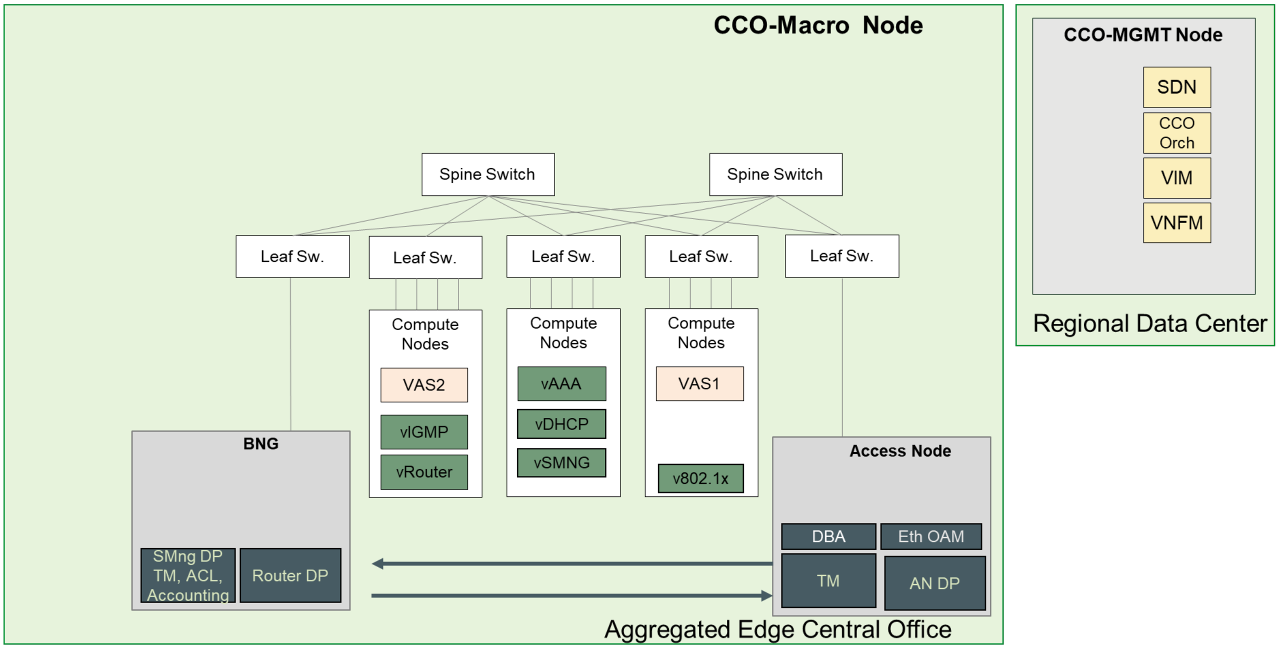

Figure 15: CCO Macro-Node transformation function migration on fabric

4.3.2.2 Geographical Penetration

This dimension captures the deployment path where network functions of the existing MSBN are migrated to locations other than where the functionality contained within the broadband node currently reside. Instead network functions are migrated to different sites (e.g., Regional DC, Near Edge Central Office, Far Edge Remote Office) that deploy the CCO infrastructure. Many times, this will occur because for reasons related to required investments and to the strategy in deploying specific virtualized network functions where this migration to happens starting from central sites and then to progress to smaller and more distributed sites in the network.

Figure 16 depicts an example of this dimension where the VAS services are deployed in Regional and Edge DCs based on certain optimization criteria for the deployment (e.g., cost, complexity, technical aspects such as latency) and Edge network functions are deployed in the Edge DC.

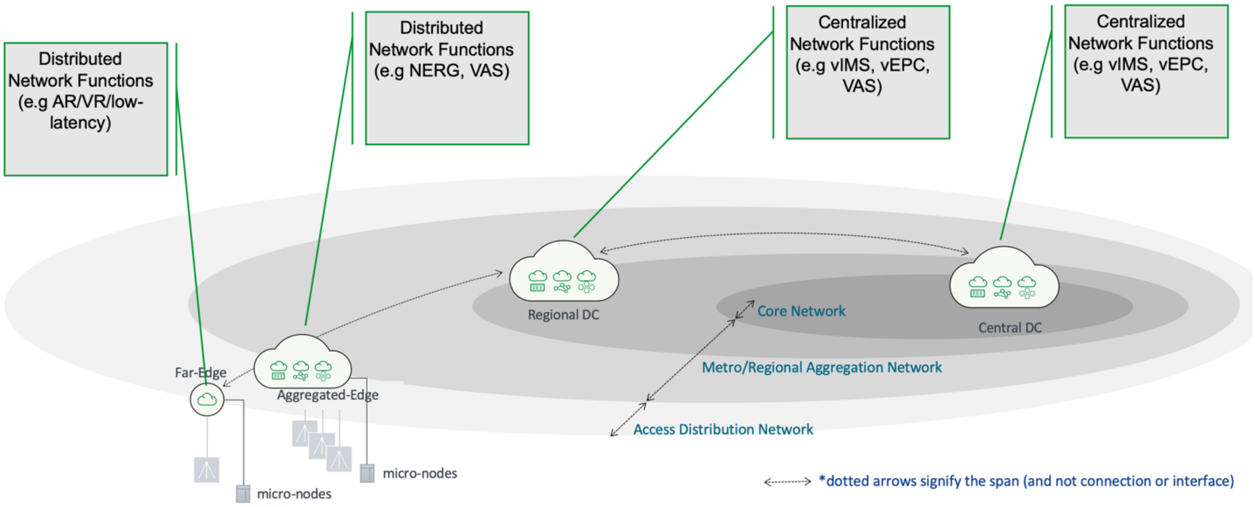
Distributed Network Functions (e.g AR/VR/low-latency)
Distributed Network Functions (e.g NERG, VAS)
Centralized Network Functions (e.g vIMS, vEPC, VAS)
Centralized Network Functions (e.g vIMS, vEPC, VAS)
Regional DC
Core Network
Central DC
Far-Edge
Aggregated-Edge
Metro/Regional Aggregation Network
Access Distribution Network
micro-nodes
micro-nodes
*dotted arrows signify the span (and not connection or interface)

Figure 16: Example of geographical distribution of Cloud CO DC infrastructure

As mentioned previously the optimization criteria for where network functions are deployed is critical to proper operation of the network function, surely when it comes to User Plane functions.
These criteria are amongst others:

- Latency requirements between the user and the Network Function (examples include Virtual Reality, Augmented Reality)
- User Density and associated scaling
- Throughput requirements

Note that the expectation is that the resources or “size” of such an Edge deployment would likely be quite smaller with respect to the more centralized deployments like the Regional DC. However, the Edge DCs would be much more numerous than Regional DCs in order to meet the expected quality of experience across the Service Provider’s network.

The penetration of DC-like infrastructure at sites with different roles in the network hierarchy lends itself to designing each site with the appropriate size (e.g., DC, Mini-DC, Micro-DC) which matches the steps depicted for the CCO Macro-Node transformation dimension and leverages the built-in scalability of this kind of infrastructure.

It can also be fore-seen that Far Edge locations do not need a full-scale NFVI to deploy the necessary functions at that location. The Near Edge locations would have to start with a fully scalable fabric as soon as it is decided to deploy virtualized network functions in that location. For more information see Annex A.

4.3.2.3 Service Migration and Enhancements

Services in the MSBN are comprised of a set of network functions that are deployed as part of the network control plane and/or user plane. In many cases, these network functions may be decomposed and distributed across different geographical locations. For example, there is ongoing work in the Broadband Forum to specify the separation of the BNG control plane network functions from the BNG user plane network functions. There is also consideration being given in many places to decomposing the functions of BNGs or other subscriber service network elements so that different user plane or control plane services can be delivered via different network elements. Any of these network functions can be virtualized and/or can be controlled via the SDN management and control functions described in the preceding sections. When network functions are separated, different combinations will have different tolerances for latency and/or latency variation between the cooperating network functions. If these network functions are virtualized then the orchestration systems that administer the life-cycle and placement of the network functions and the SDN management and control systems that control and monitor the network functions need to be aware of both the communication constraints and the network infrastructure operating parameters in order to maintain operability and effective service delivery.

Additionally, the operational requirement is that any migration to utilizing virtual systems has to proceed in incremental steps and there are common techniques that can be used to perform such migrations. In a simple scenario where new augmenting services can be added independent of the existing MSBN, the network functions can be deployed using the Cloud CO NFVI infrastructure's user plane leveraging the capabilities within and the connectivity provided by the existing MSBN network elements.

In other scenarios, where the desire is to move existing services to use the virtualized network functions, those network functions can be deployed alongside the existing MSBN. In both these scenarios existing subscribers would need to be unaffected. Depending upon operational constraints and goals, existing subscribers can be migrated to take advantage of the new services and underlying network functions over time. The implications of both scenarios are that the management and northbound control plane interfaces for the existing MSBN need to be retained until such time as the migration of all subscribers and all services is complete.

4.3.2.4 Considerations for Subscriber Migration and Coexistence

As discussed in section 4.3.2.3, subscribers will have an existing MSBN subscriber configuration side by side with the Cloud CO subscriber’s configuration and the subscriber will use functionality in both the existing MSBN and the Cloud CO to receive the same set of services as well as any additional services provided by either the existing MSBN or Cloud CO infrastructure.

By ensuring that subscribers are available during the migration period to both infrastructures, risks associated with migrating the subscriber can be somewhat mitigated until the network function resources used for the subscriber’s existing services are re-allocated once the subscriber has been migrated.

In order to provide the necessary subscriber configurations needed for the existing MSBN and Cloud CO infrastructures, the Service Provider needs to consider whether to converge their subscriber management functions (e.g., Policy, Authentication, Identity) or provide replication of the subscriber management function within each infrastructure. Just as the considerations for other transformation activities, the determination of whether to converge or replicate the subscriber management functions has similar considerations that were discussed in section 4.1

5 Migration / Coexistence Use Cases

This section of the Technical Report describes example uses that depict the migration of network functions from the existing MSBN toward the Cloud CO domain or coexistence of network functions in both the existing MSBN and Cloud CO domain.

5.1 Management Plane Use Cases

This section of the Technical Report describes use cases where management network functions that are either new or part of the existing MSBN are deployed in the Cloud CO domain.

5.1.1 Migrating the Access Segment's Management Plane

For many deployments, Service Providers can decide to migrate portions of the existing MSBN based on their business needs as was discussed in section 4.1 For the SDN-enabled access network, the first migration phase of the Access network segment can include introducing SDN technologies (e.g., Access SDN M&C, BAA layer), as depicted in Figure 3 to support service automation with programmable capabilities of the existing access nodes initially and then introduce disaggregation of selected network functions considered in a future phase.

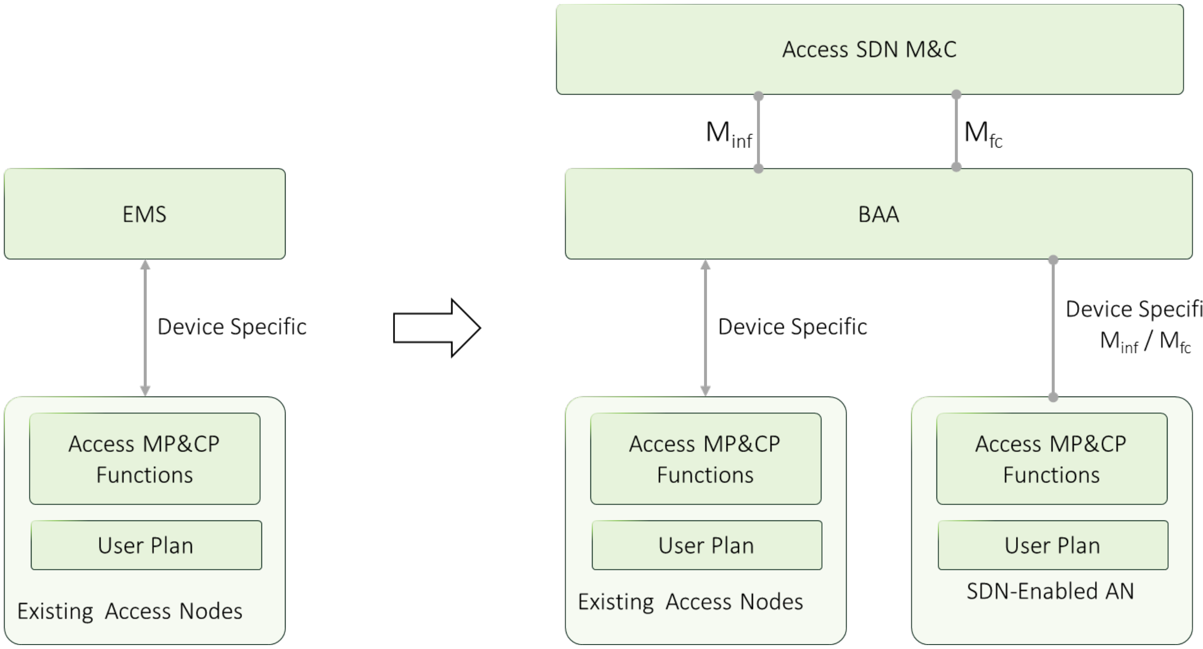

Figure 17: Migration from an existing Access segment toward an SDN Access Network

The deployment architecture for this SDN-enabled phase is depicted according to the CCO architecture and its interfaces. Migrating to this phase, the management and control interfaces are evolved to support software and automation based protocols (e.g., NETCONF) and modelling languages (e.g., YANG). In addition, the BAA layer is introduced to provide the always-on, digital representation of the access nodes (AN) needed within the SDN ecosystem and to provide the appropriate abstraction and adaptation of the access nodes.

During the initial introduction of SDN technologies phase, two types of ANs can be deployed:

- ANs that are NETCONF capable can be managed through standard interfaces defined by CCO SDN elements (e.g. NETCONF/YANG)
- Existing access nodes can be managed through the BAA layer with an adaptation between device specific and Minf/Mfc interfaces as defined in TR-413, or they can be updated to support NETCONF by software replacement if applicable.

The enabling elements for the migration of the access segment towards an SDN-enabled access network are:

1. Use of evolved management and control interfaces supporting a software and automation-based protocol (e.g., NETCONF) and a powerful modeling language (e.g., YANG).
2. Use of standard YANG Data Models developed by the IETF, BBF and other organizations, rather than vendor specific data models, to enable multi-vendor interoperability across the CCO interfaces by design.
3. Introduction of the BAA layer and the advantages offered by this always-on digital representation of ANs, telemetry data off-loading and its mediation and adaptation capabilities.
4. Support the need to migrate element and network management functionality from EMSs, OSSs and other Service Provider deployed systems into the Access SDN Manager and Controller and the associated Domain Orchestration functions as discussed in section 4.2.
5. If foreseen in the plans of the Service Provider, provide a more streamlined network automation through the adoption of a cross-segment SDN Manager and Controller which relies on a BAA-Universal Access Manager as described in the OB-BAA System Description [https://www.broadband-forum.org/marketing/download/OB-BAA-002.pdf] that acts as an Access Network Domain specific management and control element.

These enabling elements are somewhat loosely coupled together and could be seen as incremental steps themselves toward an SDN-enabled access network via multiple coexistence and migration paths as depicted in Figure 18.

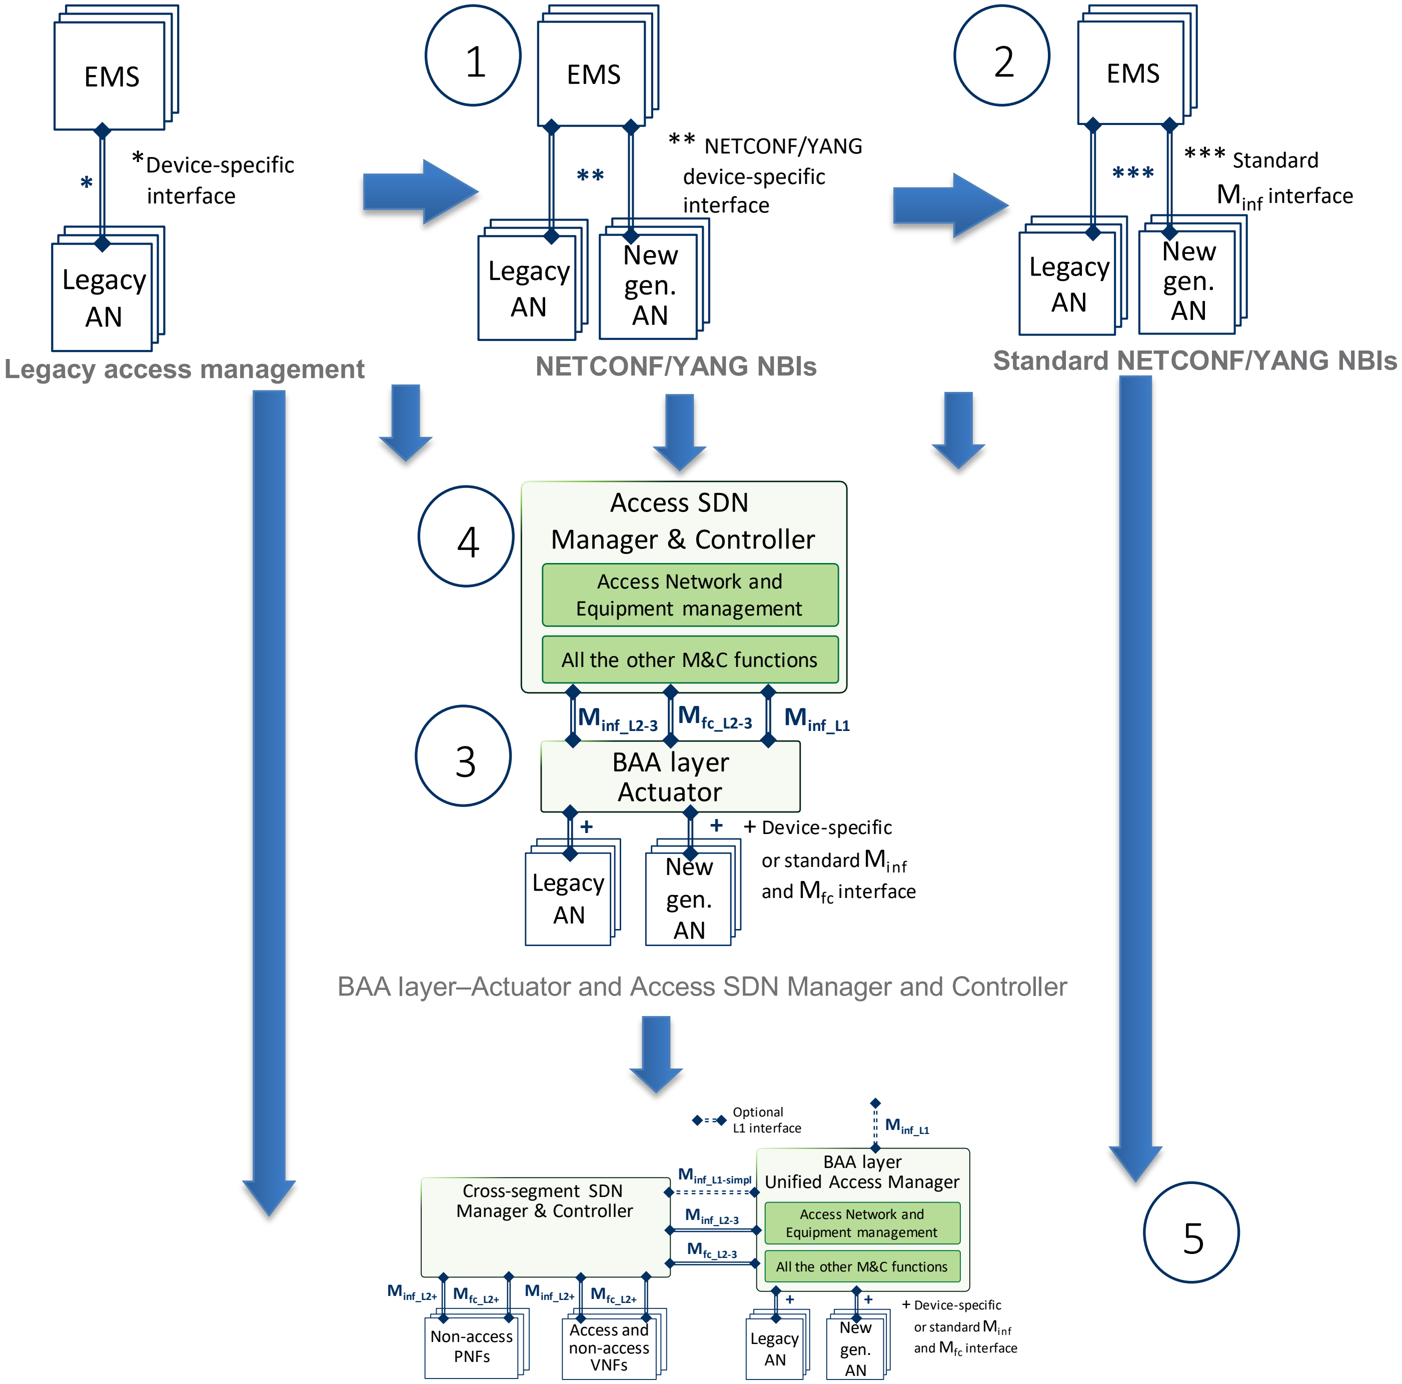
EMS
*
*Device-specific interface
Legacy AN
1
EMS
**
** NETCONF/YANG device-specific interface
Legacy AN
New gen. AN
2
EMS
***
*** Standard M_{inf} interface
Legacy AN
New gen. AN
4
Access SDN Manager & Controller
Access Network and Equipment management
All the other M&C functions
M_{inf_L2-3}
M_{fc_L2-3}
M_{inf_L1}
3
BAA layer Actuator
+
+
+ Device-specific or standard M_{inf} and M_{fc} interface
Legacy AN
New gen. AN
Optional L1 interface
M_{inf_L1}
BAA layer Unified Access Manager
Access Network and Equipment management
All the other M&C functions
Cross-segment SDN Manager & Controller
$M_{inf\_L1-simpl}$
M_{inf_L2-3}
M_{fc_L2-3}
M_{inf_L2+}
M_{fc_L2+}
M_{inf_L2+}
M_{fc_L2+}
Non-access PNFs
Access and non-access VNFs
+
+
+ Device-specific or standard M_{inf} and M_{fc} interface
Legacy AN
New gen. AN
5

Legacy access management

NETCONF/YANG NBIs

Standard NETCONF/YANG NBIs

BAA layer–Actuator and Access SDN Manager and Controller

BAA layer – Universal Access Manager and Cross-segment SDN Manager and Controller

Figure 18: Transformation of the management plane for the access network

In step 1 of Figure 18, ANs are managed via EMSs using device and vendor specific protocols and data models. Step 2 introduces newer ANs that can be managed using NETCONF with standards defined YANG models. The BAA layer is then introduced in step 3 providing the necessary adaptation from AN's device specific representations along with the always-on digital representation, telemetry data off-loading and adaptation of the ANs with device specific interfaces. Step 4 then introduces the Access SDN M&C function that is specific to the access domain. Step 5 replaces the Access SDN M&C with a cross-segment management and control entity and an enhanced BAA layer called a Universal Access Management (UAM) as described in the OB-BAA System Description [https://www.broadband-forum.org/marketing/download/OB-BAA-002.pdf that incorporates the access segment domain knowledge defined in the Access SDN M&C.

5.2 Control and User Plane Network Function Use Cases

This section of the Technical Report describes use cases where control and user plane network functions for selected services and network functions that are either new or part of the existing MSBN are ultimately deployed on Cloud CO infrastructure.

5.2.1 Network Enhanced Residential Gateway (NERG)

A Cloud CO domain supports the TR-317 [4] NERG deployments by providing virtualization and disaggregation options of the RG for Residential Services. TR-317 defines the NERG benefits, use cases, architecture and requirements.

Consequently, there might be requirements from Service Providers to evolve their networks to accommodate NERG deployments. So when Service Providers start to migrate their networks to a Cloud CO infrastructure, deploying the NERG's VNFs on Cloud CO infrastructure could be one scenario, as shown in Figure 19.

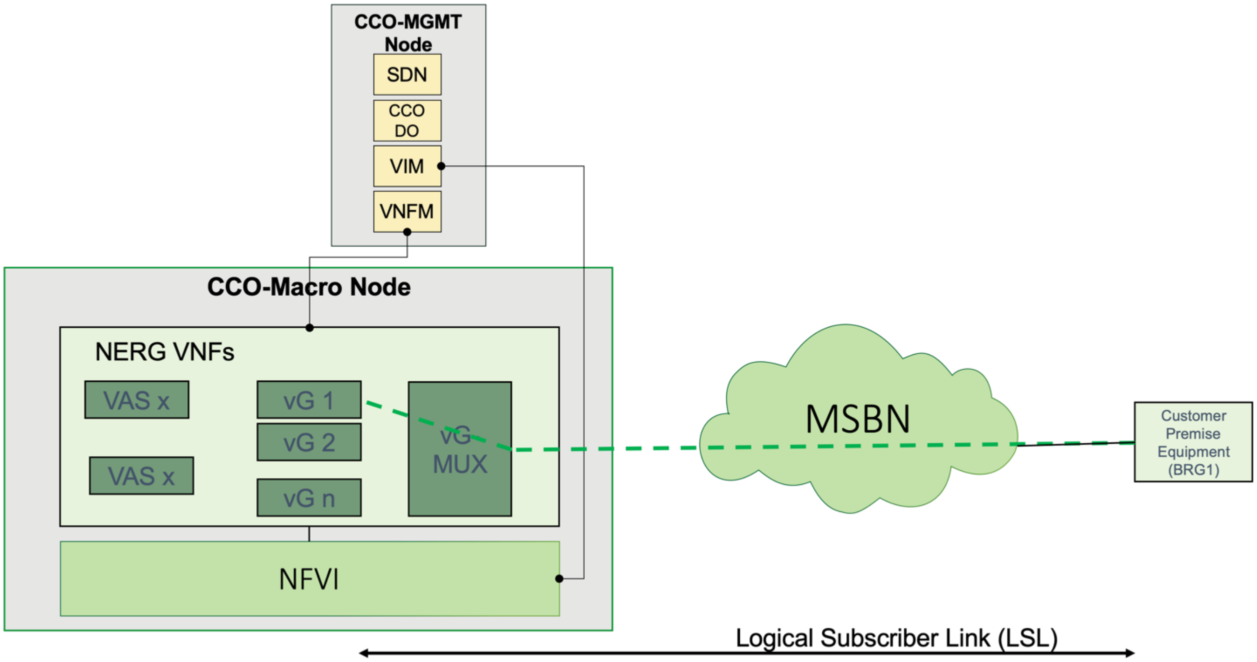

Figure 19: NERG Deployment inside Cloud CO

In this migration phase, the Access Node, as well as BNG, can be maintained as existing entities without the need to be disaggregated and distributed across directly into the Cloud CO domain and can still managed by existing management systems. They also function as the transport network for the NERG deployment.

As specified in TR-317 [4], the NERG disaggregates the RG into two components:

- A Bridged Residential Gateway (BRG), that acts at layer 2.
- A virtual Gateway (vG) that provides the network functions for the subscriber services. Examples of functions delivered by the vG include IP addressing, Network Address Translation and Value Added Services (VAS). The vG is the default IP gateway for the Devices in the customer premises.

The BRG is still located at the residential customer premises and connects to the vG through the MSBN. The vG is virtualized and deployed in the CloudCO domain..

The BRG and its respective vG are connected by a Logical Subscriber Link (LSL), which carries the subscriber L2 traffic (except the in-customer premises traffic, which is bridged locally by the BRG).

TR-317 [4] also defines a vG_MUX component which is a network function that maps L2 traffic between a subscriber's BRG and subscriber's unique vG, ensuring traffic isolation between NERG customers. In this use case, the vG_MUX is hosted on the NFVI and can be disaggregated across the Cloud CO domain.

At the same time, value added services (VAS), such as firewalling can also enabled as additional VNFs. For subscribers that have subscribed to these value added services, the vG will steer IP traffic flows from the vG to the appropriate VAS components.

The CCO Macro-Node can be co-located with the CCO MGMT-Node at the regional or central Data Center or can be located at the Aggregated Edge Central Office. Drivers of distribution will depend on:

- Small scale (trial) versus large scale NERG deployment. It can be useful to start deploying all NERG VNFs centrally, such as for a NERG trial, while backhauling all LAN traffic over the existing MSBN.
- The amount of L2 LAN Broadcast, Unknown Unicast, and Multicast (collectively known as BUM) traffic that can be aggregated at the vGMUX VNF. Distributing the vGMUX VNF will result in less aggregated L2 LAN BUM traffic. Centralizing will also mean hair-pinning all that traffic over the MSBN.
- Number of NERG subscribers per Central Office. If the number is high, a lot of LAN traffic that needs vG treatment (e.g., IGMP) might negatively impact a centralized deployment. In this case it makes sense to deploy a CCO Macro-node inside the Aggregated Edge Central Office.

5.2.1.1 Migration Process for a NERG Deployment

When deploying the NERG, Service Providers can migrate in stages toward a full deployment of the NERG as defined in TR-317 [4] using network functions in the existing MSBN infrastructure (e.g., RG, Access Node, BNG) and then introduce the Cloud CO NFVI infrastructure in the location where VAS network functions would be deployed. Using this migration strategy, the risk of introducing this solution is reduced as more of the existing MSBN is used and less new technical obstacles are introduced into the network. However, there are disadvantages to this type of migration strategy since the migration would take longer, the cost of the migration in terms of deployment costs (e.g., cost of testing during each migration stage) can increase as well as the fact that the Service Provider will still need to deploy higher cost Residential Gateways for initial deployments. These increases in the migration costs and the longer deployment timeframes need to be considered prior to selecting a migration strategy.

The following are the steps that could be used for this example of a migration process for a NERG deployment:

Step 1 : Service Enhancement
In order to provide new VAS services based on the existing MSBN's RG, the Service Provider updates software of RG to support the LSL network function of the NERG. The NFVI to support the VAS services is introduced at the network location where the subscriber network and VAS functionality will be deployed as show in Figure 20. The result of this migration step is that the subscriber's existing services and new VAS services are provided by the existing RG and new VAS VNFs.

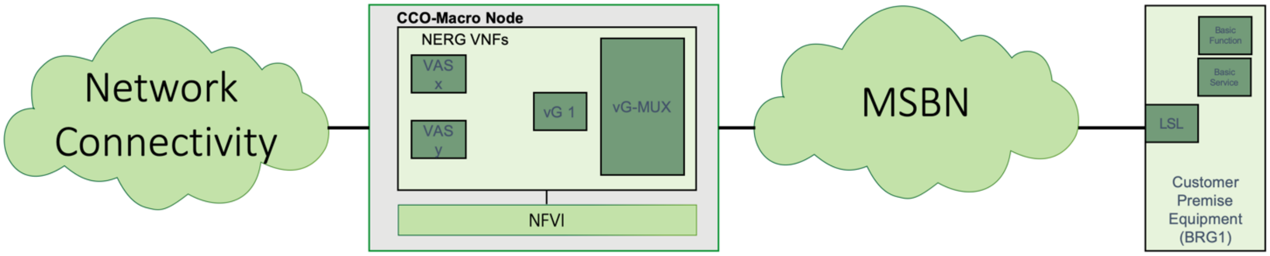

Figure 20: NERG Migration – Step 1

Step 2 : Function Migration
Once the Cloud CO NFVI is introduced, the RG has been upgraded, VAS services deployed and the Layer 2 connectivity is established between the RG and the vRG VNF, additional network functions for existing services can be migrated from the RG toward the vRG until the RG becomes the Layer 2 device as defined in TR-317 [4].

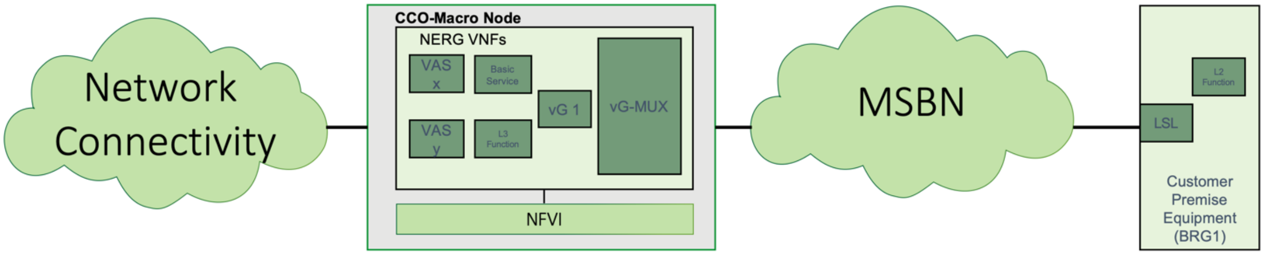

Figure 21: NERG Migration – Step 2

5.2.1.2 Coexistence process for a NERG deployment

The expectation is that the migration from existing RG deployments toward NERG deployments will take a considerable time as Service Providers have already made the capital expenditures (e.g., RG) and deployed the operational processes to support the subscriber's existing services and the Service Provider will want to realize the return on investment in these expenditures. Effectively this means, the existing MSBN infrastructure will exist until all subscribers have been migrated to the NERG. This co-existence scenario assumes that a full CCO Macro-Node has been deployed inside the Aggregated Edge Central Office, effectively co-located with the existing MSBN (AN and MS-BNG) equipment.

Step 1 : New Subscriber Deployments
In this step, subscribers that require the new VAS network functions or NERG functions are deployed using the NERG Cloud CO infrastructure while existing subscribers remain on the existing MSBN.

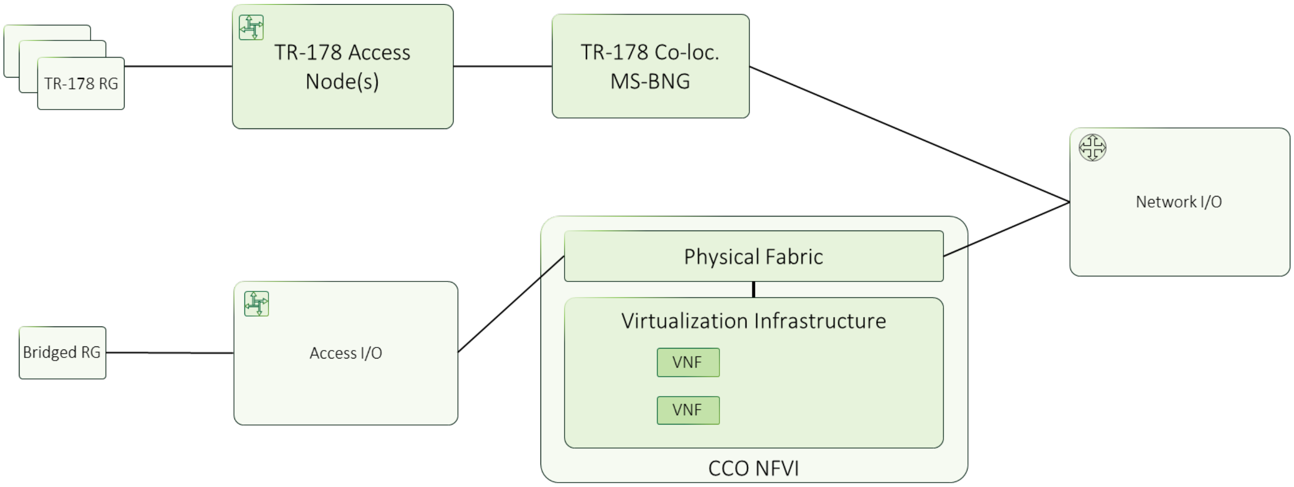

Figure 22: NERG Coexistence – Step 1

Step 2 : Existing Subscriber Migration
Once the Cloud CO NFVI is deployed at a specific location, other existing subscribers can be migrated toward the Cloud CO NERG deployment either by replacing the existing RG with a lower cost bridged RG or via software upgrade of the existing RG.

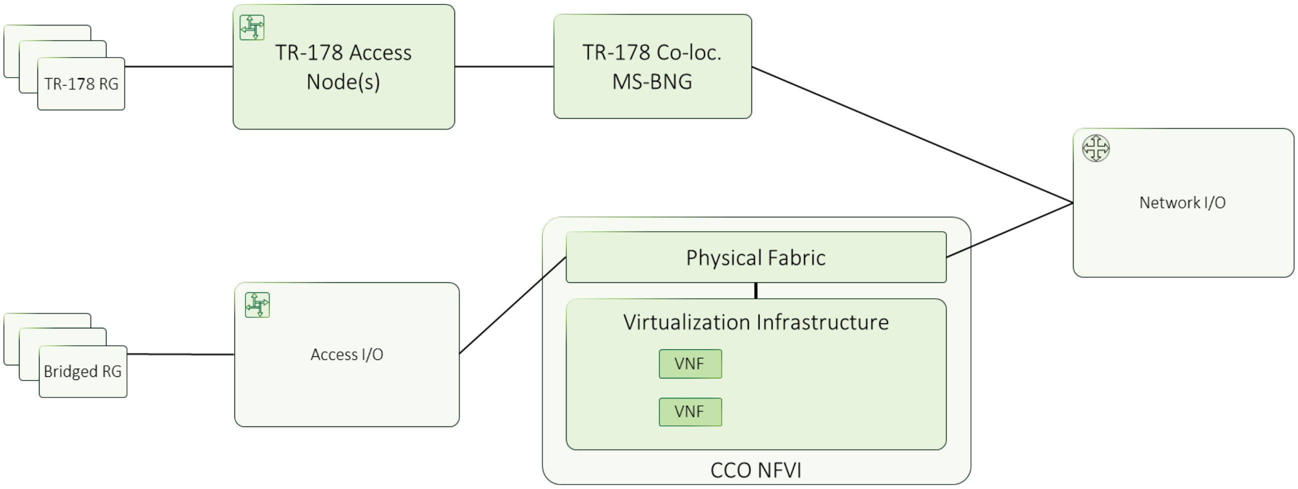

Figure 23: NERG Coexistence – Step 2

Step 3 : Decommission Existing MSBN
Once all subscribers have been migrated away from the existing MSBN, the network elements associated with the subscriber's services can be decommissioned or reused for other purposes.

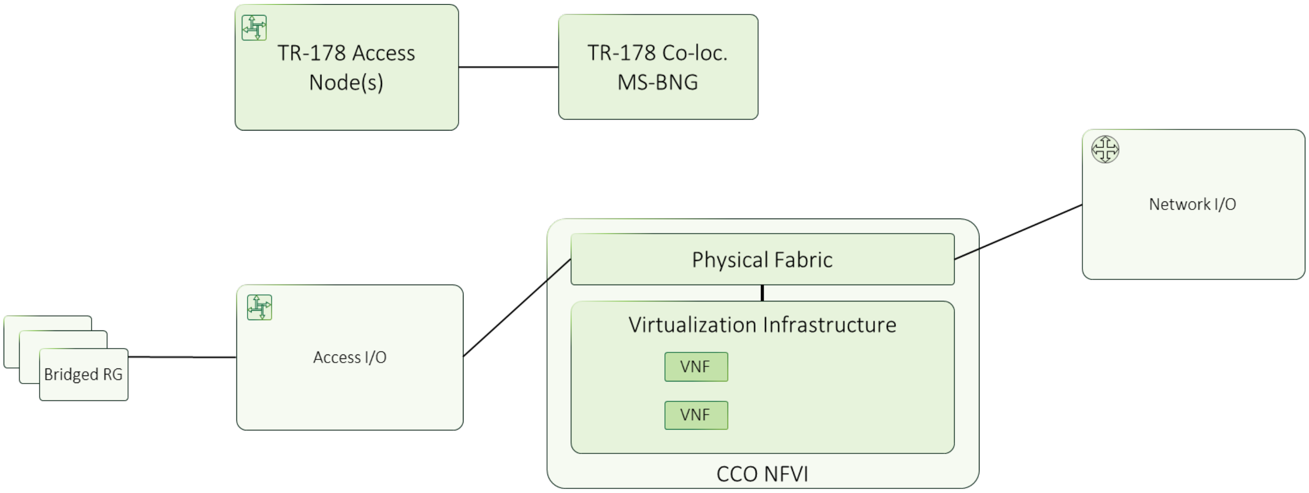

Figure 24: NERG Coexistence – Step 3

5.2.2 Broadband Network Gateway (BNG)

The migration of BNG Control and User Plane functionality from the existing MSBN toward the Cloud CO results in a number of possible deployment scenarios and migration sequences. Future issues of the Technical Report will address the BNG Control and User Plane migration scenarios.

Annex A: Design and Scaling of the NFV Infrastructure and Fabric

A.1 Introduction

One of this Technical Report's main recommendation is to start building out NFVI Infrastructures in (existing) Regional and Central Data Center locations and building out NFVI Infrastructures at the edge i.e. in the Central Office when needed (mostly for User Plane and Control Plane Network Function transformations). In any case, when an NFVI Infrastructure is built, it needs to be able to scale out efficiently. The term 'scaling out' here refers to the possibility of adding more compute and networking capacity without any impact on existing services that are carried over the infrastructure.

A.2 Scalable Fabric Design

A Data Center interconnects virtual functions, applications, users and data. It therefore needs to be able to connect to any physical equipment that allows those users, functions and/or applications to connect to their data, as well as connect the compute and storage resources needed to host these functions, applications and/or data. The underlying network topology within the data center supporting these data centers are often referred to as "fabrics."

In modern Data Center architectures, Leaf-Spine fabrics are deployed because of their unique properties around the ease of scaling out of:

- Bandwidth and connectivity between Networking Functions, whether they are physical or virtual without impacting existing services.
- Compute capacity without the need for reconfiguration of the fabric or the need to incur downtime.

In this design, compute nodes and/or PNFs connect to one layer of switches, referred to as 'leaf'. Typically, a rack contains a pair of leaf switches, to which all PNFs and compute nodes inside the rack connect to redundantly. Adding racks with additional pairs of leaf switches takes care of scaling out the compute capacity, and the PNF connectivity.

A layer of 'spine' switches in introduced and all leaf switches are wired to all spine switches with equal bandwidth links, thereby creating a 'leaf-spine' fabric. If one can ensure that all leaf uplinks are capable of forwarding traffic simultaneously the bandwidth scaling out requirement is satisfied as well. (There are various ways of achieving this, such as leveraging IETF NVO3 style overlays on the compute hosts and/or PNFs and making all links in the fabric routed. Other ways of achieving this 'Equal Cost Multi-Pathing' behavior is out of scope of this document).

In this topology there is always the same bandwidth, and connectivity model, between any two leafs in the fabric. In this way the amount of bandwidth between Leaf Switches is easily expanded, by adding additional Spine switches to the topology without the need to re-wire the infrastructure, or to momentarily break the physical connectivity.

The below figure shows how to add leaf-leaf capacity by adding an additional spine switch.

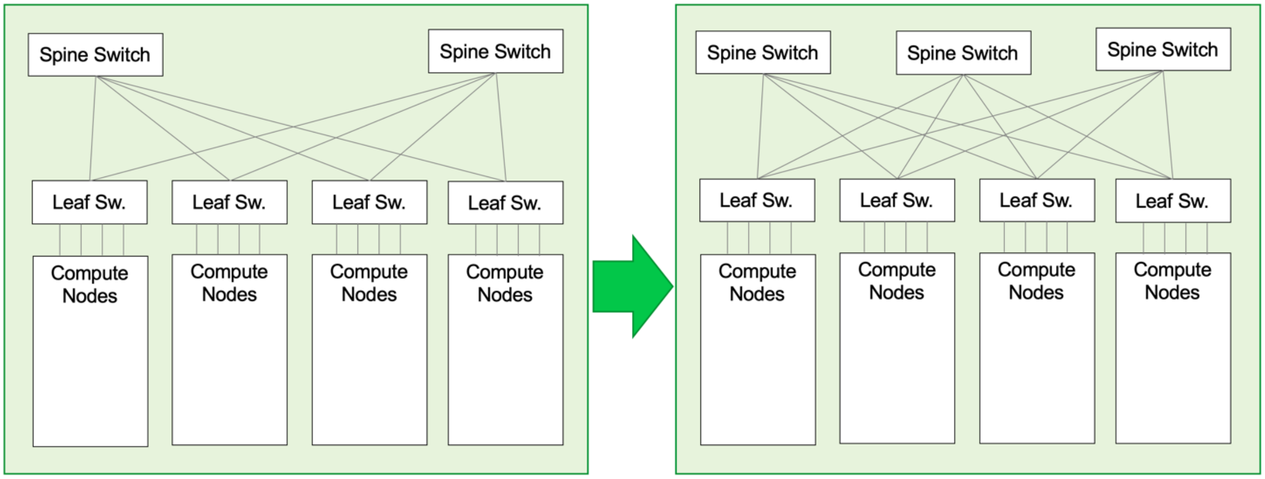

Figure 25: Scaling a leaf-spine fabric

A.3 A Minimal Fabric for the Cloud Central Office

This document also refers to a 'minimal fabric'. This is the minimal amount of leaf and spine switches needed to ensure that adding PNFs, compute nodes, or increasing bandwidth requirements do not result into any incurred downtime.

In the Aggregated Edge location (i.e., Central Office), it makes sense to wire the Access I/O into a different pair of leaf switches than the Network I/O, as this allows to scale them independently when additional capacity and/or throughput is needed. The rest of the leaf switch ports can be used to connect compute nodes or other PNFs.

In order to interconnect these two pair of leaf switches, two spine switches are needed as a minimum, for redundancy reasons, assuming that this can satisfy the inter-leaf capacity requirements. We end up with the below figure for a Minimal Fabric for CCO Macro-Node.

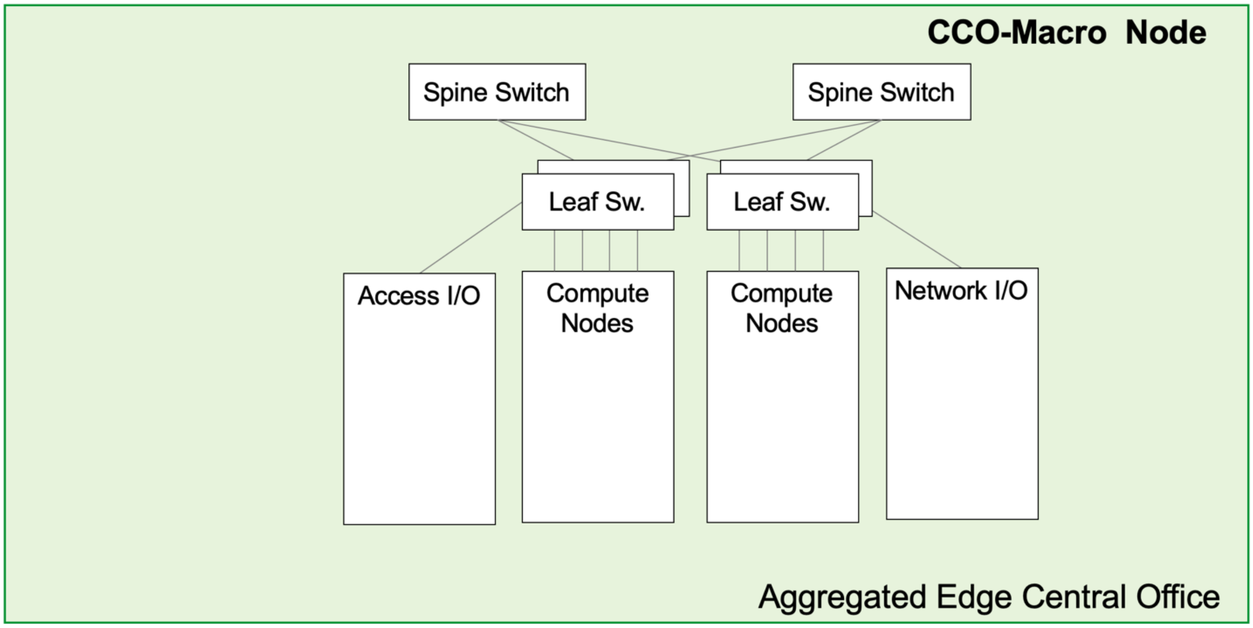

Figure 26: CCO Macro-Node with Minimal Fabric

The below figure shows a scaled-out fabric through adding additional leaf switches.

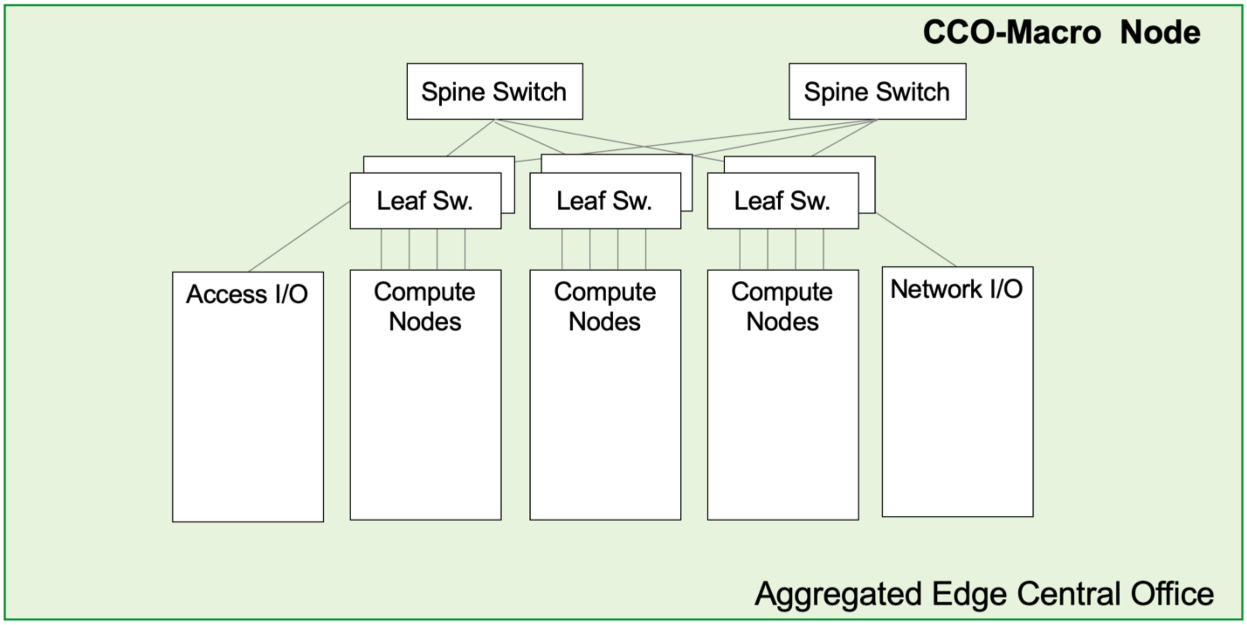

Figure 27: CCO Macro-Node with Expanded Fabric

A.4 Far Edge and Customer Premises NFVI deployments

While certain Far Edge locations (e.g., Enterprise campus, large multi-dwelling units) might require fabrics for their deployments, in other remote locations such as Residential or small Customer Premises locations where the connectivity requirements are modest and the number and requirements of VNFs are low, introducing a fabric might not make sense from a scaling perspective. In these cases, integrated networking and compute deployments are often done at these locations that do not require a fabric.

Annex B: VNF Deployment Use Cases

B.1 Small Scale VNF Deployment (1) use case by leveraging existing infrastructure

In this use case, a Service Provider needs to deploy a very limited number of VNFs that can be accommodated by a small compute infrastructure. This compute infrastructure would be connected to an existing BNG node as shown in the figure below. The VNFs could be managed by a centralized Cloud CO Management Infrastructure (a CCO MGMT-Node).

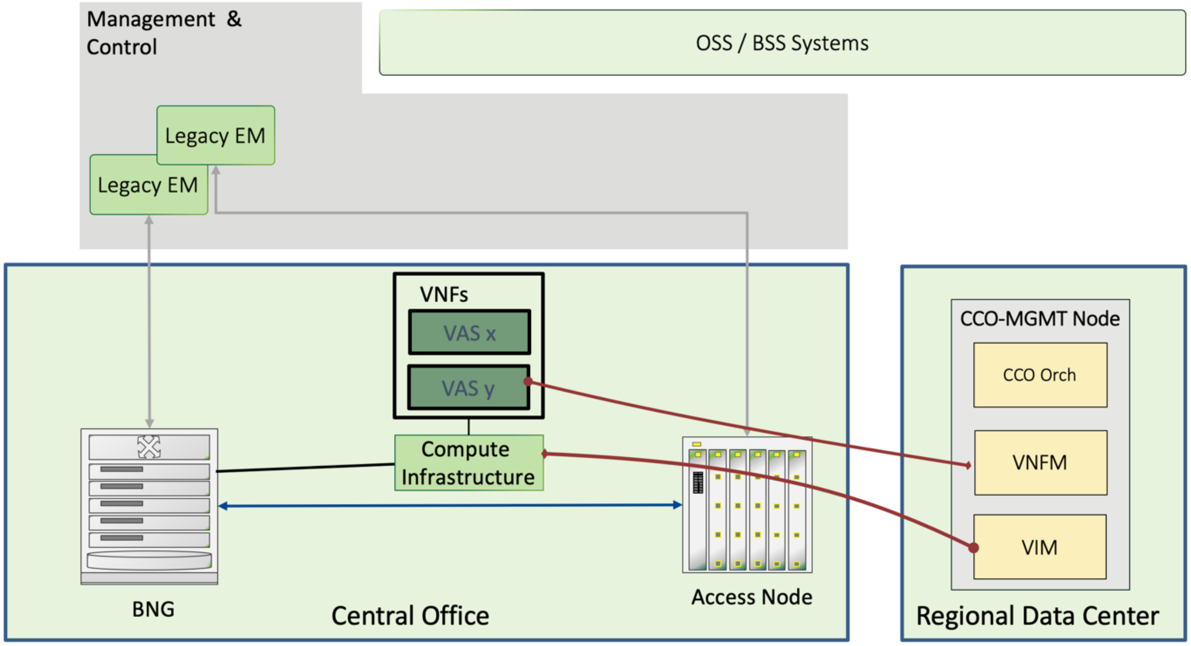

Figure 28: Small Scale VNF Deployment (1) use case

As the use case is limited to the inclusion of VAS services not supported in the existing BNG, the existing BNG could be used with no changes, provided that it is able to redirect VAS traffic towards the compute infrastructure, on a per subscriber basis. The traffic from a small number of subscribers would be redirected to the VAS.

The redirection of traffic at the BNG for the use case can be orchestrated through the BNG EMS from an End-to-End Service orchestrator as this insures that no upgrade is needed on the BNG itself.

In this use case, the access node and BNG remain unaware of the Cloud CO, without the need to be refactored, and thus avoiding the need to upgrade existing Access Node and BNG management systems.

This could enable the start of the migration towards a Cloud CO architecture with VAS services of high value which are applicable to a small fraction of subscribers enabling "trial" scenarios for specific VAS services.

B.2 VNF Deployed after introducing SDN Control of Access and Edge.

In this use case, the Service Provider needs to deploy a number of VNFs who initially can be accommodated in a few compute nodes. However, in order to enable scaling of traffic redirected to those VNFs, the Service Provider introduces a minimal fabric that interconnects the Access Node, BNG and compute nodes. Minimal Fabrics are deployed inside Cloud CO architectures because of their:

- Properties around the ease of scaling out of bandwidth between VNFs, or bandwidth between Physical and VNFs.
- Ability to reduce downtime.

For an explanation of a Minimal Fabric, see Annex A: as depicted in Figure 1 below: 4.1

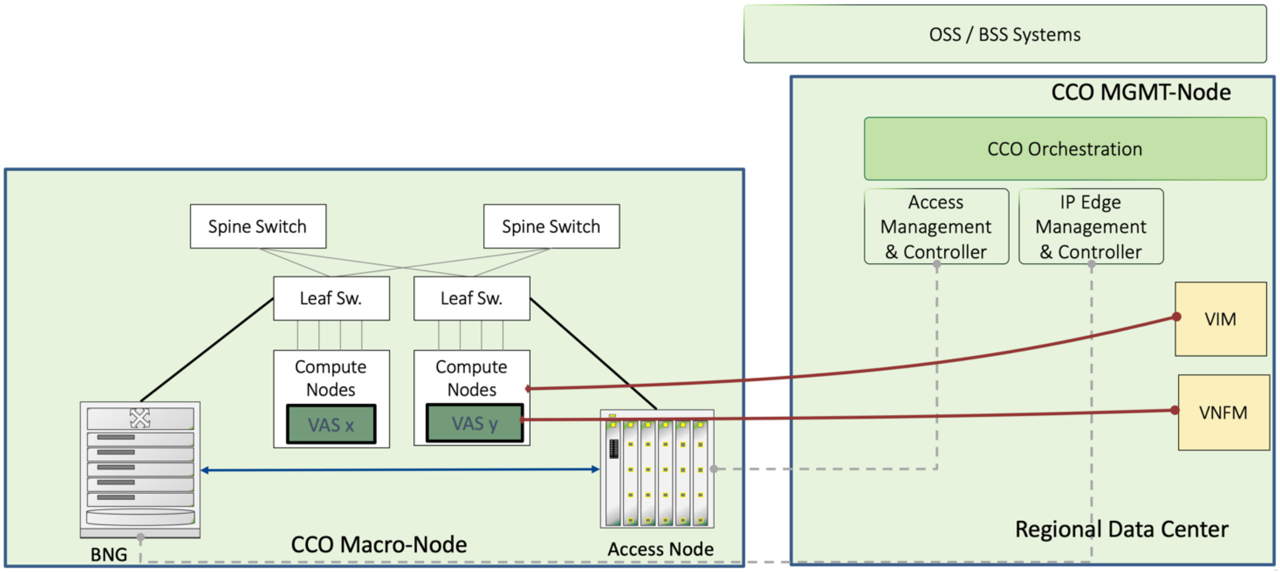

Figure 29: VNF Deployed on a Minimal Fabric use case

As traffic would need to be redirected from both the Access Node and/or BNG, the Access Node and BNG functions would need to be upgraded to allow management using the SDN Management and Control functions that are part of the Cloud CO Management framework.

End of Broadband Forum Technical Report TR-408