Virtual Central Office

Building a Virtual Central Office (VCO) with open source communities and components

October 2017

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Central Office (CO) virtualization is a hot topic for communications service providers (CSPs) today. Service providers face increasing competition in global markets, and are highly motivated to improve service agility and operational efficiency in order to improve customer experience and lower costs.

In the US alone there are more than 10,000 central offices with a broad range of equipment and configurations which have traditionally been managed locally and manually. This provides tremendous potential for optimization through network virtualization, software defined networking (SDN), Network Functions Virtualization (NFV), and standardization through open source-based, interoperable solutions.

This whitepaper is an exploration of how open source communities and components available today, such as OPNFV and OpenDaylight, can be leveraged to construct an effective and viable VCO solution. In fact, this architecture was shown to work in just three weeks after assembling the hardware and software stack in a data center and demonstrated live at the 2017 OPNFV Summit in Beijing. It highlighted the Virtual Customer Premise Equipment (vCPE) use case for both business and residential services including service assurance. See the demo brief for details specific to the demo.

This paper explains in detail the industry trends leading in this direction, the need for this reference architecture, and the architecture details as discussed and agreed upon by several CSPs and vendors using open source hardware and software components. These include the onboarding of VNFs, subscribers and the delivery of content and services through OpenDaylight, OPNFV, OpenStack, Open Compute Project (OCP), and open source orchestration software for residential, business, and mobile services.

Services Delivered via Virtualized Central Office

A Virtual Central Office architecture must be capable of delivering the following services.

- Residential Services – Full suite of residential services with transport and services elements, supporting a virtualized service edge and different models of R-vCPE functional partitioning
- Mobile Services – Build a virtualized mobile environment supporting virtualization of both the RAN and the core
- Enterprise Services – Build an Enterprise Services model for both managed and unmanaged services using the thick and thin Enterprise vCPE paradigms

Goals of Virtual Central Office (VCO)

The goal of the Virtual Central Office project is to produce an OpenDaylight-based reference architecture that, when combined with other functional elements (such as NFV and Orchestration software stacks) can support the delivery of Residential, Business and Mobile Services.

NETWORK AND TECHNOLOGY TRENDS AND THEIR IMPACTS ON VCO

As the name implies, the primary target locations for the VCO are “Central Offices,” which in the context of the telecom network operators are the locations that either directly or indirectly (i.e. through the access aggregation networks, with the equipment further distributed on outside plant and/or customer sites) enable the subscriber access to network services.

Figure 1 below depicts the primary target location of the VCO in the context of the primary network access technologies.

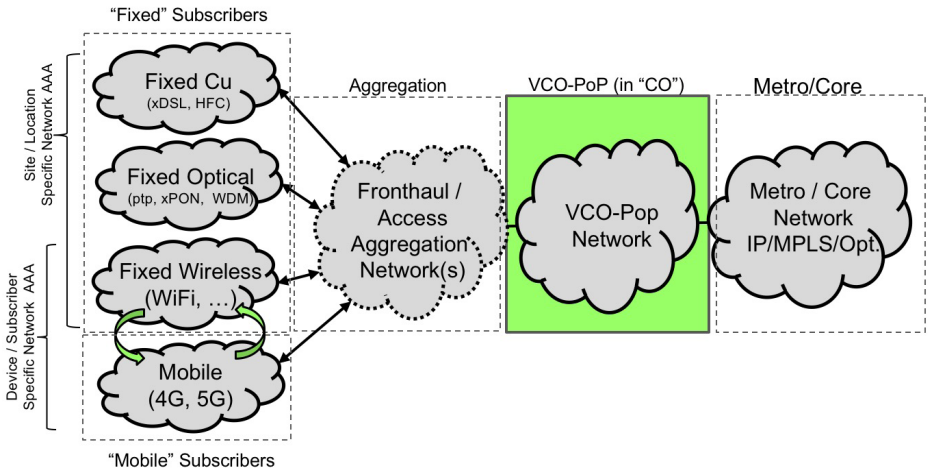

Figure 1 – VCO target location in the network context

VCO locations typically use multiple access technologies to serve diverse customers in residential, enterprise and mobile market segments, but this is not a requirement from the VCO design perspective. The goals of the VCO include simplification, convergence and integration of the access network technologies within a common architecture, as well as enablement of service delivery models that span different access types.

The common architecture of the VCO is enabled by leveraging a combination of ongoing trends, specifically:

- SDN (control and data plane separation)
- NFV (virtualization of network functions, by decoupling hardware from software and extending software-centricity via open protocols and APIs)
- Service and network orchestration; self organizing, autonomous/autonomic management
- Open protocols, interfaces and APIs, preferably with open source implementations
- Open information / data models, combined with standardized modeling languages (e.g YANG) and open source tools
- Open source software, and
- COTS and/or open, multi-source “white box” HW (for networking, compute and storage)

Figure 2 below depicts the “tiered” structure from access to the core data centers. The optimal location for the VCO points of presence are expected to be at the consolidated central office locations (VCO-C), which are close enough to terminating access network equipment to allow economical servicing of the associated subscribers within reach of a passive fiber plant, and simultaneously being able to meet the target latency constraints of C-RAN deployments and/or MEC services. However, it is also possible to deploy VCO PoPs in more distributed manner in Tier 1 CO locations, resulting in increased number of smaller sites. (VCO-D)

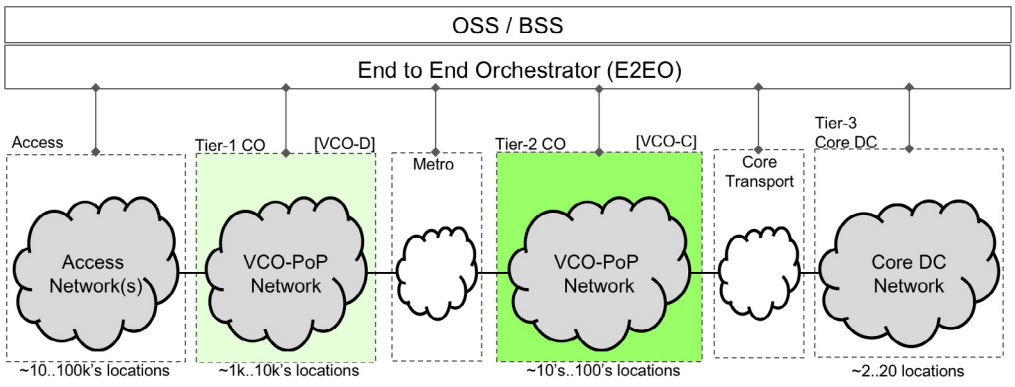

Figure 2 – Tiered core data center access structure

Historically, telco central offices have been located based on the constraints on the physical copper loop lengths, which generally means that the number of locations is determined primarily by the topological constraints imposed by the maximum distance (typically 2-4 km depending on the operator / deployment), and size of the office is determined primarily by the type of the area served (business, residential, mixed use etc.) and associated population density (urban, suburban, or rural). Fixed-mobile operators have often used CO locations to serve both access architectures by consolidating associated equipment to these locations. Wireless-only operators do not necessarily have copper loop constraints, but depending on the access network architecture, they may have latency constraints that set the maximum fiber span length between CO location and antenna sites, especially in C-RAN deployment architectures.

The access portion of the network is typically based on L1-L2 services, implemented either as physical point-to-point topology from the CO to the subscriber in dedicated medium cases (copper loops, pt-pt FTTx), or “logical” point-to-point topology from CO to subscriber in cases of shared medium and/or remote active element based architectures (copper loops from OSP/buildings, xPON, HFC). Traffic separation in logical point-to-point topologies over shared media are generally implemented with some technology specific L2/tunneling mechanism (ATM, VLAN, VXLAN, MPLS, various L2VPN tunnels, etc.). Typically, individual subscribers cannot directly communicate with each other using access network facilities, and are not able to receive or access other subscribers’ traffic (either implicitly due to topological design, or explicitly through use of encryption in shared-medium cases).

Regardless of the access technology, the common goal is to increase the access bandwidth per subscriber, while maintaining or decreasing cost levels.

SDN- and NFV-based architectures, combined with orchestration-driven, automatic management, are seen as key enablers for cost reduction (both CapEx and OpEx). Further cost reduction opportunities are expected from migrating to open source-based software and open / “white-box” hardware designs. From a network architecture perspective, these opportunities arise from location consolidation and sharing, as well as simplification of the access networks through reduction of the service- and technology-specific elements and migration towards simpler, more passive outside networks.

Since video traffic already dominates IP network traffic (between 70 and 82% of total depending on the estimate, with one-way streaming the large majority of total), optimization of video delivery costs is a key consideration, regardless of the access network technology. Mechanisms such as local content caching at the (IP) edge and

multicast are expected to be key technologies required to reduce growing, video-driven costs. These mechanisms are likely to drive providers to distribute the IP edge to consolidated access COs, as the backhaul network between access and IP edge then becomes an intra-building fiber.

The key trends on "wireline" access networks, regardless of the specific architecture/technology used are:

- All services are being migrated have been migrated to common IP/MPLS/optical based metro/core network, possibly with legacy service adaptation/interworking functions located in COs
- Reduction of the Cu loop length to enable higher access bandwidth through the Cu plant - implies that the Cu access equipment must be moved to the outside plant and/or buildings, and connected to CO locations through fiber based infrastructure
- Elimination of the Cu plant, and migrating/building fully fiber based access network (FTTH, FTTB), possibly with short Cu drops to connect to user equipment in the endpoints (through Cu Ethernet, MoCa or such interfaces)
- FTTx architecture with wireless (only) access for users—"fixed wireless" type services; this fully eliminates any Cu portion from the access network through shifting the service access to wireless only
- Fiber architecture migration with no plant level Cu use enables fully passive outside plant (i.e. no powered/managed equipment required in outside plant)
- Fiber architecture migration also allows longer loop lengths (typically 20km from CO for xPON architectures), which in turn enables potential retirement of majority of the COs (this is called CO consolidation)

The key trends on the mobile access networks are:

- All services are being migrated have been migrated to IP, and utilize common IP/MPLS/optical metro/core network
- Historically the mobile core network elements are located in more central locations, instead of COs (most extreme case being two core function locations for the whole network)
- Expectation of continuing reduction of cell sizes, resulting in a substantial increase of the cells in a given geographic area, especially in dense urban/sub-urban deployments

- Move from distributed RAN towards CRAN (i.e. move the base station based RAT signal processing to more centralized resource pools using CPRI and in future NGFI/nFAPI)
- 5G targets ultra-low latency services, which require support of service deployments from locations that are physically relatively close to edge (often quoted 1ms target for UE-RAN portion)—this is one of the key goals of "Mobile Edge Computing," but in addition to MEC and low-latency services, low latency is also a requirement for CRAN features
- Low latency, edge services, and CDN enablement all drive the trend of forcing the mobile packet core datapath functionality and associated network elements to be increasingly de-centralized, e.g. through dual-anchoring so that non-IMS traffic can be terminated much closer to the subscriber

The key implications of these trends in the context of the VCO architecture are:

- Locations should be as close to customer as allowed by latency constraints, fiber availability and fault-domain capping, but no closer (moving the locations closer than necessary negates the benefits from the resource pooling, location consolidation, CDN cache hit rates, etc.)
- We use the 200us one way latency and 20km fiber spans as suggested in COMBO project deliverables [COMBO] as basis of dimensioning, resulting on 5x-10x increase potential of the VCO location—customer distance, as compared to "traditional" Cu plant optimized locations.
- While distance increase enables the decrease of the locations, how it translates to exact reduction ratio is highly dependent on the fiber deployment topology rather than simple geometry. We use the 10x decrease ratio as target for the expectation on location consolidation opportunity, based on the multiple operator's reported reduction scenarios
- 10x consolidation ratio directly translates to 10x larger deployments in per-location, and associated increases on required resource pool sizes (i.e. 10x more subscribers, 10x more access network interfaces/traffic, etc.)
- Access networks represent huge capital expenditure for operators, and evolve slowly—this implies that anything new needs to support incremental deployments in "brownfield" scenarios, and that some level of legacy compatibility will be required.

- In addition to above, while equipment replacement can be generally justified in more aggressive schedules, the physical plant represents major cost, and its state (in terms of e.g, fiber availability) and topology are major drivers on what can be economically feasibly deployed (e.g. it is likely that use of WDM technologies, increased data rates etc. represent lower incremental cost than deploying more new fiber in fiber-poor scenarios, even if the associated endpoint costs are higher)
- For wireline access, xPON (including WDM PON), pt-pt fiber, and fiber rings (with or without WDM) are expected to be the key interfaces - we do not consider Cu access enablement from CO locations to be a requirement for VCO (other than ability to be able to interface to legacy Cu serving CO/OSP/building located access multiplexing equipment)
- VCO architecture must be able to support 4G/LTE and 5G CRAN deployments; note that the fronthaul interface may change from CPRI to NGFI/nFAPI for 5G deployments, and there may be more functional split/fronthaul options for 5G
- Mobile network support requires support of the tight timing constraints (phase/time and/or frequency) which means that the underlying fabric must support the associated timing transfer mechanisms. 5G requirements are currently under study and are dependent on the base station functional split on CRAN. For now, we expect that requirements as detailed in PTP telecom profile are needed.
- Consolidation of the access technologies enables new opportunities for further simplification of the network, while simultaneously potentially enabling new/optimized service models. COMBO project deliverables describe some of such opportunities.
- Another factor on fixed-mobile consolidation opportunities is trend on the reduction of the cell site sizes on mobile, and the desire to move towards “wireless local loop” architectures in wireline side to eliminate the need for operator supplied/installed per-subscriber equipment
- For TDM based services, we do not consider that they are required to be addressed by VCO, other than ability to interface with any associated legacy equipment
- For mobile networks, only 4G/LTE based architectures are considered, with expectation of focus on enablement of 5G networks in near-term
- Interfaces to core/metro networks are expected to be based on IP/MPLS over fiber (with or without WDM or other optical transport equipment)

Summary

Central Office plays a critical role for a telco, as it is the gateway or the primary interface to end customer where the user services are hosted and terminated. It is home to the copper/fiber termination and now the home for mobile gateways, Radio Access Network (RAN) termination components and hosted services.

Central Offices are characterized by the following:

1. They terminate a wide variety of subscriber line technologies, from copper/DSL to fiber GPON and in some cases Cable/CMTS. They are also used for placement of Mobile Gateways due to their proximity to users/subscribers.
2. Central Offices follow a strict hierarchical model for connectivity with strict serving areas and subscriber bases following rigid models.
3. Central Offices serve a variety of access speeds, from Kbps for low power, narrow band connections to very high speed (multi-gigabit rates) connections for business customers.
4. Wide variety of hardware, such as switches, routers, gateways, and peering devices
5. Central Offices have many bare metal servers
6. High source of CAPEX and OPEX for a telco
7. There may be hundreds to tens of thousands of central offices in a given country

In summary, Central Offices today are complex, rigid and have a mix of technologies that are neither uniform nor simple.

On the other hand, virtualization is fast becoming the norm for Telecommunications Services Providers as they build new services and scale existing ones. Given the complexity of Central Offices and the problems Virtualization can solve, Central Offices are ripe for virtualization.

Previous attempts to virtualize the Central Office have resulted in a proprietary or a limited approach by building components (SDN controller and Orchestrator) from scratch to address specific issues. With recent developments in orchestration tools and advances in the functionality of OpenDaylight SDN controller, capabilities now exist to build a complete Virtualized Central Office.

Broadly, Central Offices have two categories of functions: user-associated functions and service-associated functions. Subscriber Line Termination and Subscriber Services Management are examples of user-associated functions. Service-associated functions are functions that define or describe service characteristics such as QoS, routing functions, and Mobile Gateways.

User Associated Functions	Service Associated Functions
Subscriber Line Termination • Copper Line Termination • Optical Line Termination - PON/GPON • Hybrid Fiber Co-ax • Cable Modem Termination Services	**Wireline Service Functions** • Residential and Business Services - Self Service Portal • IP Routing Functions • QoS Management • Storage and Usage Quota Management • Video and Content Cache • Local Information Services • Mail and File Store • Subscriber Value Add functions
Subscriber Management Capabilities • Subscriber Gateway Functions • Authentication and Authorization • Event and Subscriber Information Logging	**Mobile Service Functions** • Mobile Gateway and Switching Center • Traffic Offload Gateways • Access and Aggregation Switching/ Routing equipment

Table: Central Office based services and functions

USE CASES

Virtual Central Office is used to deliver Residential, Enterprise and Mobile services to consumers or subscribers. Each service is specific to a customer/ subscriber base with distinct requirements and characteristics.

The reference architecture developed by the Virtual Central Office project intends to address the following use cases

Residential Services Use Cases

- Internet Service
- Email and web hosting services
- Personalized content
- Video on Demand
- Access to OTT services
- VPN service
- Any device, any access
- Bandwidth guaranteed services
- Voice over IP with E911
- Quota management
- Parental controls
- Game servers and services
- Security
  - Email scrubbing
  - Personal firewall
  - Virus scanning
  - Malware protection
- Data backup from NAS or home computers and devices
- Value-added services
  - Home monitoring
  - Community/shared portal
  - DropBox shared storage

Enterprise - Business Services Use Case

The following Enterprise or business use cases are addressed with the Virtual Central Office architecture.

- **vCPE - Thin CPE model.** In this model there is a network interface device (NID) to terminate connections at the enterprise site, and all services are provided from the Virtualized Central Office. The role of the NID or the Thin CPE is to encapsulate and send all traffic from the site to the VCO. This model is sometimes also referred as the "Cloud CPE" model. The Thin CPE model may also be used for residential applications.
- **vCPE - Thick CPE Model.** In this model there is a compute node at the enterprise site that runs localized applications and services. This compute node also runs virtual network functions such as firewall and vRouter. These functions are orchestrated from the Central Office tools. These VNFs may be service chained together or to other VNFs residing in the VCO. This use case is the basis of the branch office consolidation scenario as shown below in the figure below.
- **VNFaaS -** VNFs can be hosted at the VCO and/or the "thick" vCPE, and provided as a service to the enterprise via managed or unmanaged services.

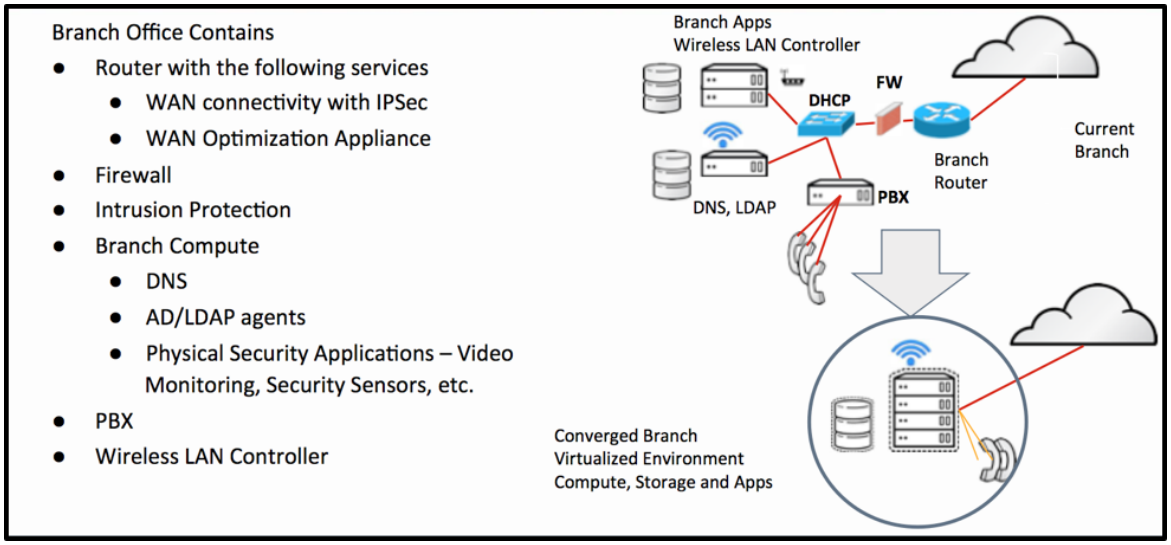

Figure 3 – Example Branch Office Configuration - Before and After

Branch Office Simplification

Branch office architecture today consists of many components (nodes and functions). A typical branch office IT infrastructure may contain some servers for branch office applications, a wireless LAN controller (WLC) to control/manage WiFi access points, a firewall, an IP PBX system (VoIP Call Manager), video and security functions, an IP Router, VPN and some DHCP, DNS and LDAP capability. Some of these capabilities, such as firewalls and routing, may be part of managed services (provided by a carrier or other third party) , while others are managed by the organization's corporate IT, such as applications and servers.

All of these capabilities can be easily collapsed onto a few servers by virtualizing them and running them as virtual machines or containers on the server. Through the use of Network Function Virtualization (NFV), the branch office architecture can now be provided via a "branch-in-a-virtual-box" solution.

Mobile Services Use Case

The following are examples of use cases for Mobile Services can be hosted at a Virtualized Central Office:

- **Network Slicing (5G Use case).** Provide a slice of the mobile infrastructure for a Virtual Network Operator or for a tenant
  - For mobile virtual network operator
  - For public safety and security
  - For enterprise private mobile networks
- **Cloud RAN or Virtualized RAN.** Ability to virtualize the radio components for centralized control and better utilization of resources. (Refer to details on C-RAN).
- **Multi-access edge computing.** Placement of compute nodes as close as possible to the edge of the network or close to the data acquisition points: sensors, mobile devices, virtual reality devices or self-driven vehicles
- **Machine-to-machine communications -** Build vEPC based networks for rapid scaling to address use cases such as connected car, smart city infrastructure
- **VoLTE and IMS -** Deploy Voice over LTE and IP Multimedia Subsystem for voice and video services including presence
- Seamless Mobility
  - Roaming from WiFi to LTE and back
  - Roaming from Enterprise to Macro networks and back
- WiFi Calling
- Increasing Coverage Deployments using small cells and U-LTE, using unlicensed spectrum

GENERAL REQUIREMENTS FOR VCO

The services described in the previous section translate to the following requirements for architecture or products deployed at the VCO.

- Public or Private IPv4 and IPv6 Addressing
- Unicast and Multicast
- IPAM - DHCP, DNS, etc.
- Subscriber Management and Policy
  - AAA (Authentication, Authorization and Accounting)
  - Service lifecycle management (from activation to de-activation)
  - Self-Management services (may be in form of management portal access - possibly linked to subscriber authentication/authorization state)
  - Quota management (access rate enforcement, based on service tier)
  - QoS
  - Mobility - Any device/Anywhere
- Configuration of encapsulations such as VLANs, VXLANs, GRE and Keyed IP Tunnel/L2TPv3
- Time Service (NTP)
- OAM and troubleshooting
- CPE/NID remote test
- Transparency of data path
- Device management
- IP Quality of Service to handle Voice, Video and Data
- Video and content cache for better quality of experience
- VPN capabilities/pass through to services or even home devices
- Virtual Network Functions: Router, Firewall, IP Services, DPI, Policy Enforcement and bandwidth control functions
- SIP, SIP Proxy for VoIP and Video.
- Service chain configuration and management
- Virtual Network Function Configuration and Management
- Security Configuration and Management
- Routing configuration and management

The following transport types must be supported with the above needed functionality.

- PON Access (GPON, XGS-PON, NG-PON2 etc.)
- HFC/Cable-Coax Access (DOCSIS 3.0, 3.1 etc.)
- DSL Access (ADSL2+, VDSL2, G.Fast etc.)
- LTE Access

The following Subscriber access link termination technologies must be supported.

- PON (GPON, XGS-PON, NG-PON2 etc.) access with vOLT and vBNG
- Cable Access with rPHY and vCMTS (vCMTS has both control plane/data plane and is approximately functionally equivalent to the combination of vOLT and vBNG)

Virtual CO must have the ability to terminate all the above access types. In order to scale horizontally some of these access termination devices can be directly connected to the leaf Spine Fabric—such as a Virtualized OLT (Optical Line Termination Device) or a Virtualized CMTS, whereas others may be terminated via a gateway function residing as a VNF or a dedicated node connected to the Spine switch.

Tenancy Requirements

Residential services are generally hosted as a single-tenant environment with free access between users and services separated only by an administrative region. In case there is a need for multi-tenancy because multiple regions are hosted in the same DC or virtual network operator models are used then there is no requirement to share VNFs between those “Virtual Network Operators” or “Regions.” Hence they operate as a single tenant each from the DC or fabric perspective.

However, for business or enterprise services, VCO must be able to host multiple tenants on the same infrastructure. This implies the virtualized infrastructure layer must be capable of dealing with private tenant addressing, keeping them separate without violating security policies, jeopardizing provider routing and provide virtual and private networks and services. This also means VCO must be capable of hosting multi-tenant virtual machines.

High Availability Requirements

In addition to supporting redundancy in hardware components, any virtualized software components must make use of the hardware redundancy to provide a highly available environment.

The following HA requirements must be supported.

- Non-blocking fabric design, using a Clos fabric to eliminate bottlenecks in the system
- Redundant components for the hardware, servers and switches—such as power supply etc
- System redundancy to support Service Availability targets of >99.99% and higher
- Live software upgrades and patching without affecting service
- Clustering capability for
  - SDN controller - OpenDaylight
  - Other orchestration controllers
- Hot-swap of hardware components, such as disks for storage
- RAID, striping and other mechanisms for data storage redundancy
- Backup and restore capabilities

While some of these requirements are out of scope of this architecture, such as backup and restore, disk array configuration etc., they are listed here for completeness.

Specific Requirements for Mobile Services

The services listed above result in the following requirements for delivering Mobile Services from a Virtual CO.

- vBBU or Hosting of the Remote Radio Head for C-RAN applications using CPRI or NGFI/nFAPI.
- Hosting of Virtualized Mobile Packet Core Components - PGW, SGW, MME, SGSN, GGSN, HLR, VLR etc
- Hosting of Virtualized SAE Gateways - vHNB, ASN Gateway
- Hosting of Virtualized IMS components (x-CSCF, TAS, SBC, BGF etc) for Voice and Video over LTE
- Multi access Edge Computing infrastructure & MEC services hosting
- MEC Gateway for IoT/M2M and Edge/Fog Computing models
- Virtualized Subscriber Services Components or (S)GiLAN components such as Traffic Optimization, Caching, firewall services, DPI etc
- Placement of Mobile and Radio Management tools for SON
- MVNO Components - Gateways, Billing tools etc.
- Roaming and Peering Gateways
- PSTN Handoff using virtualized SBC

SDN Controller Requirements

Given the requirements stated above in the use cases, the SDN controller (in this case, OpenDaylight) must have the following capabilities.

- Topology discovery of connectivity internal and external to the system, including underlay and overlay topology management
- Configuration of switches and routers in the data path at the Virtual Central Office through APIs, based on policy driven by discovery/change events
- Configuration of the underlay fabric (leaf and spine switches), and configuration of fabric (leaf/spine CLOS), initially using EVPN, using Openflow or LISP in future
- Traffic placement, balancing, and fabric resource use management
  - Configuration of QoS
- Configuration of overlay service chaining using any of the available methods: NSH, Segment Routing, Neutron Port Mapping or VLAN and VM Stitching
  - Configuration of VNF FG at the enterprise site and at the VCO
- Management of all southbound interfaces - BGP, OpenFlow, LISP, OVSDB or NETCONF for configuration of physical and virtual network elements (including VPP)
- Configuration of Open Virtual Switch (OVS) and/or Vector Packet Processing (VPP)
- IPv6 Addressing - ability to allocate, translate and process IPv6 addresses to subscribers depending on the Service Providers addressing schema
- Public, Private IPv4 Address - Usage of public/private IP address for the infrastructure and for subscribers
- Configuration of encapsulations such as VLANs, VXLANs, IPSec, GRE and keyed IP tunnel/L2TPv3
- Configuration of security policies
- Configuration of vRouter and routing policies
- Configuration of multi-tenancy environments

Service Assurance and Reporting Requirements

- Monitoring and assurance services for performance, fault tolerance, control and management; requires reliable, low latency asynchronous event and telemetry (time series data) interfaces and services support for all managed entities
- Services scale requires configuration of compute and network components to address dynamic load in a VCO. While the compute can be instantiated (initially via Openstack but longer term also via e.g. Kubernetes), network components must be ready and controllable via the ODL controller.

Orchestration Requirements of Virtual Central Offices

This architecture whitepaper is intended to demonstrate the ease of integration between the many components required to create a network service with readily available open source components, including OpenDaylight. Furthermore, orchestration of virtual devices must demonstrate significant improvement in service agility as compared to physical devices. For the purposes of this document, the term "orchestration" focuses on the ability to manage the deployment lifecycle of virtualized functions that comprise a network service, including the ability to manage network connectivity between virtualized functions and connectivity to legacy devices. The following are the detailed requirements for the orchestration layer.

- The components should be compliant with standards, including MEF, ETSI, and TM Forum
- The orchestration layer must be capable of fulfilling all the use cases detailed above; this implies that the orchestration layer should be agnostic to the type of VNF being managed
- The orchestration layer should be capable of managing the operational lifecycle of virtualized functions, such as:
  - Onboarding – the ability to manage and maintain a catalog of Virtual Network Functions (VNF) and virtual Network Services (NS)
  - Instantiate - the ability to dynamically deploy VNFs or NSs, including the ability to automate the deployment of the VNFs and NSs
  - Scale – similar to instantiation, the ability to dynamically scale a VNF or NS, and the ability to automate the scaling of the VNF or NS in a proactive (e.g. capacity planning) or reactive (e.g. network demand spikes or abatement) manner
  - Heal – the ability to automatically react and respond to faults and SLA degradation
  - Terminate - the ability to stop VNFs and NSs and recover virtual resources
- Networking is a key component of service provider services, therefore the orchestrator must be able to facilitate advanced networking capabilities such as VNFFG and SFC to enable service differentiation
- The orchestration layer must be capable of managing across locations and administrative domains in order facilitate both centralized (e.g. residential vCPE) and distributed (e.g. enterprise vCPE) deployment models

As automation is a key requirement of orchestration, the following integration requirements are also noted:

- Ability to integrate with northbound OSS/BSS or end-to-end orchestrators in order to automatically construct an end-to-end service, from customer request to order fulfillment
- Ability to integrate with service assurance and/or analytics functions to automate instantiation, scaling, and termination based on proactive or reactive metrics

OPENDAYLIGHT-BASED VCO ORCHESTRATION ARCHITECTURE

Components and Overview

The VCO orchestration architecture below is aligned with/based on the architectural frameworks specified in in the following standards organizations:

- ETSI – ETSI NFV Management and Orchestration
- ETSI – ETSI Generic Autonomic Network Architecture (GANA)
- TM Forum – TM Forum ZOOM
- MEF – Third Network Service Orchestration

The overall management architecture is influenced by the operator requirements / reference architectures for the management and orchestration SW stack, such as ONAP (a merger of AT&T's ECOMP project and Open-O), Verizon SDN & NFV reference architecture, and SKT's COSMOS and ATSCALE architectures. The goal is to align the functionality and relevant lower level APIs with the emerging open source orchestration project alternatives as much as possible, including OSM and ONAP orchestration stacks.

The overall software stack and the dependencies between the associated management, orchestration and control components are shown in the figure on the next page.

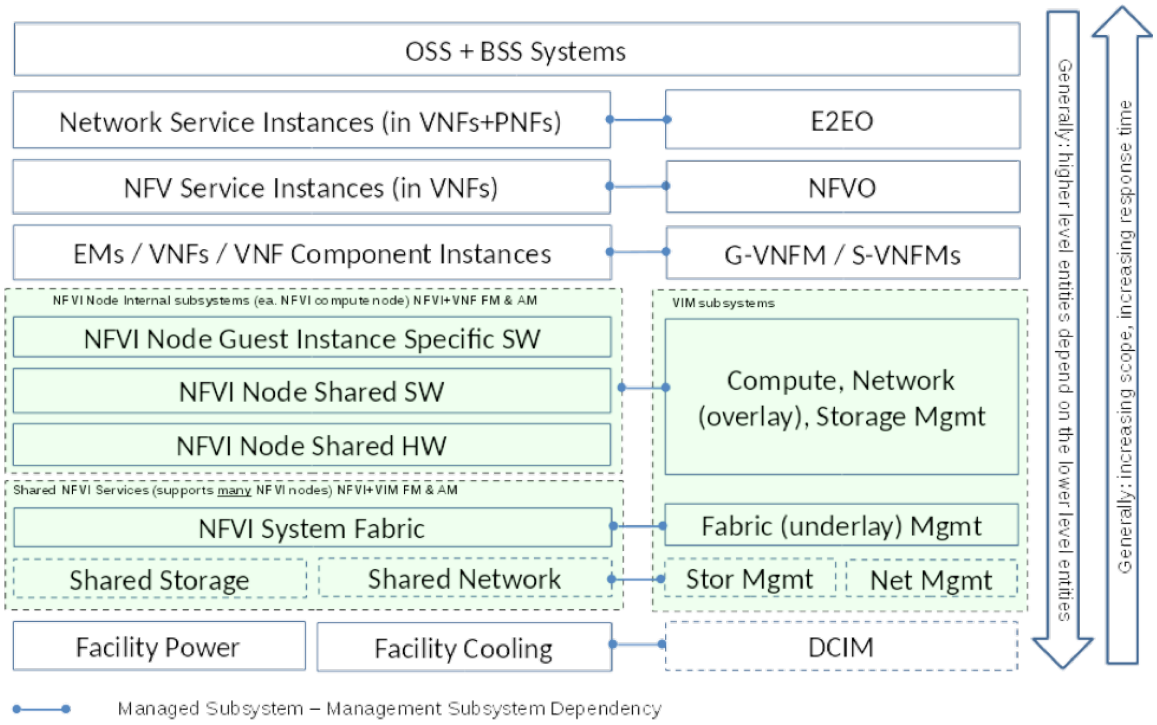

The high-level orchestration and management relationships in the context of the network-level deployments are represented by the figure below. OSS/BSS, E2EO and NFVO are expected to have network-level views, and therefore need to provide the interfaces to all assets within and up to the single operator's network domain. Orchestration across multi-operator domains is not considered as part of this document. This is not meant to imply that the implementation of the associated systems is fully centralized; in fact, for resilience and availability reasons complete centralization would be undesirable.

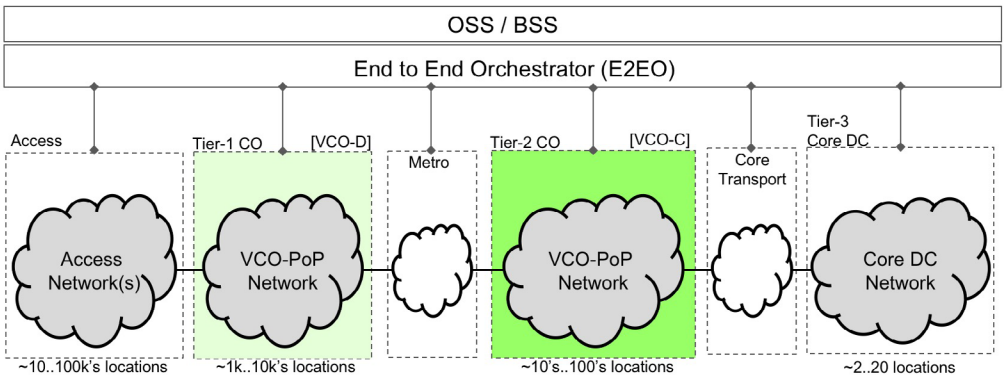

The VCO and associated management software stack is expected to build on the NFV management stack. It includes the associated functional components and reference points in the ETSI NFV reference architecture (NFVI, VIM, VNFM, and NFVO). In addition, VCO clarifies the role of the SDN controller in the management stack, and addresses the interfaces required to manage the physical NFVI fabric, as well as the decomposed physical network elements required to implement the VCO functionality.

For end-to-end service orchestration, the full MANO stack should include orchestration functionality that supports a holistic view of the network elements and services required to support business services. The architecture provides for the MANO to be loosely decoupled from the specific implementations of physical or virtual devices and interfaces that are used to control and manage the associated network elements, regardless of whether such interfaces are implemented through SDN controller(s). This broader management function is referred as End-to-End Orchestration (E2EO) in the above and subsequent pictures.

A more detailed diagram of the NFV reference architecture per ETSI is given in the figure below.

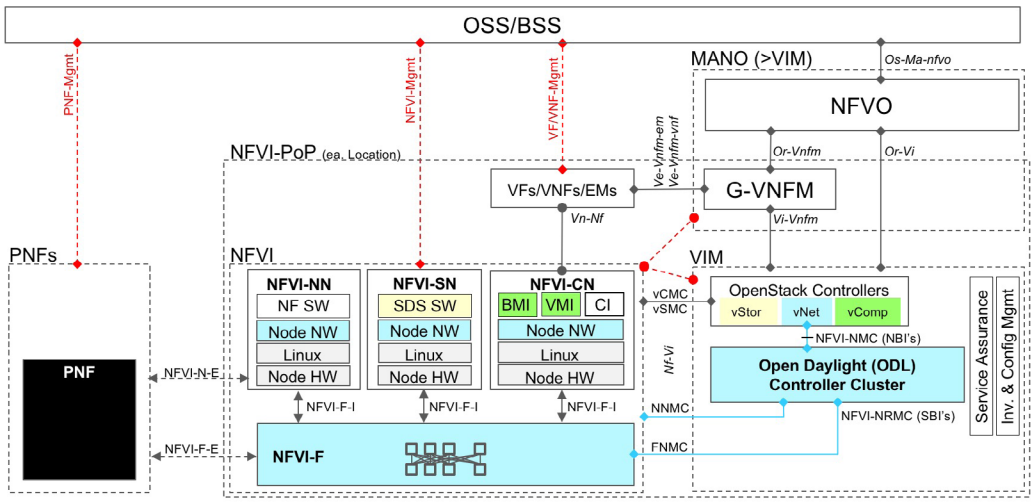

In most implementations, the goal is to migrate the functionality and interfaces from OSS systems towards automated orchestration systems. While this transition is expected to be long, the target architecture would look approximately like the following picture with respect to interfaces from physical elements and NFV towards orchestration and OSS/BSS systems. Note that in a realistic implementation, interfaces from some elements will experience an extended time period during which some components and subsystems are connected to OSS and some are connected to orchestration systems.

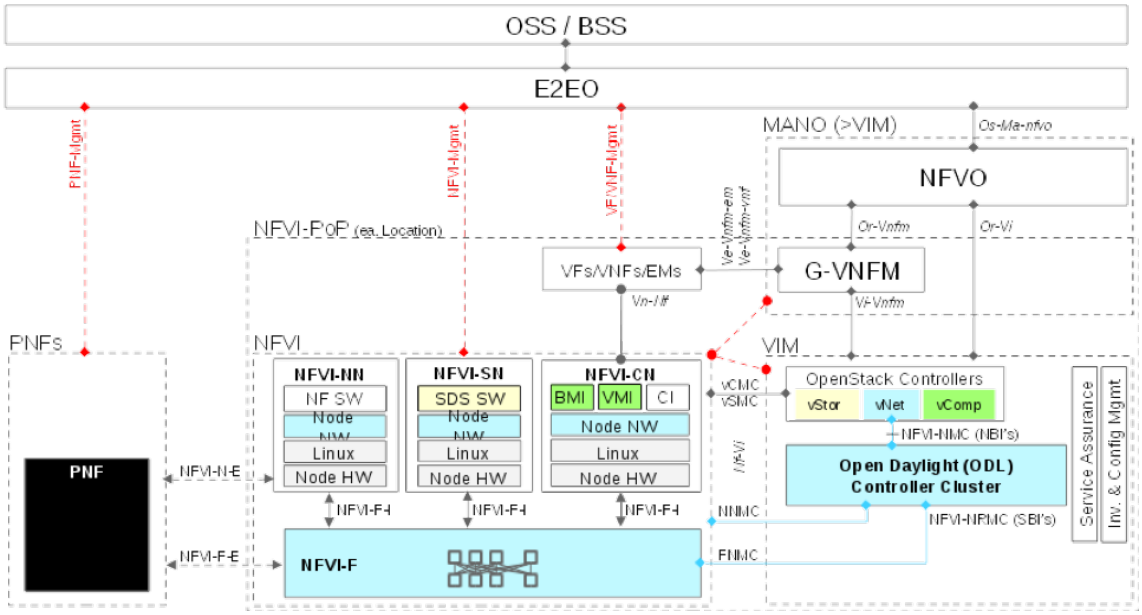

When implementing the SDN controller, the objectives are to support control and management of the NFVI/VCO PoP physical (“underlay”) and virtual (“overlay”) networks, as well as other the decomposed physical and associated virtualized elements. The controller is expected to expose the associated configuration and management functions through its northbound API(s), and support associated controller-resident services.

NFV Domain

NFV domain orchestration is expected to adhere to the ETSI NFV Management and Orchestration (MANO) reference architecture and associated specifications.

The NFVO is expected to manage all aspects of the virtual domain, including:

- All lifecycle events of VNFs, such as:
  - Onboarding
  - Instantiation
  - Scaling
  - Management (query, monitor)
  - Termination
- All lifecycle events of virtual Network Services, such as:
  - Interconnection of VNFs into a multi-Network Function service
  - Creation and maintenance of VNF Forwarding Graphs
- Resource orchestration: interfacing with compute, storage, and network controllers to allocate resources for the Network Service
- Interfacing with networking domain orchestration and other domain orchestrators for advanced use cases
- Interfacing with northbound systems such as customer portals, customer orchestration, OSS/BSS, and health monitoring systems over the Os-Ma-nfvo interface for automated service creation and service assurance

VCO Node Types and Roles

As described before, the VCO builds on the NFV reference architecture. Therefore, it implicitly includes the various nodes required to implement NFVI and associated management functionality, specifically:

- VIM controller nodes (including nodes hosting VCO-PoP local SDN controllers)
- Physical fabric nodes
- Network nodes, regardless of role and implementation ("GW" nodes, LBaaS nodes etc., whether physical, based on specific SW instantiated in general purpose compute node, or hybrid, i.e. "compute-like" node with specific HW acceleration functionality)
- Storage nodes (compute nodes with storage resources and SDS software)
- Compute nodes (general purpose pooled compute resources)
- Compute nodes with specific acceleration HW support (otherwise generic, but may affect desirable role assignments and/or eligibility for the guest use)

In addition, the following node type/roles are identified in the context of the VCO requirements:

- Access network interface nodes (xPON, xDSL, switches with longer range interfaces,...)
- Core network interface nodes
- Service specific support nodes (e.g. CPRI interfaces)
- Timing source nodes (e.g. GPS with 1588v2 grandmaster and/or NTP support)

All VCO node types, regardless of the functionality are expected to interface to the VCO fabric. Also, it is preferred to support direct interface from the VCO fabric to external physical network elements, even if they are not otherwise directly integrated in VCO system.

Virtualization Technology Alternatives

The following figure depicts the key potential alternatives that can be identified in support of the virtualization, packaging and management of application component instances that can be made available in the compute nodes.

The expected priority list is based on the current state, maturity and applicability of these in the context of the NFV/networking applications, as follows:

1. Bare metal
2. Virtualized (application component instances in VMs)
3. Containers in VMs
4. Containerized (application component instances in containers)
5. Mixes of Virtualized/Containerized in the same system (in different hosts)
6. Mixes of Virtualized/Containerized in same host
7. VMs in containers

Note that the bare metal nodes can be used as a basis of the application deployment and lifecycle management processes of the physical nodes, or nodes with specific hardware (such as switches), as long as they can be sufficiently identified to be able to assign the appropriate roles and deploy associated SW images to them. It is expected that at least in the near term, the control processors embedded in such integration candidate nodes needs to be IA based, and they need to otherwise "look like a server" with additional application specific HW from an integration perspective (e.g. no support for ARM, PPC, MIPS or other architectures).

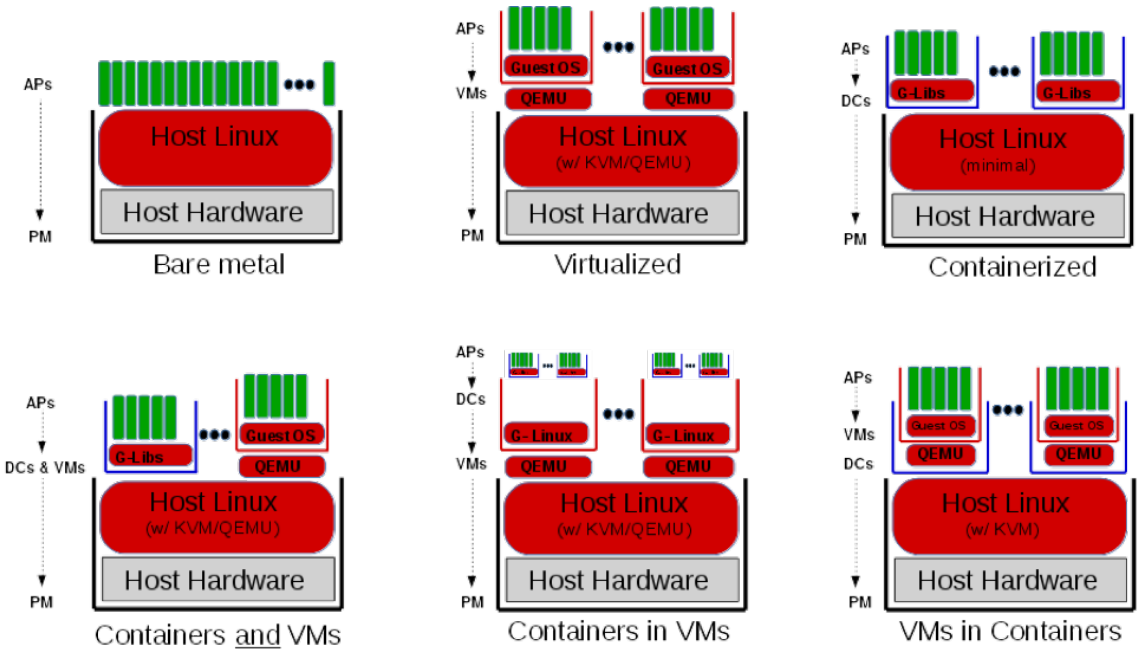

PM=Physical Machine, VM=Virtual Machine, DC=Docker Container, AP=Application Process

SDN Control Domains and Their Relationships

SDN controllers can range in scope, the breadth of which affects their southbound and northbound interfaces, interactions with each other, as well as the entities that needs to be involved in the E2E service orchestration and control-related operations.

A common design separates the control domains from the underlying network segments, and associates one set of controllers to each of the resulting domains. Higher-level network orchestration and control functions must interface to these separate network control domains and their respective network segments, and connect the network segments, interfaces and services on separate domains together to form and manage end-to-end services. This architecture is depicted in the figure below in the context of the multi-tiered network topology discussed earlier.

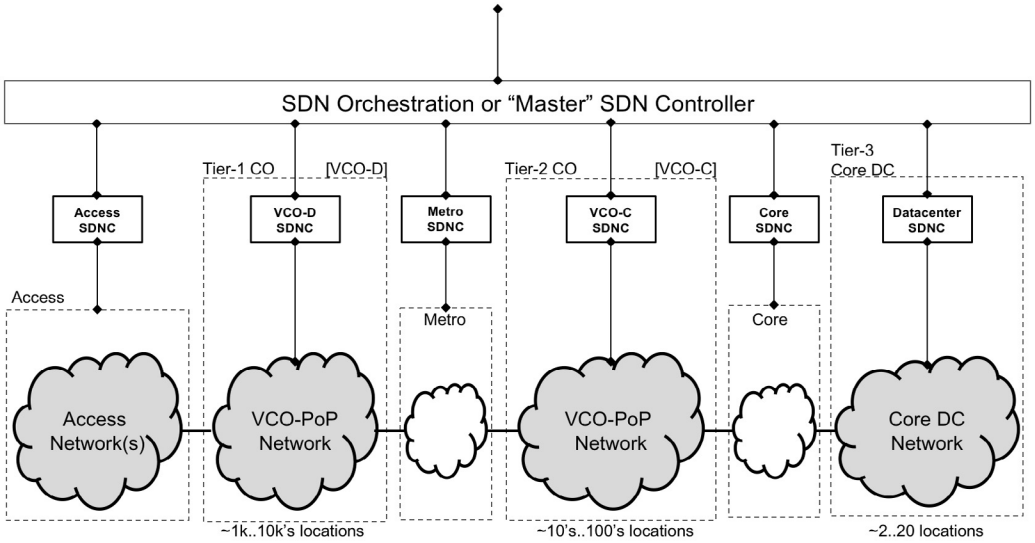

There are various consolidation options available to reduce the number of controllers and control domains. For example, the control functions related to preceding network segment can be associated with the next higher tier controller cluster. In such an architecture, the number of control domains is reduced, while the scope of the each control domain is increased, as shown on the next page.

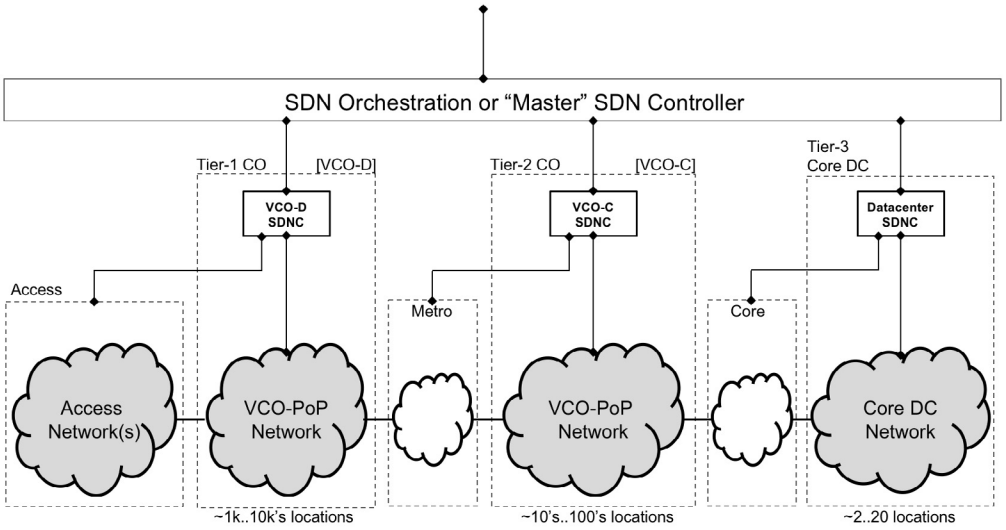

Various hybrid control topologies are also possible, where the controllers for the domains are kept separate, but are connected hierarchically (e.g. access network controller has a separate controller cluster, but instead of connecting directly to network level controller/ SDN orchestration function, it is connected through the associated VCO-POP/VCO-POPs SDN controllers).

Since one of the main goals of VCO is simplification of the access network (up to the fully passive optical access network architecture), we need to consider the impact of decomposing the SDN control domains and associated interfaces. The figure below depicts various entities that could potentially be controlled by the VCO SDN controllers.

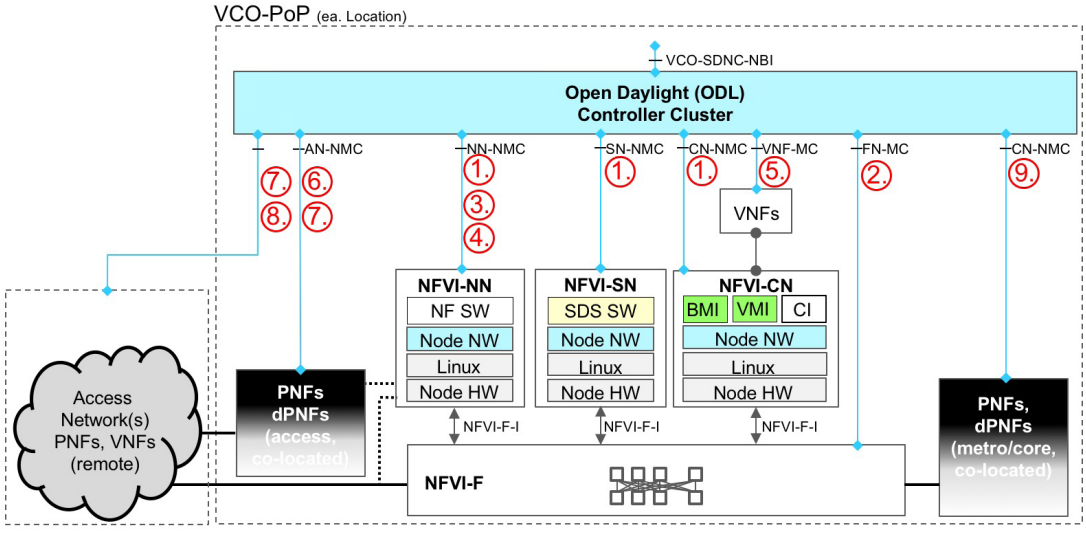

The following potential reference points can be identified, and their associated implementation priorities (based on frequency of use in VCO use cases) in the context of the overall VCO objectives are discussed on the next page.

1. NFVI nodes (compute, network, storage) “overlay” network
2. NFVI fabric “underlay” network
3. NFVI network functions control (gateway, LBaaS, FWaaS,...)
4. Edge network functions control (network functions that are not considered to be NFVI network services, e.g. network functions that are not supported through OpenStack APIs)
5. VNF controls (VNFs northbound interface is SDN controllable interfaces - application specific)
6. Decomposed physical network functions (dPNFs)
7. Physical network functions (PNFs), possibly with “legacy” control / management plane interfaces
8. Remote access network PNFs
9. Core network-facing functions - including decomposed (dPNFs) and PNFs

VCO Network Topology

VCO core physical network topology is typically based on leaf-spine fabric architecture, composed of physical, commodity Ethernet switches, and fabric control and management functionality exposed through OpenDaylight controller’s southbound interfaces.

All VCO nodes, irrespective of their type, attach to the common system fabric. Node types and roles can vary, and include infrastructure services nodes (such as controller nodes, shared storage nodes, networking nodes etc.), general purpose compute nodes (i.e. nodes used to host Virtual Network Functions and other virtual functions), access nodes to support interfaces to access network subscribers or elements (e.g. xPON, xDSL, access pt-pt optical interfaces, and associated adaptation functions). Node types can also vary in terms of capabilities, such as attachment speeds and numbers, and the presence of hardware acceleration functions. In any case, anything that attaches to the system fabric has to do so with the Ethernet-based interfaces, and utilizes the set of provided fabric services.

For the physical network functions that are not decomposed or otherwise directly integrated with the VCO control and management infrastructure, two attachment options are available:

1. **Direct fabric attachment through local Ethernet interface(s)**, which is supported as long as the fabric service model can support the associated interface protocol stack and hand over traffic to next virtualized or physical network function
2. **Attachment through a specific node.** This would typically happen due to the incompatibility of the fabric service with the PNE interface, or performance considerations that require interface-specific hardware functions, such as acceleration.

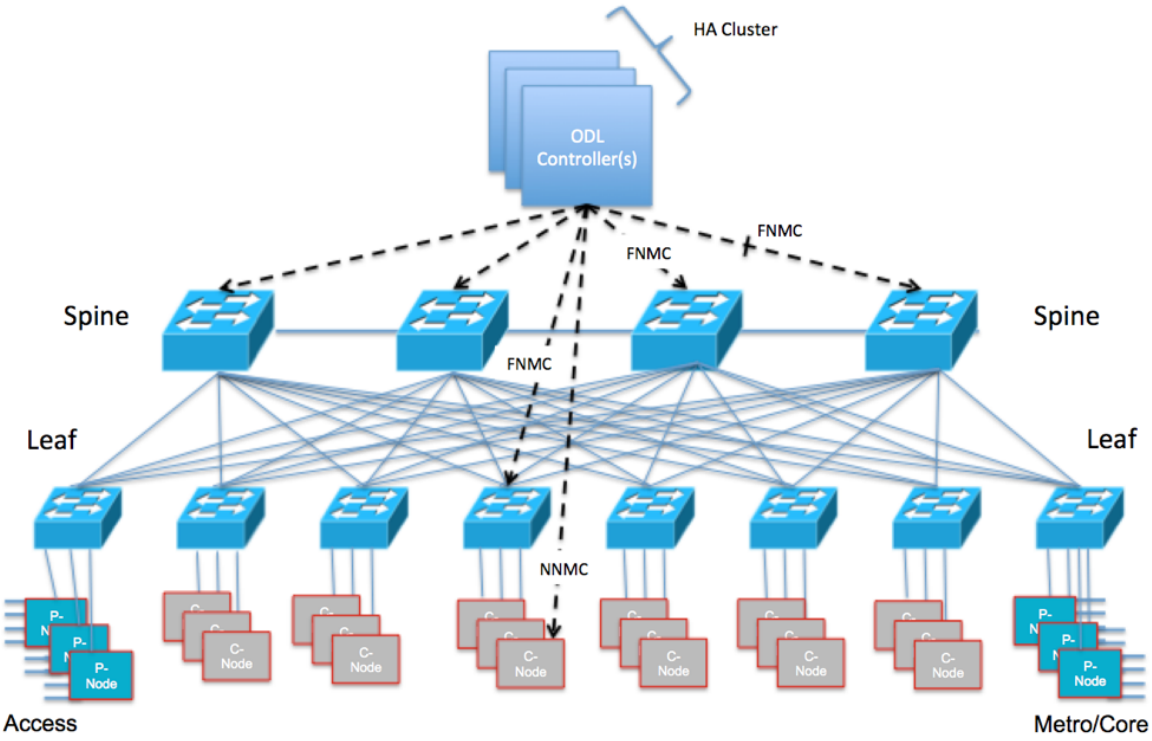

Virtualized Central Office Topology

The general purpose compute nodes used to support the overlay topology are implemented primarily on the nodes using software switches such as OVS or VPP. Virtual overlay topology is expected to be a primary mechanism for the VNF component instances, whenever performance constraints allow it. In the fabric interface, the NFVI overlay implementation will initially be based on a meshed set of tunnels using VXLAN encapsulation (and established using BGP,i.e. EVPN/VXLAN) between the virtual switches for L2 overlay services. The L2 overlay service models in the NFV case are generally expected to be based on the Ethernet MEF service models, i.e. ELAN, E-Tree and E-line service. Note that the service model available for the virtualized functions (VFs/VNFs) is generally considered to be independent of the underlying implementation, including the associated tunneling mechanisms. However, the interworking requirements with external elements and networks, VCO physical / decomposed elements, virtual switch bypass mechanisms, and acceleration functions impose additional constraints on the implementation mechanisms. For container-based VCO instances, we can use the same networking models as for VM cases, due to internetworking both internal (e.g. connecting in-container and in-VM VNFs/VNFCIs), and external (e.g. external networks, esp. towards access networks) to the system.

Control and Data Plane - Full Stack

The VCO architecture uses a single ODL controller cluster to configure OVS (in servers), as well as physical network devices and Virtual Network Functions as shown below.

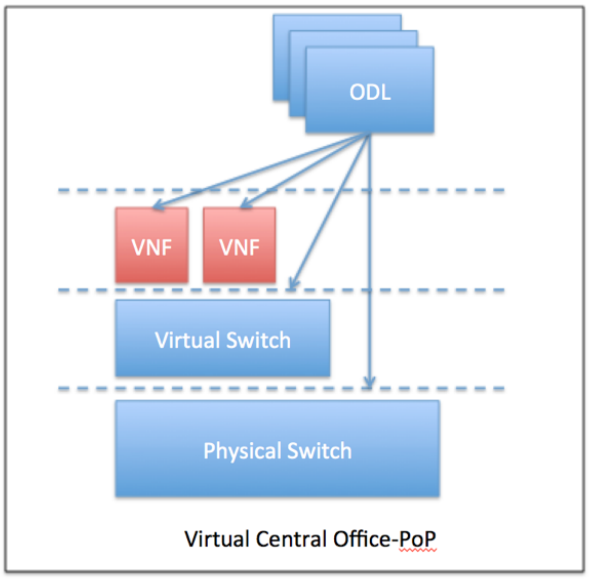

The ODL controllers are clustered for high availability and provide a single interface to manage the entire virtualized central office infrastructure. ODL controller has all the necessary plugins and protocols to interface to physical devices such as the leaf-spine switches, virtual switches such as the OVS and the virtual network functions such as the firewall, load balancers, virtual routers etc.

For most managed network functions, there is a need to expose at least the interfaces required to configure, control and retrieve the events and metrics from the MNF, as shown below. The arrow illustrates the direction of the main dataflows, but all interfaces are expected to have bi-directional protocol adjacency relationships due to the need to support reliable information transfers.

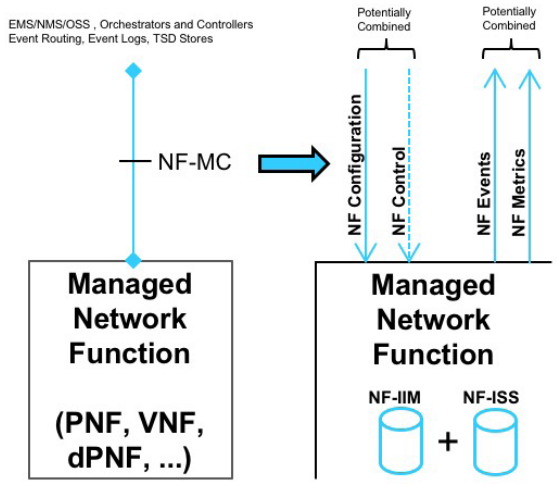

The plugins and interface options that are currently supported by ODL controllers are described below.

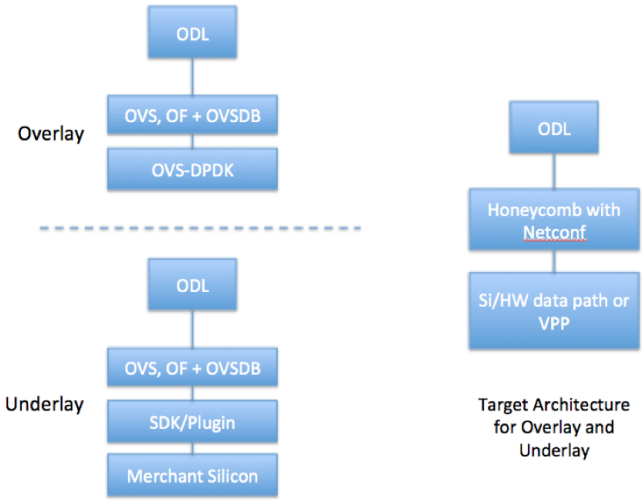

Figure 4 – Data Plane and control plane interfaces for Overlayand Underlay.

The OpenDaylight controller can be deployed in two different ways to configure and manage the underlay and the overlay inside the VCO. For the underlay, configure the leaf-spine switches, switch ports and Open vSwitch using native plugins or OpenFlow agents using OVSDB and OpenFlow. In the overlay case, the controller needs to configure the OVS on the server and the accelerated data path such as DPDK or SR-IOV using the Openflow protocol and and the OVS configurations protocol, OVSDB. The same ODL controller can be used to configure both the overlay and the underlay.

If another SDN data plane such as VPP is used, then a single ODL controller can be used to configure all components, the hardware datapath in the leaf and spine switches as well as the software data path on the servers using VPP. In this model, network devices can be modeled using YANG, and then Netconf is used by the controller to configure those devices.

This model simplifies the architecture considerably and the same stack applies to the leaf switch, spine switch, server as well as hardware nodes such as a vOLT or a physical device that supports the same data plane.

Disaggregation of Services

There are two ways of virtualizing and instantiating residential, business or mobile services. One option is to virtualize the entire bundle associated with the business service and the other option is to disaggregate the components in each of the services and reuse them by stitching components together using Service Function Chaining. The second approach allows the operator to choose components and build services that are flexible and dynamic. Disaggregation of services also allows growth of components independently without inhibiting the overall service model. It also allows the operator to optimize the components and reduces duplication. For example, if a VNF is bundled with DHCP, Subscriber Management, QoS and Routing as would be the case for a vBNG (Virtualized BNG), then any of the components in the vBNG would need to be separately instantiated for mobile or enterprise services. Design or security policies may make such separation the preferred choice, but there may be cases where better optimization can be achieved between services using the same components. In such deployments, disaggregation is valuable.

Composable Services

Composable services using virtualized components is the preferred approach when building residential, business or mobile services for the Virtualized Central Office. These components can be composed into flexible services using the service chaining capabilities of the OpenDaylight controller.

Forwarding Graphs and Service Chaining

Service function chaining is necessary to stitch together Virtual Network Functions (VNFs) into a service characterized by a Service Level Agreement. It can be provisioned in one of two ways:

1. Use the ODL controller to create the VNFFG and provision the data path.
2. Use the Network Service Header (NSH) as defined by IETF SFC WG. This requires classification capability and metadata description for the creation of dynamic service chains.

Either of these approaches is valid. However, for NSH to work, a classification function is needed to identify the forwarding equivalence class either via a dedicated classifier or via OVS extension. In addition, packets must be routed based on NSH information to appropriate VNFs.

For this whitepaper we will assume that SFC is done using OpenDaylight VNFFG capability. Below is an illustration of the service chain setup using an OpenDaylight controller cluster.

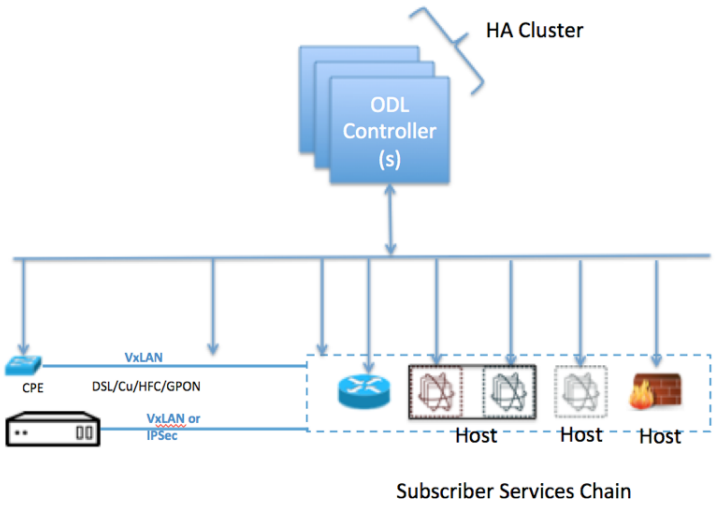

Residential Services Model

Timing and Synchronization Services

Timing synchronization shall be based on IEEE 1588v2 PTP, and meet the requirements of the PTP Telecom profile. Note that 5G systems may have more strict synchronization requirements, but this depends on a functional split in CRAN/vRAN implementations as well as other aspects of 5G RAN system design that are still under active discussion on standardization bodies.

We will address the synchronization needs in detail as we get to the RAN use cases, and as the associated standards progress.

HIGH AVAILABILITY

High availability is needed to ensure service availability and continuity. Enabling mechanisms are distributed throughout the system, at all levels, and the redundancy of the hardware and software subsystems eliminate single points of failure (SPOFs) that can affect the service.

Control mechanisms for these redundant components are distributed to associated control and redundancy management systems. The control processes should be implemented at the lowest level that has sufficient scope to deal with the underlying failure scenarios, with escalation to higher-level mechanisms when failure can not be handled at lower levels. This ensures that the fastest mechanisms are always used, as the remediation control loop times generally increase when the decision logic is moved up in the stack or into broader control domains.

High availability alone is not sufficient for telco services. In addition to high availability of the service, service continuity after failures and failovers is also required. Service continuity requires associated state protection mechanisms to be implemented for the stateful components, e.g. using state checkpointing. State protection mechanisms are the responsibility of the associated state "owner" processes, whether they are part of the infrastructure or part of an application. In other words, SDN controllers would be responsible for their internal state protection, while applications implemented as VNFs would be responsible for their internal state.

For the redundancy, the target is to decrease the amount of redundant resources required. This implies that the direction of the redundancy structures is away from dual-redundant (typically active-standby 1:1 prevalent in the "box" implementations) towards either N+1 or N:1 schemes. This is enabled by the reduction or elimination of the direct physical attachment associations with the use of the virtualized functions, a move towards micro-services types of architectures, and the ability to add, remove and relocate service instances dynamically, based on offered load and resource utilization state. The infrastructure support for such mechanisms is essential for success, and the new, more dynamic configuration of the applications and services has many implications for both network services as well as the whole control, orchestration and assurance software stack.

The availability and continuity mechanisms use “intent”-driven interfaces. The intent specifies the desired state, connectivity or other aspect of the system from the service user’s perspective. The orchestration and control systems determine the implementation of the request, using the network resources available at the time of new request, and subsequently ensure that the intent continues to be met while the service is in use. Intent statements are expected to be used at multiple levels of the system: for example, a VNFD specifies the application topology and related configuration rules. Together with the instantiation parameters, the VNFD represents the initial intended configuration, which can be subsequently modified (e.g. due to scaling) through intent changes.

The associated orchestration and control subsystems are always aware of both the intended configuration and the actual configuration, and can autonomously work to drive the system to intended state should there be deviations (either due to failures or due to intent changes). In addition to declaring basic service specifications, the service interfaces at the top of the orchestration functions can contain e.g. QoS/KPI SLO parameters, which can be used by the assurance-related control mechanisms in the controllers to meet the requested SLA. Intents are closely related to policies, which can be used to couple the intents to the implementation in the context of the underlying physical and virtual network infrastructure. It should also be noted that intents do not necessarily need to originate solely from service requests driven through the orchestration APIs, but can also be determined from service provisioning processes driven from subscriber or site authentication. In such case, the intents for the associated service configurations come from the subscriber database, and is retrieved and instantiated at the time of the subscriber attachment to the network.

Network Topology HA Considerations

In the core of the VCO network, the topology is expected to always support sufficient redundancy with multiple paths to any network destination, redundant interfaces, standard OAM mechanisms and control plane protocols. Topologically, the core of the network is generally mesh, ring or composite. IP routing protocols, MPLS mechanisms and link/path fast failure detection are all standard components of managing connectivity failures.

In the access network, redundancy mechanisms are more technology-specific, and also vary by the market or application. For example, redundancy and access connection redundancy are typically not supported for residential customers due to extra cost, while for large and critical business locations, redundancy may be supported all the way to the customer premises. Redundancy mechanisms need to be investigated in the context of the specific access technology, especially for mechanisms at the level of the

subscriber interface. If the access network uses elements in the outside plant as the primary implementation, then typically the associated access-aggregation network is implemented in redundant ring/mesh topology similar to core/metro networks, but the implementation may be focused on Layer 2 and lower-layer mechanisms instead of IP/MPLS due to the service model differences. If IP is used in the access network, then Layer 2 connectivity needs to be tunneled over IP, or MPLS is used to support customer separation. Technology-specific HA implications will be addressed in the access network sections of this document.

Service- and rate-agnostic WDM elements with Add-Drop capabilities can provide basic access network connectivity in terms of point-to-point wavelengths between two specific endpoints, especially at the lower levels of the network (i.e. in distributed access-aggregation scenarios), and to provide connectivity between different CO sites. Multiple projects are presently defining these interfaces, with open implementations for the associated physical equipment.

ACCESS NETWORKS

PON for FTTH access use cases

The first use case selected is xPON access for the FTTH applications. The system configurations are limited to single-family homes or other 1:1 relationships between the ONT and subscriber. Solutions optimized for multi-dwelling units are specifically out of scope of this use case. This section provides a brief overview of the elements of PON in the context of the FTTH use case, and discusses the implications and priorities with respect to the implementation in the VCO environment.

There are multiple types of the PON from several standards organizations; these vary mostly in terms of data rates, power budget, and split ratios. Examples of the relevant current PON standards include GPON, XGS-PON, NGPON2 (ITU-T) and EPON, 10GEPON (IEEE). The initial FTTH PON implementation is expected to be based on GPON [G.984.x], as its data rates are generally sufficient for the residential access use cases, and it is presently widely deployed for such applications by telcos, which translates to a wide selection of hardware elements (at both component and "box" levels) available to support it.

The figure below shows the representative physical elements of the GPON system implementation, including the interconnect to the VCO Fabric (connected to leaf switches in this case).

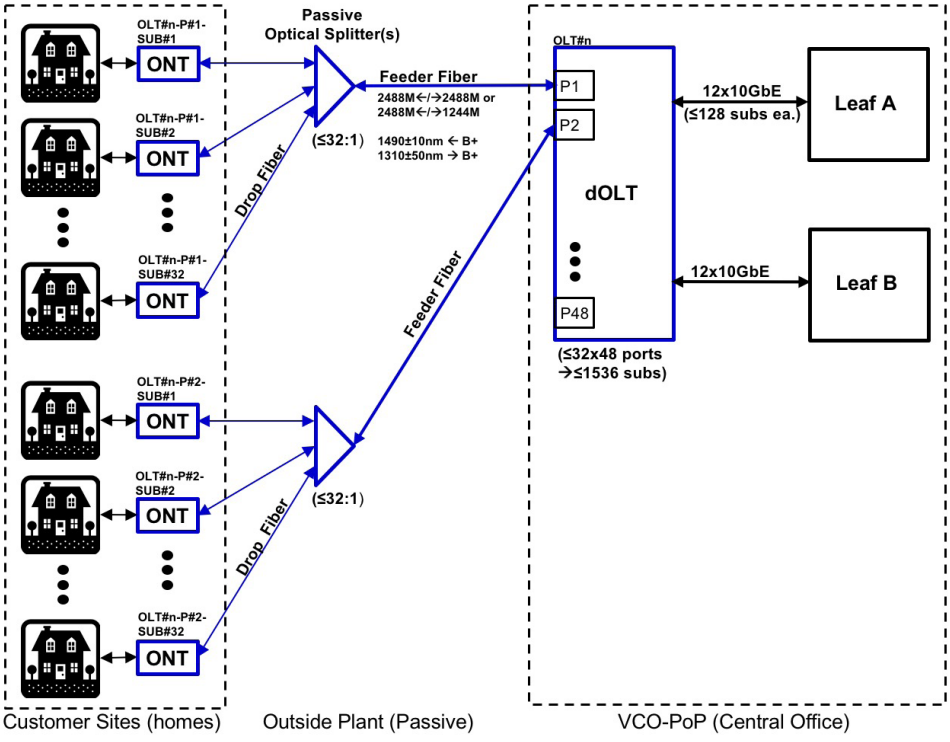

The key characteristics of the physical system in this configuration are:

- Decomposed OLT (dOLT), which implements the physical termination of the PON access lines within the central office
- dOLT supports up to 48 GPON interfaces, at a maximum rate of 2488 Mbps downstream and 1244 Mbps downstream. Other rate combinations are not not considered.
- GPON interface optics is B+, with target max. reach of 20 km between vOLT and ONT
- dOLT interfaces to leaf switches through dual-redundant 10Gbps Ethernet interfaces (12x10Gbps for each leaf); redundancy configuration of these links is active-standby for each 10G pair
- There is no inherent oversubscription between access and fabric side of the box; however, on downstream direction the 10G rate exceeds the individual GPON port rate by factor of 4, meaning that it is possible to have major congestion points on the port egress queue unless the traffic is shaped somewhere upstream in the system
- The design target is for up to 32 split ratio on each GPON interface. This translates to 48x32=1536 subscribers per box. This is low enough that VLAN tags can be used for the subscriber site traffic separation on dOLT-fabric interface.
- ONTs are located in the customer premises, and the initial plan is to start with the simplest possible device that can implement a usable ONT (GPON to 100Mbps/1Gbps Cu Ethernet NID functionality only)
- More complex ONTs with greater functionality (e.g. wireless access points, MOCA or POTS interfaces, etc.) can be integrated in the ONTor implemented through other site-specific CPE devices, and integrated to the overall VCO system as required after the basic connectivity related functionality is in place
- For the additional services, the expectation is that they are built on top of the underlying IP access services. TDM and RF Video Overlay implementations are explicitly considered to be out of scope of the VCO project's goals.
- Both broadcast and on-demand video services are expected to be implemented through the IP connected "set-top-box" or other TV connected devices (game consoles, DVD/Blu-Ray players) or directly through integrated applications on smart TVs and other devices
- In the above configuration, the whole outside plant is non-redundant. The biggest SPOF is dOLT, which translates to service interruption for up to 1536 customer sites upon failure.

Since the sufficient redundancy support is identified as the key design target for all elements and subsystems of the VCO project, the following configuration describes how more of the GPON system redundancy could be implemented to reduce the failure domain size, based on the protection mechanisms outlined in the GPON specifications.

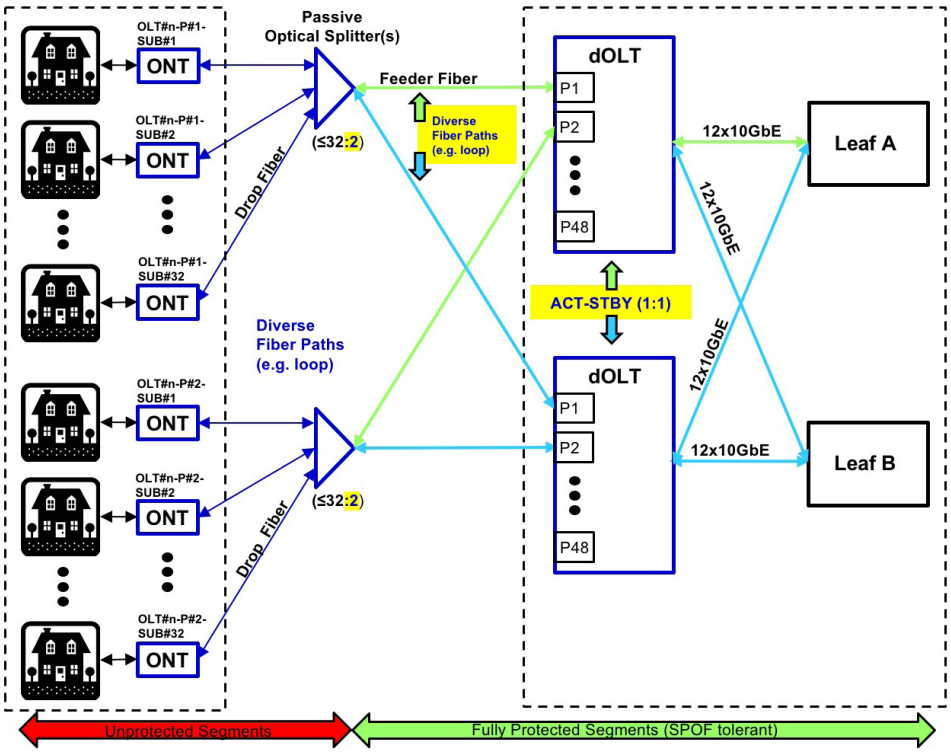

The configuration utilizes most of the same components as the non-protected configuration outlined above. This configuration is specified as “Type-B” protection in the GPON family of specifications. The key differences from non-protected configurations are:

- dOLT is dual-redundant, with a pair in active-standby redundancy configuration
- Redundant dOLTs are hosted on the same VCO-PoP
- Splitters are 32:2 splitters instead of 32:1 splitters
- The OSP path is expected to exhibit sufficient path diversity to avoid trunk cable being a single point of failure
- The combination of the increase of port count in splitters and potential extra fiber length caused by path diversity increases the losses on the OSP, which need to be taken into account on the power budget calculations. This may reduce the maximum dOLT-ONU distance that can be supported, and may reduce the maximum split ratio.
- ONTs are same as in unprotected configuration, as is the drop topology to the subscriber associated splitter
- All components up to the splitter are free of single points of failure. The splitter would be a failure point for up to 32 subscribers (with target 32:2 split ratio), but as it is a passive component its failure rate is expected to be negligible
- Drop fiber and ONTs are not protected, and therefore are single points of failure. However, any failure on these elements is limited to a single customer.
- This configuration effectively reduces the maximum failure size to 1 subscriber for most failure modes, at the cost of dOLT duplication and path diversity at the OSP

The above redundancy configuration assumes that the dOLTs are in the same site. The following figure illustrates the target physical configuration in that scenario from the single PON interface perspective. It also shows an alternative topology option where the redundant dOLTs are hosted on the adjacent VCO sites, which would have potential implications on the control and management plane.

The following figure illustrates the expected L2 encapsulations and associated mappings between the Ethernet and GPON interfaces at the high level for the (single family) FTTH use case, which is the initial focus with respect the GPON use. The details on the expected behaviors are documented in the BBF TR-101 and TR-156 specifications, and are widely supported by both ONU/OLT and OLT side GPON chips/implementations.

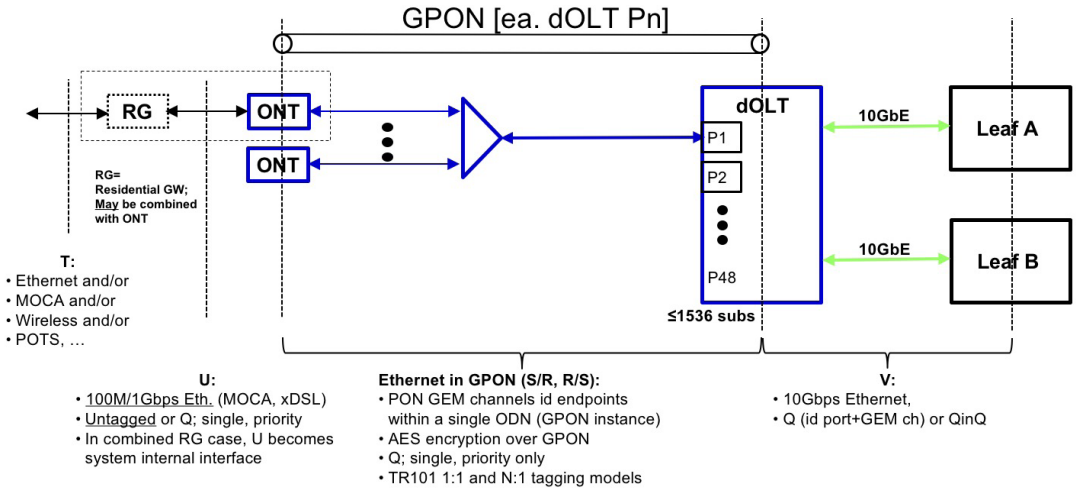

Initially, the RG functionality is expected to be outside of the ONT/ONU implementation, but even external services may still have effects on the encapsulations as well as mappings and protocol interactions on the GPON section of the system to support IP multicast-based delivery of “broadcast” video channels without separate RF overlay wavelength. As the PON downstream direction is inherently broadcast-and-select architecture, there is no need to copy channels separately to each customer. However, as the typical channel line-up can well exceed the GPON interface bandwidth, it is not feasible to broadcast all streams to all PON subscribers, which necessitates the support for multicast and associated IP group membership management mechanisms within the GPON system components and potentially also in subsequent fabric switches. In any case, the starting point is a single VLAN tag per ONT at the ‘V’ reference point, which is sufficient for the baseline IP service delivery across the PON access network from the vBNG/vBRAS function to each distinct ‘U’ reference point. Even without considering the potential additional tags for multicast group traffic, a single dOLT consumes nearly half of the VLAN space, so leaf switches need to support either additional Q tags or some other suitable encapsulation (VXLAN, MPLS or other tunneling) to be able to distinguish between the multiple associated dOLTs and their associated customers/services.

The software interfaces, functionality, decompositioning and associated requirements for the active elements for the GPON access portion are discussed next. This assumes that the dOLT is implemented based on the OpenCompute Telco project's Open GPON specification, or its functional equivalent.

The key elements of the GPON dOLT that require software support in the VCO target configuration include:

- X86 CPU subsystem for overall management (Linux + PON management stack)
- BMC subsystem for management of HW, esp. power and cooling resources (OpenBMC)
- GPON SOC (Microsemi PAS-5211 [PAS5211]; each SOC supports 4 GPON interfaces)
- Physical GPON interfaces (in SPF cages),
- 10G Ethernet interfaces (in QSFP cages); copper interface (passive DAC cable) is used to connect dOLT box to the leaf switches, therefore there is no need to manage the optical interface options here (at least initially)

The key aspects of the GPON ONT that require software support in the basic ONT VCO target configuration are:

- Site level AAA support (assumed to be 802.1x based, i.e. no PPPoE etc. legacy mechanisms that require additional encapsulation layers for all traffic)
- PON MAC / interface management through OMCI (DBA, security, ...)
- Ethernet interface management
- Support for the "box" and its environmental functions: power and battery management, "dying gasp", etc.

This configuration assumes ONT implements only the basic Internet access service.

SUMMARY

As demonstrated in the above whitepaper and on the keynote demo stage at the 2017 OPNFV Summit in Beijing, it is possible to architect a VCO using open source communities and components like the OpenDaylight controller and available orchestrators.

OPNFV and OpenDaylight, along with other open source communities, have come together using open source components to build the the VCO hardware and software stack. Gaps in current OpenDaylight functionality, such as the fabric management and fabric control capability, are actively being filled by ODL members and fabric management and EVPN capabilities will be available in future OpenDaylight releases. OPNFV projects such as Doctor and Barometer continue to explore what's possible for telemetry and end-to-end service assurance. There is considerable industry interest in further development of the open source VCO model to help meet the most pressing CSP virtualization use cases.

REFERENCES

[ATSCALE] SK Telecom ATSCALE White Paper, Vision on Future Telco Infrastructure, July 2016, SKT: http://www.netmanias.com/en/post/reports/10843/sdn-nfv-sk-telecom/atscale-white-paper-sk-telecom-vision-on-future-telco-infrastructure

[COMBO] Convergence of fixed and Mobile BrOadband access and aggregation networks, EU project deliverables, 2013-2016: http://www.ict-combo.eu/index.php?id=4_deliverables

[COSMOS] SK Telecom COSMOS White Paper, Evolving Telco Data Center with Software Defined Technologies, SKT, August 2016: https://developers.sktelecom.com/upload/opdc/document/20160902/20160902185733711.pdf

[ECOMP] AT&T ECOMP (Enhanced Control, Orchestration and Policy) white Paper, AT&T, 2016: http://about.att.com/content/dam/snrdocs/ecomp.pdf

[GANA] GANA - Generic Autonomic Networking Architecture, ETSI White Paper No. 16, 1st Edition, ETSI, October 2016: http://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp16_gana_Ed1_20161011.pdf

[OpenGPON] "AT&T Open GPON NFV OLT Line Card Specification", V1.2.1 http://files.opencompute.org/oc/public.php?service=files&t=66cf54e664c0317fa59f8e20e562bcc7&download

[OPENO] Open Orchestration Project (Open-O) website: https://www.open-o.org/

[ODL] OpenDaylight Project website: https://www.opendaylight.org

[OPNFV] OPNFV Project website: https://www.opnfv.org; OPNFV Summit VCO Demo Keynote Video: https://www.youtube.com/watch?v=Ow4KGaGU3uw; OPNFV Summit VCO Demo Booth Video: http://www.telecomtv.com/articles/opnfv-summit/poc-opnfv-community-demo-virtual-central-office-15783

VCO Solution Brief: https://www.opnfv.org/wp-content/uploads/sites/12/2017/09/OPNFV_VCO_SolutionsBrief_Oct17.pdf

[OSM] Open Source Mano project website: https://osm.etsi.org

[PAS5211] https://www.microsemi.com/products/optical-networking/ftth-pon/pas5211-gpon-olt

[VZSDNNFV] Verizon SDN-NFV reference architecture, Version 1.0, February 2016: http://innovation.verizon.com/content/dam/vic/PDF/Verizon_SDN-NFV_Reference_Architecture.pdf