Transforming the Telco Central Office with NFV and SDN

The data revolution is transforming the consumer and business landscape with intelligent transportation, smart cities, home automation and IoT devices. The common theme: data is everywhere and is growing exponentially.

Over the years, network service providers have invested in multiple technology transitions in mobile, broadband access and business services. This has led to disparate SP networks, each with its own proprietary architecture, leading to enormous OPEX and CAPEX budget bloats. More importantly, these silos have crippled their ability to introduce new services or scale out existing services in an agile manner. With increased bandwidth demand and reduced revenue per bit forecasted, service providers are under pressure to reduce their TCO and transform their infrastructure to drive innovation, all while being more agile in service delivery.

With the advent of Network Function Virtualization (NFV), service providers are at the cusp of yet another major technology transformation to drive experience, services and innovation to cater to the new business models.

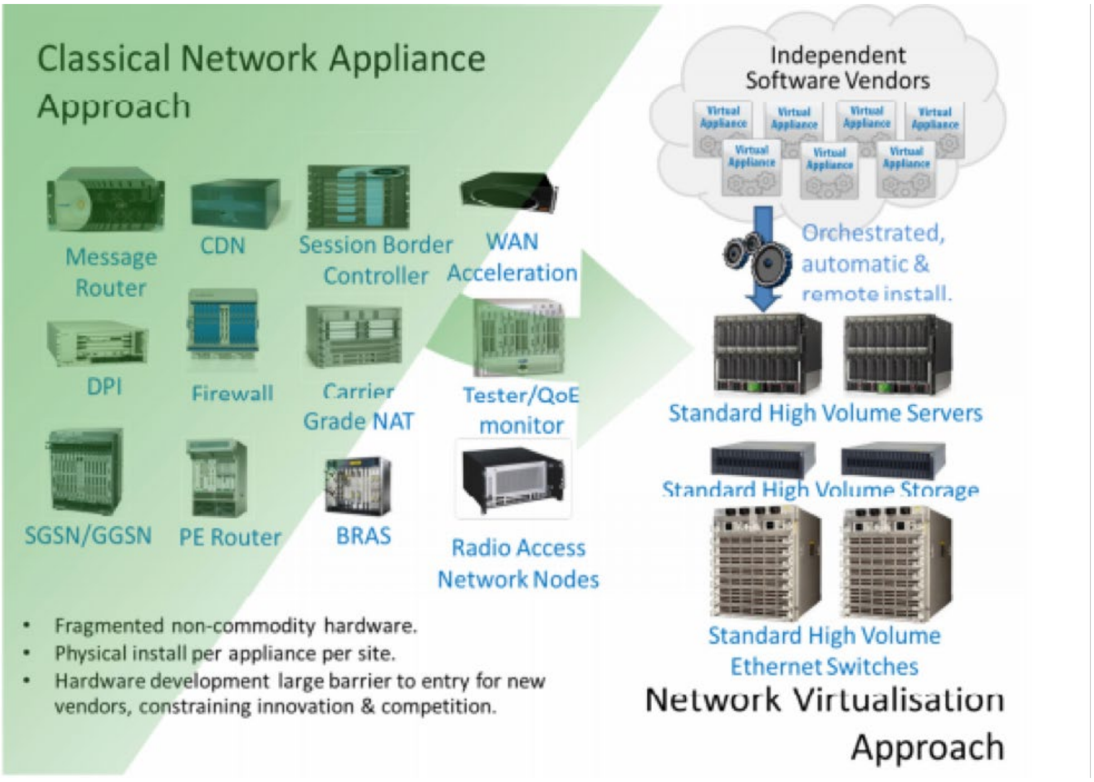
Classical Network Appliance Approach
Message Router
CDN
Session Border Controller
WAN Acceleration
DPI
Firewall
Carrier Grade NAT
Tester/QoE monitor
SGSN/GGSN
PE Router
BRAS
Radio Access Network Nodes
Fragmented non-commodity hardware.
Physical install per appliance per site.
Hardware development large barrier to entry for new vendors, constraining innovation & competition.
Independent Software Vendors
Virtual Appliance
Orchestrated, automatic & remote install.
Standard High Volume Servers
Standard High Volume Storage
Standard High Volume Ethernet Switches
Network Virtualisation Approach

Figure 1 : ETSI Vision of Network Function Virtualization (https://portal.etsi.org/NFV/NFV_White_Paper.pdf)

NFV, driven by virtualization and defined by ETSI, allows the user to decouple the control, bearer and management plane functions, all running on commercial off the shelf (COTS) compute hardware, instead of deploying expensive point products. Many, if not all of the specialized network functions, such as firewalls, load balancers, mobile RAN functions, subscriber policy management, even PEs and CPEs, can now be delivered as a virtual function. These are referred to as Virtual Network Functions (VNFs), which drive the new architecture and deliver:

1. Improved capital efficiencies, compared to dedicated hardware implementations
2. OPEX efficiencies based on a simplified and consistent operational model
3. Service agility in the form of an open and programmable automation framework
4. Flexibility and scale; placement of VNFs, independent of hardware, allows for horizontal scalability and largely decouples functionality from physical location
5. Standardized and open interfaces allow for interoperability between virtualized network functions, the infrastructure and associated management entities

These attributes help create an agile service delivery model for SPs for all architectures. This is an effort similar to what cloud providers have done in the datacenter to level the playing field on the service platform, allowing different VNF vendors to differentiate with best-of-breed software. However, creating and delivering new services using VNFs requires, at a base level, some form of orchestration (sizing, capacity, automation, etc). Software Defined Networking can also be leveraged to further integrate and automate the end-to-end system. Note that while SDN is complementary to NFV, the two are often deployed together to deliver a fully automated solution.

The NFV Impact - SP Central Offices become Cloud Datacenters

SP-based business services are often distributed out to the Central Office (CO) Edges where the traditional collection of physical function hardware exists. Using this distributed and "proximity to end customer" advantage, SPs can now transform their COs into a cloud-based architecture, based on an NFV architecture, which is composed of distributed datacenters.

This has many impacts and benefits, including:

- **Common network for multiple services:** This results in a uniform Telco cloud "blueprint design" that can be replicated at various sizes across a number of interconnected COs, as well as centralized DCs, delivering multiple services. This also drives down capital expenditures on hardware, power, and space in the CO. As a result, OPEX efficiency is improved due to the simplicity and standardization of the architecture.
- **Simplified operational model:** The user is able to simplify service delivery using a cloud scale architecture for a variety of use cases including SP broadband, mobile access, business services and content delivery.
- **Resilient, scaled-out, and programmable framework:** This framework is highly resilient and available, as opposed to a collection of highly specialized, brittle, and dated single points of failure.
- **Common automation framework:** A software-centric approach allows for a high degree of automation and provides a scalable operational model for multiple geographically dispersed DCs that are operated as a common pool of resources.
- **Agile service delivery:** The combination of the above attributes brings the ability to spin up services in a short period of time as VNFs. This creates agile services delivery and enhanced end-user experience, catering to new business models.

Components of the Telco NFV Cloud

ETSI defines an architectural framework for deploying NFV-based services. Operators have since taken that model and derived the necessary components to deploy production services. Essentially, to deliver a Telco NFV cloud, the main pillars that work together are:

1. The network fabric
2. The automation and orchestration layers
3. The virtual layer hosting VNFs

The NFV Network Fabric

Depending on the size of the service provider, the network may generally consist of:

- Centralized and distributed cloud fabrics (the NFV datacenter networks)
- Existing core network (MPLS-based)
- Existing (but morphing) metro edge network

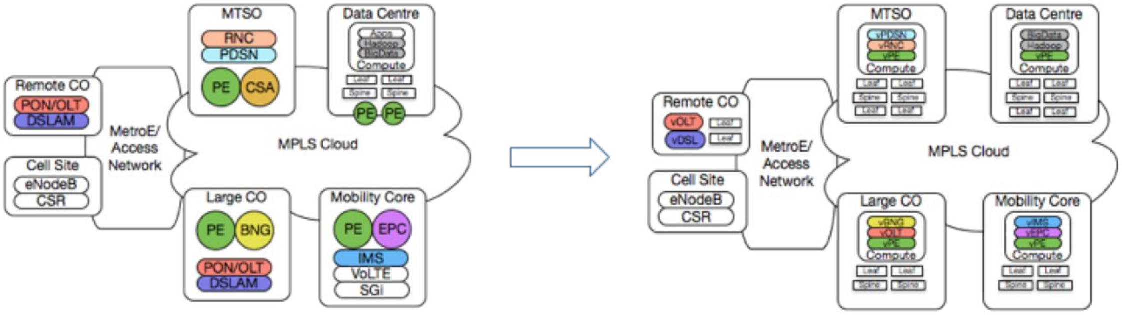

Figure 2: NFV architecture, transforming the Central Office

These networks must coordinate together to build an end-to-end service for the customer. The use of centralized or distributed NFV clouds will depend on the type of VNFs used, the scale, and the latency requirements. The following is an examination of the various roles of the NFV network fabric:

A. The Transparent Underlay IP Fabric

The ETSI NFV architecture does not dictate a specific approach to building the NFV cloud fabric. VNFs being decoupled from the underlying hardware and hypervisors can be instantiated as VMs, containers, or on a bare metal server and attached anywhere to the network. This resembles a cloud architecture, where the network becomes a fabric for VNF applications and provides a common communication medium for both east-west and north-south traffic flows. This underlay network is now an open IP fabric, providing connectivity between VNFs and the outside world. Software defined control complements this, where an SDN controller programs the connectivity between the vSwitch and the DC Gateway. The data plane connectivity between VNFs and the datacenter gateway is established using MPLSoGRE/UDP or a VXLAN overlay, which is managed by the controller.

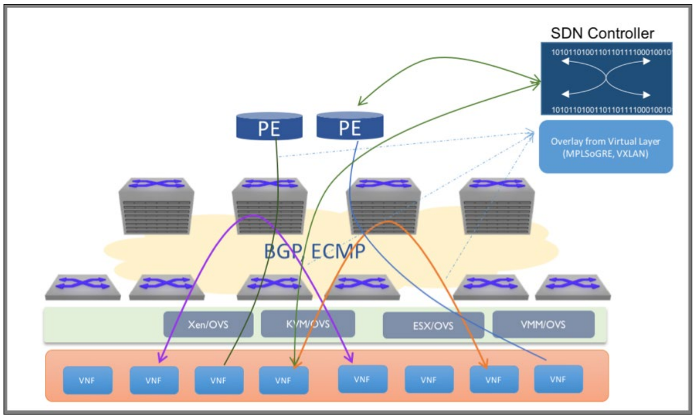

Figure 3: NFV architecture, with network fabric as transparent IP underlay

The networks of service providers are subject to a high bar of five-nines availability, low latency for mission critical applications, carrier grade (NEBS) compliance, resiliency and security. These fundamental attributes are still the principles for the underlay IP fabric as well. But especially due to a need for high performance, resilient, any-to-any connectivity for disparate workloads in NFV environments, here are some key desired attributes for an NFV underlay fabric:

1. **Scale-out design:** An ECMP design can grow over time without significant up-front capital investment, as has been proven in various cloud and web datacenter deployments. These have experienced exponential growth in the introduction of new services.
2. **Leaf-spine design:** A universal leaf architecture for multiple use cases, connectivity options (10G/25G/40G/50G/100G) and fewer tiers decreases cost, complexity, cabling and power/heat, etc.
3. **BGP:** BGP's ability to achieve massive ECMP scale, while providing rich policy control and minimal state churn, have made it a perfect fit for underlay networking.
4. **Buffering and congestion management:** Capability to absorb bursty traffic is subject to external events and delivering deterministic performance.
5. **Open and programmable:** Deliver a consistent operational model that fits any scale and improves responsiveness for managing cloud-scale workloads and workflows.
6. **Resiliency and high availability:** System, path and network availability and resiliency are mandatory, especially for this new model to be capable of 5-9s availability.
7. **Consistent operational framework:** a model that simplifies management and provisioning, enables agile service delivery while lowering OPEX costs, all while creating competitive differentiation.
8. **Telemetry and analytics:** The key to agile service delivery is the experience, as well as the ability to resolve issues and guarantee SLAs. Visibility and monitoring of the dynamic cloud requires a modern automated telemetry framework. Arista AlgoMatch™ technology with high-rate sFlow enhances Arista's Telemetry solution for rich network visibility for troubleshooting, DDOS mitigation, trending analysis or traffic steering.

B. The Hybrid IP Fabric, for network virtualization and overlays

The universal leaf needs to support various designs. One design would be a compute-facing leaf for pure underlay with the transport technology of choice (IP or MPLS). Another leaf design is an edge routing leaf, supporting VXLAN or MPLS termination/origination for incoming broadband access, bare metal, or for VNF performance acceleration technologies. The most common overlay protocols are VXLAN, MPLSoGRE, GRE, and MPLSoMPLS. Moving forward the BGP-EVPN control with a VXLAN or MPLS data plane will facilitate large-scale, standards-based, interoperable overlay signaling.

Drivers for VXLAN (for Overlay)

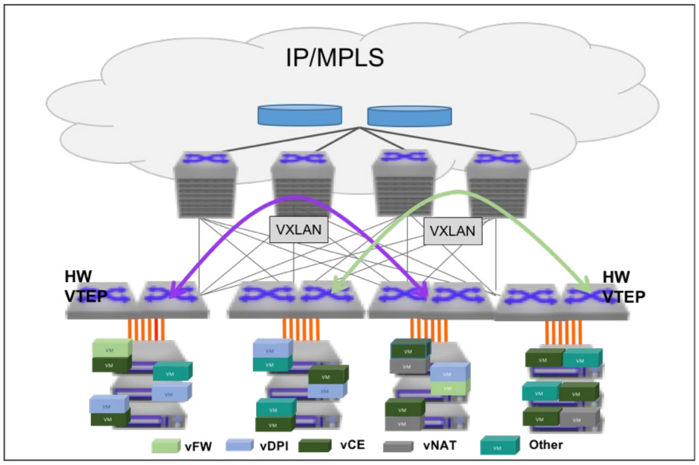

Figure 4: NFV in Carrier Central Office

One of the advantages of the NFV deployments is service agility. VNFs can be spun up and down, on demand, depending on the need to deliver a specific service by the carrier. This imposes a requirement on the network to be able to configure and deploy VNF service, regardless of its placement in the virtual infrastructure. While merging the new VM into its existing broadcast domain, traditional techniques with VLANs are restrictive and not well suited to scale and multi-tenancy requirements. VXLAN helps solve these challenges and facilitates the service deployment, workload mobility, business continuity and multi-tenancy over Layer-3 boundaries. In addition, VXLAN can be deployed on an open IP network running BGP, which is now more prevalent in modern cloud datacenters. With Layer-3 ECMP in the IP network, the user can achieve high resiliency and availability. Software VTEPs with VxLAN are also supported by such a network; however, a translation layer is needed at the datacenter edge between VxLAN and MPLS layer.

Drivers for MPLS (for Underlay and Overlay)

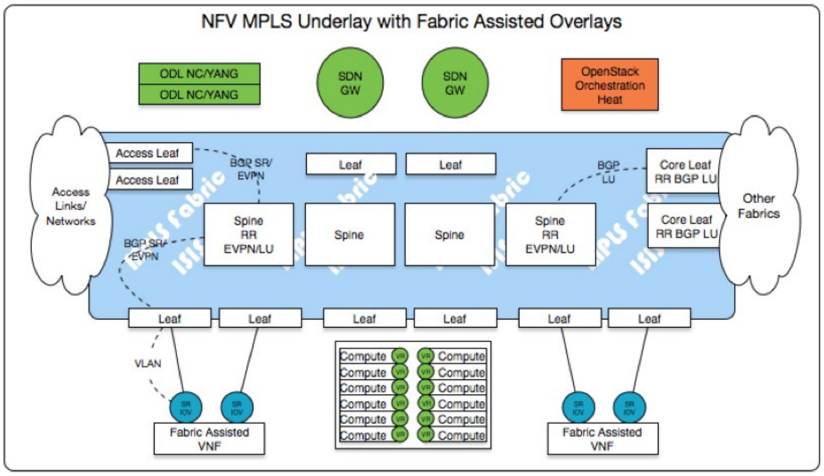

Figure 5: MPLS based Network Fabric

The prevalence of MPLS in Service Provider networks is yet another driver for the MPLS enabled fabric. Not only are their core networks and existing services built on MPLS constructs, but so are their operational and architectural development paradigms. While hypervisors will frequently leverage an SDN controller and software tunnels (pure MPLSoGRE/UDP), a number of VNFs may leverage the hardware components of the fabric for increased forwarding capacity (Fabric Assisted VNFs).

With MPLS as the overlay protocol, operators can have a natural extension of the base service offerings and avoid translation layers (VxLAN to MPLS), and derive operational efficiency by having the ability to integrate with current tools built with MPLS VPN. Solutions like MPLSoGRE, which enable MPLS in the overlay and IP underlay fabric are already prevalent in NFV deployments. Additionally, some operators are considering MPLS transport in the underlay fabric. The advantage that an MPLS transport brings is the seamless integration with the MPLS core and with technologies like Segment Routing, which delivers the ability to use controller driven traffic engineering, while preserving simplicity of design and ECMP behavior.

Tying Together NFV Fabrics, Metro and Core Networks with EVPN

BGP-EVPN further advances the L2 and L3 services on which it was built and addresses some key challenges of previous implementations: most notably, "active/active" multi-homing, tight L2/L3 integration and choice of tunneling encapsulation. Signaling Active/Active segments not only allows providers to solve a long-standing customer to provider (CE to PE) challenge, it eliminates dependence on vendor specific multihoming protocols and paves the way to n-way multi-homing support. Using BGP-EVPN as the method for signaling services within and between DCs/Fabrics provides a scalable, extensible control plane for L2 and L3 services and maintains the foundation upon which these services were originally deployed. BGP-EVPN not only provides choice of encapsulation type (which is signaled in BGP (VXLAN or MPLS), but it also enables multi-tenancy and simplifies tunnel discovery and host/workload mobility. Lastly, BGP-EVPN facilitates datacenter interconnect (DCI) services and provides the option to tie various NFV fabrics over existing networks.

The Automation and Orchestration Framework

The NFV operational framework delivers services in an agile model. This includes rapid provisioning of VNFs, VMs, networks, policy-based traffic steering, along with the ability to monitor, troubleshoot and predict faults via advanced telemetry. The complexity is amplified due to the sheer number of VNFs and the diverse functions they support, which may span many datacenters. This makes the traditional operational model of manual deployment obsolete. A strong model of automation and orchestration is needed to operationalize these deployments.

As with the NFV network fabric, the automation and orchestration framework needs the following foundational attributes:

- A scalable virtualized infrastructure manager
- A common automation framework with open/standard protocol for data modeling, transport and encoding
- A modern telemetry and analytics framework.

Service Providers have embraced open-source technologies like ODL, as their Virtualized Infrastructure Manager (VIM) platform - ATT OpenECOMP being a prime example of the framework being supported/tested by other SPs like Bell Canada, Orange etc. The primary drivers for SP, to embrace open-source, are:

- Developed and supported by a massive open-source community
- Well defined and programmable APIs
- Rapid evolution (new features every 6 months)

Additionally, Openstack/ ODL provides a standard integration point for SDN controllers, 3rd party orchestrators, and management systems. Furthermore, many VNF software providers are developing their components, ensuring seamless support in SP NFV environments.

With the scale and diversity of NFV applications, the OSS model has to evolve to efficiently manage resource allocation. This is needed to drive consistent provisioning, configuration, and image management from day-zero. The new OSS model will additionally drive capacity planning, upgrades, network bandwidth, traffic steering and finally telemetry and analytics. To achieve this, the new OSS must have an open standards based protocol and API model in place to interact with the various layers.

In a large multi-vendor NFV stack, the current model of data collection (SNMP/ CLI is not scalable. Openconfig (openconfig.net) is aiming to solve this important problem, by creating industry standard API for network elements. By creating a standard, open, vendor neutral interface, OpenConfig enables operators to leverage their automation and orchestration efforts over a large swath of NFV elements.

In the NFV approach, due to the scale of services and traffic patterns, the penalty for lack of fine grained visibility and monitoring of the application and network performance will be huge. The legacy model of polling based approach (SNMP) will not scale to predict/ detect, isolate and resolve application performance issues. The modern approach of real time state streaming is required to drive end-to-end visibility, at scale and to guarantee SLAs. The protocols in openconfig model support some degree of streaming event notifications. This is crucial for scalability and performance.

One area that has seen some recent momentum is the ETSI NFV Management and Orchestration (MANO) working group. There are now three open-source efforts working in this space - open-source MANO (OSM, osm.etsi.org), Open-O (open-o.org) and as well as AT&T's Enhanced Control, Orchestration, Management, and Policy (ECOMP, http://about.att.com/content/dam/snrdocs/ecomp.pdf)

While path forward is unclear, these initiatives are important to developing the maturity of NFV services.

The Virtual Layer (VNF Layer)

This layer provides the platform to host the various VNF functions, which includes the hypervisors, VMs and the virtual switch or router. The Hypervisor - ESTI's "ETSI GS NFV-INF 004" (NFVI Hypervisor Domain) document states, "It is one of a relatively small number of critical components which enable the objectives of NFV to be met." An interface to the VIM layer allows compute resources to be a fungible entity with automation and instrumentation. With OpenStack's dominance in the NFV Orchestration Layer, comes the prevalence of the KVM hypervisor at the Virtualization layer. Nevertheless, VMWare ESXi, Microsoft Hyper-V and Xen can also be found as hypervisors at the NFV virtual layer as well.

Virtual Switches and Routers

The point of integration between the networking components (NICs) and the hosted VNFs is a virtual switching component. The VNFs connect to the network directly or over an overlay tunnel (e.g. VxLAN or GRE). These virtual networks (VLANs or VxLAN/GRE tunnels) can be defined and controlled as a function of the orchestration system that integrates with the Hypervisor layer.

Many SDN controller implementations may interface directly with a Virtual Router (VR), which provides automated routing functions between the VMs they are attached to, and acts as a "SDN Gateway" (all of which most commonly signaled via MP-BGP). For VNFs that require high bandwidth network I/O, the implementation can leverage SR-IOV (Single Root IO Virtualization). Most VNFs leveraging SR-IOV will rely on the capabilities of the network fabric to maintain tenant separation, inter and intra-VNF reachability as well as access into and out of the Telco Cloud.

The Arista NFV Solution

The Arista NFV Solution builds upon the extensive production and operational experience that has been gained from deployments in hyperscale cloud providers, which are based on the same architectural principles. Arista NFV framework provides the open and programmable foundation for carriers to build the three pillars for NFV deployments.

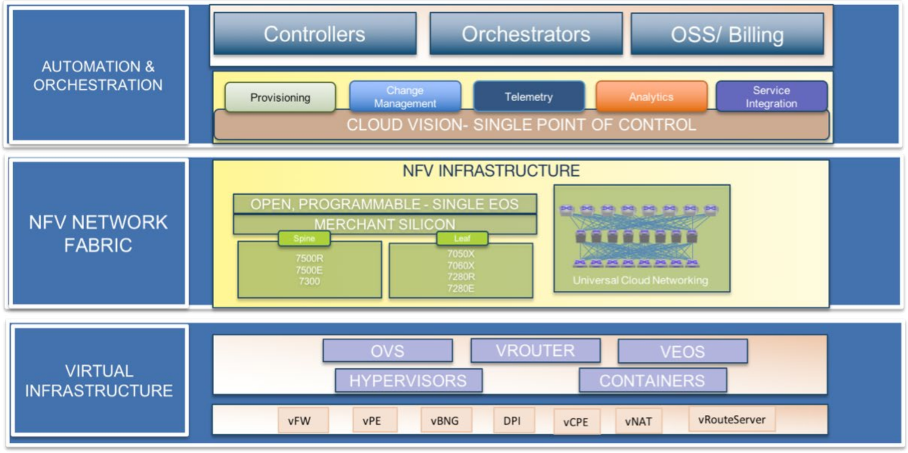

Figure 6: Arista NFV Framework

Building the Fabric: Arista's Universal Cloud Network (UCN) leaf-spine architecture, with the 7500R Series and 7280R Series delivers the best scale-out network architecture for the NFV network fabric. Key attributes such as deep buffering, NEBS compliance, 128-way ECMP, High Availability and resiliency are built-in the market leading high density 10G/25G/40G/100GbE switches. EOS with Arista FlexRoute supporting up to 2M+ route scale, MPLS, EVPN and Segment Routing solutions and CloudVision providing single point of control, automation and real-time Telemetry enabling providers to leverage cost effective and efficient methods for the modern telco fabric.

The Orchestration and Automation layer: At the core of Arista cloud networking solutions is EOS, which is purpose-built to be fully programmable and highly modular. EOS provides a rich set of standards-based, well known programmable interfaces including: Standard Linux and Linux Tools, EOS extensible APIs (eAPIs) using JSON, open-source Go, Python and Ruby based object models, Native Go and Python on box scripting, XMPP, Advanced Event Manager, SQLite Databases and EOS SDK. Arista EOS® has extensive integration with the OpenStack Neutron project, providing powerful network platform for operators to automatically provision tenant networks across the physical infrastructure.

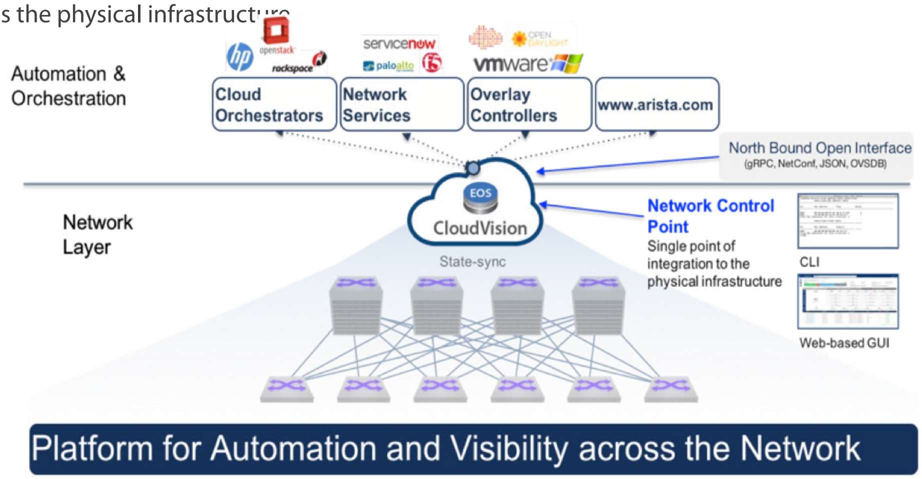

Figure 7: CloudVision - Single point of control for Network, Orchestration and Automation layers

Arista's CloudVision™ integrates with VIMs (OpenStack, VMWare, etc) and SDN controllers to provide VNF-driven network provisioning as well as VNF/PNF integrations used in a hybrid fabric. It also provides the infrastructure required to automate all functions of NFV Fabric provisioning, lifecycle and maintenance. CloudVision also includes CloudVision Analytics engines and CloudVision Telemetry Apps that take full advantage of the state streaming infrastructure of EOS and NetDB, providing customers improved visibility, faster problem isolation and correction, and a view of network performance over time. CloudVision can be used for turnkey fabric installation and operations along with a rich API, which can be integrated into any SDN controller, OSS/BSS, automation or orchestration framework.

Integration in Virtual Layer: Multiple points of integration with the Virtual Layer, such as Openstack, VMWare and Microsoft are all tightly integrated with Arista Networks' software and hardware elements (through Neutron, vSphere, NSX and OMI APIs). Arista also tightly integrates with SDN Controllers, leveraging open interfaces such as OVSDB. Perhaps most importantly, Arista integrates in environments where orchestration systems, SDN controllers and VNFs leveraging SR-IOV or other direct access methods all co-exist and must be interconnected. The Telco Cloud Virtual Layer will not be limited to a "one size fits all" methodology, making the choice of an open and flexible infrastructure critically important to successful NFV deployments.

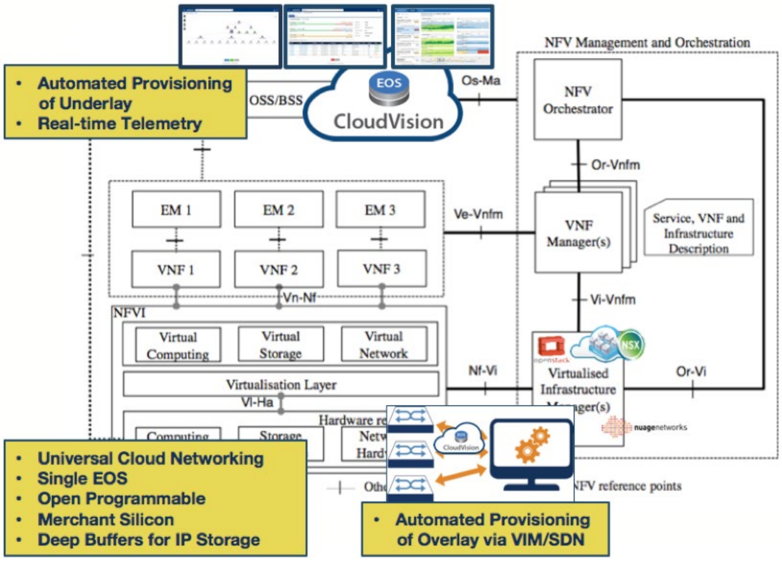

Figure 8: Summary of Arista Integration with ETSI NFV Framework

Summary

Modern NFV deployments will transform carrier networks from a siloed service delivery model to one of Telco Cloud Services. The success of NFV hinges on the critical pillars of the network, a scalable automation framework and the virtual layer integrating with existing infrastructures. NFV creates the largest evolutionary disruption of Telco services in decades, bringing massive cost and operational efficiencies.

Arista solutions combine the principles that have made cloud computing an unstoppable force. Arista's universal cloud architecture, powered by EOS with large scale Layer 3 ECMP and a CloudVision framework, frees applications from the limitations of legacy networking. This combination of network virtualization, custom programmability, simplified architectures, and more realistic price points, creates a best-in-class software foundation for delivering an NFV platform for service provider innovation today and into the future.

Abbreviations

ETSI: European Telecommunications Standards Institute
NFV: Network Function Virtualization
VNF: Virtual Network Function
PNF: Physical Network Function
NFVI: Network Function Virtualization Infrastructure
SDCN: Software Defined Cloud Networking
VXLAN: Virtual Extensible LAN
BGP: Border Gateway Protocol
EVPN: Ethernet Virtual Private Network
CAPEX: Capital Expenditure
OPEX: Operational Expenditure

BSS: Business Support Systems
NETCONF: Network Configuration Protocol
YANG: Yet Another Next generation (data modeling language for the definition of data sent over the NETCONF protocol)
REST: Representational State Transfer
JSON: JavaScript Object Notation
OVS: Open vSwitch
OVSDB: Open vSwitch Database
VIM: Virtualization Infrastructure Manager
ML2: Modular Layer 2

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office
1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062