4 Emerging Software Defined Data Center & SDN Technologies You Should Understand

(BEFORE you deploy!)

White Paper

January 2015

datavision

Future-Proofed Networking Solutions

Contents

Executive Summary

Datacenters are specialized facilities that have been used to house computer elements of all types since the early days of the computer industry. During the dot-com boom, the total number of datacenter deployments proliferated exponentially. It is estimated that there are at least 500,000 datacenters in service today worldwide. The architecture of modern datacenters continues to evolve, driven by industry trends such as virtualization, consolidation, and massively scalable architectures.

In today's business environment, businesses are facing pressure to reduce costs while at the same time creating elastic services with shortened delivery cycles. Servers represent the dominant portion of the equipment cost within a datacenter, however, network architecture and supporting technology choices can influence how well the server assets are utilized.

Recent developments with Software Defined Networking (SDN), Network Function Virtualization (NFV) and related technologies are changing the network deployment architecture and operational model as well as the economics of deploying and operating the datacenter environment for both service providers and enterprise administrators. However, the unparalleled pace of innovation in networking architectures, virtualization of networks, new overlay technologies and solutions for simplified operations, etc. have created an environment for networking professionals which is becoming more and more complex.

Navigating through all the new technology options and choosing the right network architecture has become a strategic imperative for companies that leverage datacenters as their core business, and for those for where IT is considered a strategic weapon.

This technology whitepaper attempts to provide a simplified view of many of the architectural changes and technology options in the data-center networking domain. In addition, we will also provide our view of what typical architectural issues a specific technology option attempts to resolve. For clarity, this whitepaper is comprised of four parts, each addressing different applications of SDN/NFV technology to the datacenter environment.

- **Part 1: Datacenter Physical network,**
  This section highlights major changes and key trends in context of the data-center networking environment. This includes describing new spine-and-leaf architectures, the importance of a virtual switching layer, defining Whitebox switching trends, and key design considerations from physical as well as logical network design perspective.

- **Part 2: Centralized vs. Distributed control of datacenter environment**
  This section discusses the concept of Centralization trends in the networking domain. We will explain the difference between two divergent philosophies i.e. “**Centralized** Management” vs. “Centralized Control Plane” and solutions in-between. Although the decision to use centralized control plane or distributed control plane is one of the first decision one will make (at the stage of hardware procurement), this shall be discussed in part 2 to ensure that the reader is familiar with concepts, key problem statements and available solutions before aligning and adopting a specific approach.

- **Part 3: Virtualized Network Service Functions and Service Chaining:**
  This part will discuss what kind of services are needed in a typical datacenter domain. In additional to describing those functions, we will also highlight services that can be enabled by virtual appliances and how these services can be chained together. Many of the SDN solutions in industry focus on the overall concept of service chaining and provide their own flavor of virtual services as well as the interface to allow for that specific service

- **Part 4: Network Orchestration and Controllers:**
  This part will discuss what network orchestration in the datacenter domain really means from a networking professional perspective. In addition, we will also discuss how traditional VM orchestration systems such as VMWare, OpenStack etc. are evolving from the perspective of networking orchestration.

Part 1: Physical Networks in Datacenter Domain

The section compares the emerging spine-and-leaf architecture with traditional 3-tier architecture for datacenter environments and explains the reasons why it has become the architecture of choice for all new datacenter deployments. In addition we will also introduce the virtual switching layer and its importance for overall network design. Subsequently, we'll describe key design considerations as it relates to deploying a new physical infrastructure for networking team members to implement a spine-and-leaf architecture for virtualized cloud infrastructure.

Datacenter Network Architecture: Three-tier vs. Leaf and Spine Architecture.

The traditional 3 tier architecture consists of a core, aggregation and access switching nodes connected together to create redundant and resilient communication path between the transport networks and compute servers where applications are hosted. This architecture has been the de facto standard for many years and is being used in thousands of datacenter across the globe. The strength of the three tier network design lies in its ability to carry large amounts of North-South traffic – as a result, its performance is satisfactory in client-server centric environments. However, this design is week in the below areas.

However, with the requirements of application in virtualization world (such as dynamic workload mobility and distributed workflows), application deployment guidelines as well as network traffic profile within the datacenter is changing. In new virtualized world applications are completely agnostic of underlying hardware and virtualization features allows the application VMs (such as vMotion, DRS) to get moved within the datacenter or different datacenters automatically. In order to support these requirements, network designers not only need to adopt new design principals but also need to focus on optimizing the network for increased east-west traffic flows.

The new "Leaf and Spine" architecture (sometimes referred to as distributed core architecture) was essentially created to address the problem statement described above. A spine and leaf datacenter network design is an alternative to the traditional datacenter network design that uses a flatter, two-tiered, scale-out network where each leaf switch has an uplink to every spine switch.

Following diagram provides the side-by-side view of traditional three tier architecture and new leaf-and-spine architecture.

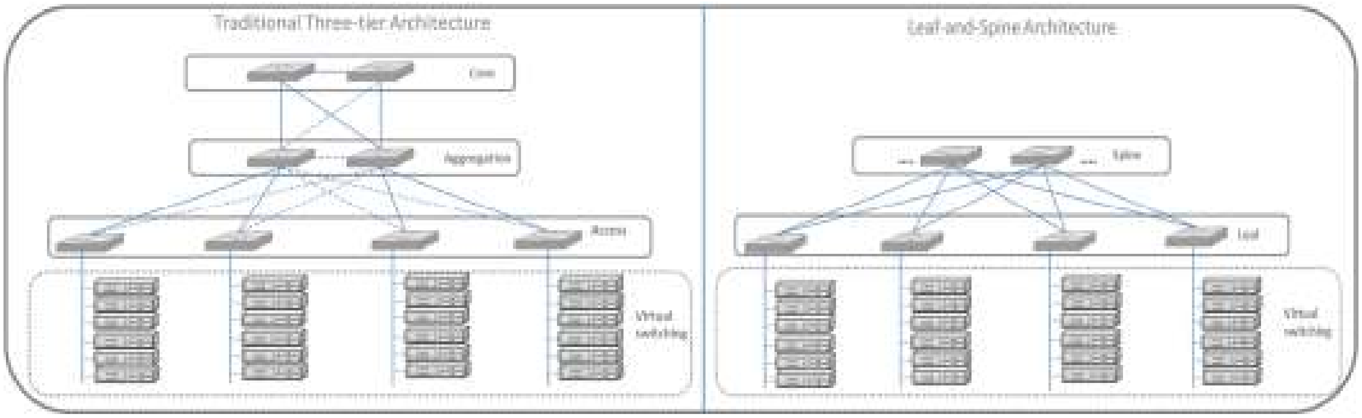

As the name suggests, Leaf and spine architecture) is a two-tier architecture and has one less switching tier as compare to three tiers in traditional architecture. These two switching tiers are referred to as Leaf switching layer and Spine switching layer.

- Leaf switching layer essentially provide the same functions as that of access layer in traditional architecture i.e. network connection to servers/compute/storage and uplink to the layers above. These nodes also referred to as Top-of-the-rack switches or End-of-the-row switches.

- Spine switching layer provide collapsed view of aggregation and core layer in traditional architecture. These nodes are high-port density Layer 3 switches that aggregates traffic from all the leaf nodes as well as connect to the WAN edge routers.

The key difference of leaf-and-spine architecture is that every leaf switches uplinks to every spine switches. As compare to traditional network, each port is working in active mode at all times instead of behaving as a backup link that only gets used when blocked by the Spanning-Tree Protocol (STP) in order to prevent network loops. In general leaf-and-spine architecture has several benefits over traditional architecture. Some of the key benefits include

- **Lower Costs**: Less number of tiers often get translated into lower CapEx requirements as well as lower OpEx due to simplified design, less space, power, cable and cooling requirements.

- **Higher Utilization and Better Resiliency**: Required for Network traffic from each leaf switch gets distributed to all the spine switches. This distribution model makes individual spine switch less important in that if a spine switch were to go down, traffic shall be automatically distributed to other spine switches.

- **Easy to Scale out**: In case one needs to scale out (such as adding several additional racks of compute infrastructure), additional spine nodes can be easily added to the existing deployment to increase the overall switching capacity for multiple racks of servers/compute/storage.

Arista Networks has added their twist to the "flatter" E-W traffic paradigm in Data center switching. The "Spline" is formed by using two switches in the middle of a row and link each server to the switch. (for redundancy, two server ports are connected to two switches.) From a performance perspective, there is now only one hop between servers.

Many companies have been re-architecting their networks to the leaf-spine design referenced in this section – the difference driving the adoption of Spline is that these organizations – with anywhere from 500-2000 servers need to segregate traffic for regulatory and/or compliance reasons.

The spline is a smaller-scale approach to leaf-spine, with all the advantages of speed and efficiency – along with the added bonus of drastically fewer stranded ports on each switch used.

Virtual switching Layer

Besides the traditional "physical" network switches/ nodes described above, virtual switching layer is one of key component of datacenter networking environment and responsible for packet forwarding to/from virtual machines. Virtual switch essentially works just like a physical Ethernet switch but this ground-up software implementation of switching logic can also be used to intelligently direct communication by inspecting the packets before forwarding them on.

Depending on what virtual switch is being used in the solution, network designers often have to work with an existing virtual switch environment. For example, VMWare implementations would have vSwitch as virtual switching layer, KVM implementations has open vSwitch as the virtual switching layer, and Microsoft has HyperV hypervisor, etc.

Some vendors embed their virtual switch implementation as part of their hypervisor platform and others include the same function as part of server's firmware. Since the features and throughput in a virtual switch varies based on implementation, it often becomes very important for network designers to be aware of capabilities of these virtual switching nodes to come up with best networking solution for the environment.

The Virtual switching layer essentially becomes very important when network design requires deployment of overlay network design, multi-tenant network design, or high throughput within the cloud environment.

“White Box” Switches

Widespread adoption of software based switching platforms such as Cisco Nexus 1000v, VMWare HyperV etc. has established software-switching as a viable and stable function whereby the switching function is divorced from underlying hardware where throughput requirements are limited.

This adoption of software-switching, along with new innovation in general purpose hardware (primary driven by Intel) is facilitating new startup companies and solutions that are looking to disrupting the data-center switching market as it exists today, is often referred to as “Whitebox Switching”. The key concept here is that the customer is buying off-the-shelf low-cost merchant silicon switching hardware, and leveraging switching software (or pre-tested Linux switching libraries) from two different vendors to create equivalent of leaf node and spine node functions. The same “white box” economics are also key to justify any business case of most Openflow-based SDN implementation that are discussed in subsequent sections of this whitepaper.

While the overall concept of “White box” has created pricing pressures on established switching vendors such as Cisco, and to a lesser extent, Arista, the feature parity, performance gaps (specifically for high through-put environment), Integrated support model, and necessary ecosystem has yet to be developed to make “White box switching” a mainstream alternative for data-center switching environment. For large deployments, complexity associated with Management and integration, along with the diminished price difference between white boxes and discounted, branded switches, does not usually offer a compelling enough reason for an enterprise to shift to the white-box mode. Having said that, there are some network consultancies that have developed real-world expertise in integrating existing infrastructure based on the “established” vendors, along with that of the emerging SDN/NFV players.

Design Considerations

The discussion above highlighted the key areas that data-center networking teams have had to focus on, along with the main aspects that influence the choice of vendors for underlying hardware. This section will highlight key design considerations for the implementation of a leaf-and-spine architecture from physical layer perspective. The architecture doesn't assume use of White box or branded switch for this discussion.

Data center scale

As discussed above, one of the key benefits associated with the leaf and spine architecture is ability to scale the datacenter infrastructure as requirements evolve. There are several factors that network team needs to consider when planning for forecasted scale requirements. This includes

- Port Density consideration when selecting the leaf and spine nodes
- Port Speed associated with Leaf and Spine nodes
- Understanding the Layer 2 as well as Layer 3 forwarding rate to determine and plan per overall fabric throughput
- Understanding the traffic profile of applications to define and implement over-subscription policies on the network.

Layer 2. vs. Layer 3 vs. Overlay

From requirements perspective, typically VM migrations and communications are typically needed within boundaries of a VLAN across multiple server pods. This in turn drives design requirement to make ensure all VLANs are already available to those server pods where VMs may get migrated or needs to communicate with. This is not always a simple task because hypervisor functions such as vMotion, etc., could move VM from one server pod to another for variety of reasons. In addition, extending the VLAN across wider segments of the data center requires the elimination or workaround for Layer 3 boundaries between server PODs. These common issues associated with Layer 3 design typically gets addressed by creating a Layer 2 overlay.

In general there are many overlay tunneling technology options (such as VXLAN, NVGRE, OTV, LISP etc.) to select from when considering creating overlays. Each vendor generally leverage one of these overlay approach as part of their Datacenter SDN solution that addresses the requirements above. It is important for network teams to understand that key focus area for many of these SDN solutions is basically software based datacenter network management solution. Each solution proposed by different vendors has advantages and disadvantages associated with these technology options as well as core software function provided by vendor in their solution.

At least for the specific requirement described above, VXLAN technology is slowly becoming a de facto standard. It is scalable, open-standard based approach for solving this specific issue where layer 2 overlays across Layer 3 underlay can be created to enable seamless communication or VM mobility across server pods within the datacenter.

Selecting an Overlay Approach

Business continuity, workload mobility, and the desire to decouple services from geographic location are all driving the advancement of overlay networking for the datacenter. Some of the technology options for overlay networking include VXLAN, NVGRE, OTV, EoMPLS, and others, each differing in how they create an overlay and hence catering to different use cases.

Business continuity, workload mobility, and the desire to decouple services from geographic location are all driving the advancement of overlay networking for the datacenter. Some of the technology options for overlay networking include VXLAN, NVGRE, OTV, EoMPLS, and others, each differing in how they create an overlay and hence catering to different use cases.

The “Underlay”, or switch-based approach, is more of a hardware-oriented approach for SDN implementation. In the underlay approach, network traffic is switched across the architecture using full IP addressing formats or VLANs. The SDN architecture, comprising of controller and North/Southbound APIs, sits on top of the Underlay network i.e. an existing data center fabric. Proprietary underlay solutions includes vendor-specific technologies like Cisco ACI, FabricPath, OTV, Juniper QFabric.

The “Overlay” approach allows the deployment of an additional tunneling and encapsulation scheme on top of the existing Switching fabric (the underlay model). This tunneling and encapsulation does not require any changes to the underlying Data Center network itself. Some networking features and functions are moved into overlays to control the data, flow or forwarding path such as server virtualization, L4-L7 load balancing, security, Openflow etc.

A hybrid overlay allows the combination of physical and virtual resources providing end-to-end visibility. In simple terms: the combination of the two approaches discussed above.

With the advantage of flexible routing with Overlays, paths are selected based on different metrics. For example, overlay selects paths based on latency whereas the underlay might try to balance load. **The reason for different approaches for routing is because the underlying**

switching network is traffic engineered to minimize overall network congestion whereas the Overlay minimizes end to end latency.

Table 2.4.3 Advantages and disadvantages of Overlay based approach.

Advantages	Disadvantages
• Overlay-based solutions can generally be implemented over existing networks with no changes required. • Decoupling of the virtual network topology from the physical network Infrastructure avoids issues such as limited MAC table size in physical switches. • Support for VM mobility independent of the physical network. If a VM changes location, even to a new subnet in the physical network, the switches at the edge of the overlay simply update mapping tables to reflect the new physical location of the VM. The network for a new VM can be provisioned entirely at the edge of the network. • Provision of communication [east–west (i.e. VM-VM) and north–south] while maintaining isolation between tenants. • Ability to manage overlapping IP addresses between multiple tenants - an important consideration to support multi-tenancy. • For controller-based NV solutions, the controller is not in the data path, and so it does not present a potential bottleneck.	• Virtual and physical networks are separate entities, possibly with separate service assurance solutions, policy management, provisioning, and control points. • As the virtual networks grow and evolve, the physical network does not automatically adapt to the changes. As a result, overlay-based solutions require a lightly oversubscribed or non-oversubscribed physical underlay network. • <u>Decreased Fabric Visibility:</u> The adoption of overlay technologies may decrease the visibility of the fabric as a whole because network constructs that exist in the overlay network are hidden from the underlay fabric, e.g. traceroute in the overlay will not report individual underlay hop counts. • <u>Increased Troubleshooting Complexity:</u> The mapping of the virtual topology on top of the physical topology needs to be investigated.

For simplicity, nodes in the overlay can be thought of as being connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network. The “underlay” or transport network provides a “service” to the overlay, and this service has well defined SLAs in terms of attributes such as quality of service and restoration times. A basic example of overlay/underlay approach would be Ethernet-over-SONET. Here underlying network is the physical is the optical network of SONET on top of which L2 functionality of Ethernet is implemented. Similarly, in SDN the approach is used to make things simpler as briefed in the preceding paragraph.

The following is a discussion of encapsulation techniques used to enable overlay networks:

VXLAN

Virtual eXtensible LAN, or VxLAN, virtualizes the network by creating a L2 overlay on a L3 network employing a MAC-in-UDP encapsulation (encapsulation of an Ethernet L2 Frame in IP). This enables the creation of virtualized L2 subnets that can span physical L3 IP networks. VXLAN enables the connection between two or more L3 networks and makes it appear like they share the same L2 subnet. This now allows virtual machines to operate in separate networks while operating as if they were attached to the same L2 subnet. This eases VM communication and allowing live migration transcending L3 boundaries as they happen to be in the same VXLAN segment.

The VXLAN has a 24-bit VXLAN Network Identifier which acts as an identifier for specific VXLAN segment during the overlap of IP and MAC addresses. This provides 16 million segments (2^24) for traffic isolation and segmentation, in contrast to the 4000 segments achievable with VLANs. The VXLAN encapsulation is performed through a function known as the VXLAN Tunnel End Point (VTEP), typically a hypervisor vSwitch or a physical access switch. Since MAC frames are encapsulated within IP packets, there is no need for the individual L2 physical switches to learn MAC addresses alleviating MAC table hardware capacity issues on these switches.

One of the reasons why VXLANs use MAC-in-UDP: All modern L3 devices parse 5-tuple (source IP, Destination IP, Source port, Destination port, Protocol type). VXLAN uses well known UDP destination port of 8472, the source UDP port can be any value. As a result, a VTEP can spread all the flows from a single VM across many UDP source ports. This allows for efficient load balancing across Link Aggregation Groups (LAGs) and intermediate multi-pathing fabrics in the case of multiple flows between just two VMs.

VxLAN is emerging as the most common method of network virtualization.

NVGRE

Network Virtualization Using Generic Routing Encapsulation, or NVGRE, allows the creation of virtual Layer 2 topologies on top of a physical Layer 3 network. NVGRE is an L2 overlay scheme over an L3 network. Similar to VXLAN, NVGRE enables the connection between two or more L3 networks depicting as if they are in the same L2 subnet. This now allows inter-VM communications or live VM migrations across L3 networks as they appear to be in the same L2 subnet. NVGRE and VXLAN have similar principles but different encapsulation modes.

NVGRE Ethernet frame encapsulation: depicting MAC-in-GRE encapsulation. A unique 24-bit ID called a Tenant Network Identifier (TNI) is added to the L2 Ethernet frame, using the lower 24 bits of the GRE Key field. This new 24-bit TNI, similar to VXLAN, now enables more than 16 million segments to operate within the same administrative domain, which is way over the 4k segments of vLAN.

The L2 frame with GRE encapsulation is then encapsulated with an outer IP header and finally an outer MAC address. A simplified representation of the NVGRE frame format and encapsulation is shown in Figure 2.4.3a below.

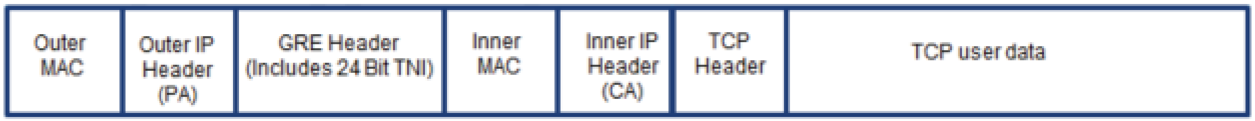

Outer MAC	Outer IP Header (PA)	GRE Header (Includes 24 Bit TNI)	Inner MAC	Inner IP Header (CA)	TCP Header	TCP user data

Figure 2.4.3a

The NVGRE endpoints are responsible for the addition or removal of the NVGRE encapsulation and can exist on a network device or a physical server. NVGRE endpoints perform functions similar to those performed by VTEPs in a VXLAN environment. NVGRE endpoints are gateways between the virtual and the physical networks. One common deployment is for the endpoint to be part of a hypervisor. The NVGRE endpoint, residing in the server or a switch encapsulates the VM traffic, adding the 24-Bit TNI, and sends it through a GRE tunnel. At the destination, the endpoint de-encapsulates the incoming packets and presents the destination VM with the original Ethernet L2 packet where the VM originating the communication is unaware of this TNI tag. With reference to the figure 2.4.3 the inner IP address is called the Customer Address (CA). The outer IP address is called the Provider Address (PA). When an NVGRE endpoint needs to send a packet to the destination VM, it needs to know the PA of the destination NVGRE endpoint.

Microsoft's initial implementation of NVGRE relies on L3 vSwitches whose mapping tables and routing tables are downloaded from the vSwitch manger via CLI based shell scripting.

OTV

Overlay Transport Virtualization is a Cisco proprietary protocol. OTV is optimized for inter-data center VLAN extension over the WAN or internet using MAC-in-IP encapsulation. In an OTV network, the OTV edge device is responsible for encapsulation and de-encapsulation of the OTV header and IP header and exists primarily on physical switches or routers.

It uses stateless tunnels to encapsulate L2 frames in the IP header and does not require the creation or maintenance of fixed stateful tunnels. OTV encapsulates the entire Ethernet frame in an IP and UDP) header, making the provider or core network transparent to the services offered by OTV (Figure 2.4.3b).

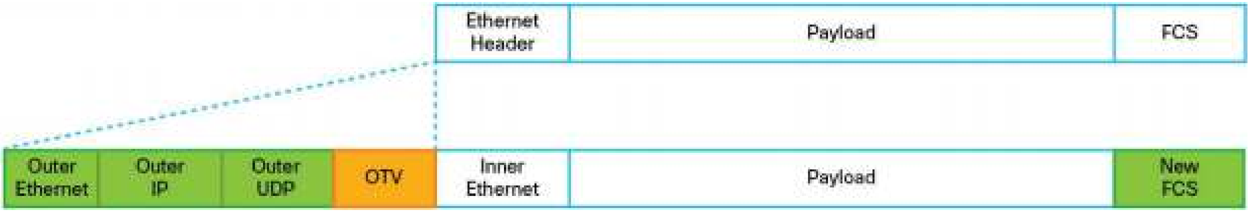

Figure 2.4.3b

In OTV, flooding of unknown destinations is prevented by its use of IS-IS routing protocol extensions. MAC-in-IP encapsulation supports L2 VANs over any transport as long as the transport can forward packets. Since OTV is both core and site transparent, no changes to the Layer 2 design of the sites are needed.

OTV can support multiple concurrent overlays and provides two levels of load balancing:

- Within the core: OTV headers are defined to allow the core to hash traffic based on five-tuples and distribute traffic over multiple paths to avoid polarization of encapsulated traffic.
- Within the site: OTV enables effective load balancing of flows across the multiple edge devices available in an all-active multihued deployment. Load balancing follows equal-cost multipart (ECMP) rules based on the information provided by the OTV control protocol.

OTV uses Ethernet over Generic Router Encapsulation (GRE) and adds an OTV shim to the header to encode VLAN information. The OTV encapsulation is 42 bytes, which is less than virtual private LAN service (VPLS) over GRE. The encapsulation is performed entirely by the forwarding engine in hardware.

Cisco Nexus ® 7000 Series Switches (M series) support OTV.

LISP

Locator/ID Separation Protocol is a Cisco proprietary protocol. LISP employs IP-in-IP encapsulation that allows end systems to keep their IP addresses even while moving around the network. LISP is designed to address the challenges of using a single address field for both device identification and topology location. In modern data centers, the mobility of endpoints should not result in a change in the end-host addressing, but simply the location of the end host. LISP addresses the problem by uniquely identifying two different number sets: routing locators (RLOCs), which describe the topology and location of attachment points and hence are used to forward traffic, and endpoint identifiers (EIDs), which are used to address end hosts separate from the topology of the network (Figure 2.4.3c)

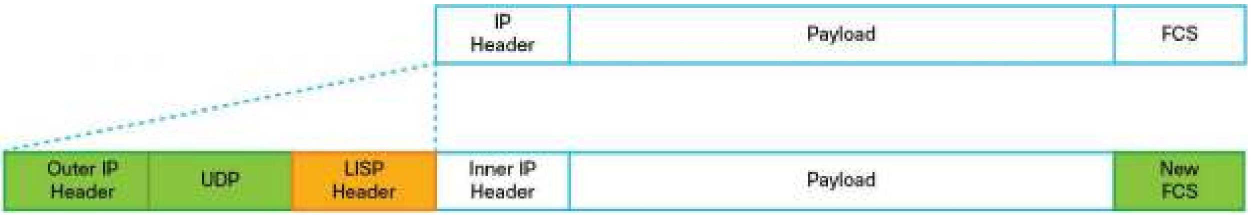

Figure 2.4.3c

LISP provides true end system mobility while maintaining shortest path routing of packets to the end system. IP endpoints such as VMs can be relocated anywhere regardless of their IP addresses. LISP also supports multi-tenancy with L3 virtual networks by mapping VRFs (Virtual Routing and Forwarding is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time.) to LISP instance IDs.

LISP is currently defined as a Layer 3 overlay scheme over a Layer 3 network, and it encompasses IPv4 and IPv6 for both the underlay and the overlay. Similar to other encapsulation schemes described previously, LISP provides a mechanism to help ensure virtual segment isolation through the addition of a 24-bit instance ID field in the LISP header, allowing more than 16 million virtual segments to be instantiated

MAC-in-MAC encapsulation

This scheme has been used to encapsulate end-user or customer traffic in the provider's MAC address header. Two implementations of MAC-in-MAC will be discussed, Shortest Path Bridging MAC-in-MAC (SPBM) and Provider Back Bone (PBB). Both SPBM and PBB follow the model of provider (transport) network that connects two or more customer (access) networks.

In an SPB network, packets are encapsulated at an edge node either in MAC-in-MAC 802.1ah frames and transported only to other members of logical network. IEEE 802.1aq supports unicast and multicast, and all routing is on symmetric shortest paths IEEE 802.1aq includes

SPBM functionality. IS-IS is the routing protocol that routes data by determining a best route for datagrams transmitted through a packet switched network. SPBM hardware switches are currently available from several vendors including Avaya and Alcatel-Lucent.

PBB network is a L2 bridged network that uses MAC-in-MAC encapsulation to transfer user L2 traffic between two or more L2 networks that are located at the edge of the PBB network. It can be noted that PBB network includes all networks that use MAC-in-MAC encapsulation including SPBM. PBB network typically includes a Back-bone Edge Bridge (BEB) and a Back-bone Core Bridge (BCB). BEBs, also known as provider network edge nodes function as devices that enable transfer packets to/from interfaces within/outside the PBB network. BCBs, also known as provider core nodes enable transfer of packets between interfaces that are within the PBB network. Cisco provides PBB solutions.

STT

Stateless Transport Tunneling is another overlay technology for creating L2 virtual networks over L2/3 physical network within a data center. Conceptually, there are a number of similarities between STT and VXLAN; the tunnel endpoints are typically provided by hypervisor vSwitches, the VNID is 24 bits wide and the transport source header is manipulated to leverage multipathing. STT encapsulation differs from VXLAN in two ways; it uses stateless TCP-like header inside the IP header that allows tunnel endpoints within end systems to take advantage of TCP Segmentation Offload (TSO) capabilities of existing TCP/IP Offload Engine (TOE) server NICs.

One area that STT specifically addresses is the size mismatch between Ethernet frames and the maximum transmission unit (MTU) supported by the underlying physical network. Most end-host operating systems today set the MTU at a small size so that the entire frame plus any additional (overlay) encapsulations can be transported over the physical network. This setting may result in a potential performance degradation and additional overhead compared to frames that can be transmitted with their desired maximum segment size (MSS). Host-based overlay networks address many of the challenges posed by rigid underlay networks and their associated protocols (Spanning Tree Protocol, etc.), but the overlay network needs to be integrated with the physical network.

A major and unfounded assumption about host-based overlay networks is that the underlying network is extremely reliable and trustworthy. However, an overlay network tunnel has no state in the physical network, and the physical network does not have any awareness of the overlay network flow. A feedback loop is needed from the physical network and virtual overlay network to gain end-to-end visibility into applications for performance monitoring and troubleshooting.

The initial implementation of network virtualization using STT from VmWare/NSX are based on Open vSwitches and a centralized control plane for tunnel management via downloading mapping tables to the vSwitches.

Part 2: Centralized vs. Distributed control plane for Datacenter Networks

With the hype of SDN/Openflow and the inherent need for existing vendors in data-center switching space to respond to customer demands, simplification and flexibility have become key areas of innovation for many vendors. Before going into the details of what new solutions are available in this space, it is important to describe the difference between centralization and distributed control plane architecture.

In general, Centralized control plane architecture refers to a model where control plane of a router or a switch is physically separate from data-plane fabric. The Control-plane acts as a centralized layer that has complete topological view of the entire forwarding plane it controls. In addition control plane is typically responsible for programming forwarding rules on data-plane elements by instructing the data-plan to "match" incoming packet and perform defined "action" on the packet.

During the latter half of 2012, Open Networking foundation (ONF) came up with this conceptual architecture along with the first draft of openflow protocol that is used to communicate between the two planes. At this time the openflow protocol has evolved to its 4th release. While openflow in general had its fair share of visibility in recent years, it is important to understand that openflow by itself does not solve any new unique networking issue other than bringing simplicity by having centralized node for configuration and policy management.

This concept has spurred new waves of innovation in the networking industry which are no longer limited to the data-center environment. Currently, there are four types of innovations we see in the market:

- Limited feature Open-flow SDN controller that leverages low-cost merchant silicon switching fabric.
- Full featured Control plane from traditional vendors that leverages low-cost general purpose x86 platforms as data plane.
- Management-focused SDN Controller
  Open API environment on switches/routers with an integrated control architecture
- Centralized Control plane vs. Distributed control plane.

The following are examples of Controllers that have emerged in the market (2013-14). [**Apologies in advance to all the Marketing and Engineering talent behind these products for the brevity. This info is meant as a basic introduction and peek at your reference architecture. There is a LOT more to these products than there is room to mention here.]**

Tail-f Network Control System (NCS) (www.tail-f.com)

NCS provides a single network-wide interface to all network devices and all network applications and services, as well as a common modeling language and datastore for both services and devices. A transaction engine handles transactions from the operations at the service layer to the actual deployment of configuration changes in the network.

Designed to be a generic solution, NCS supports the implementation of network applications and service on a wide variety of networking devices, both traditional hardware-based devices and virtual software appliances. The extensive library of Network Element Drivers that are used to abstract the physical network from the controller, make NCS one of the leading platforms in the SDN arena today.

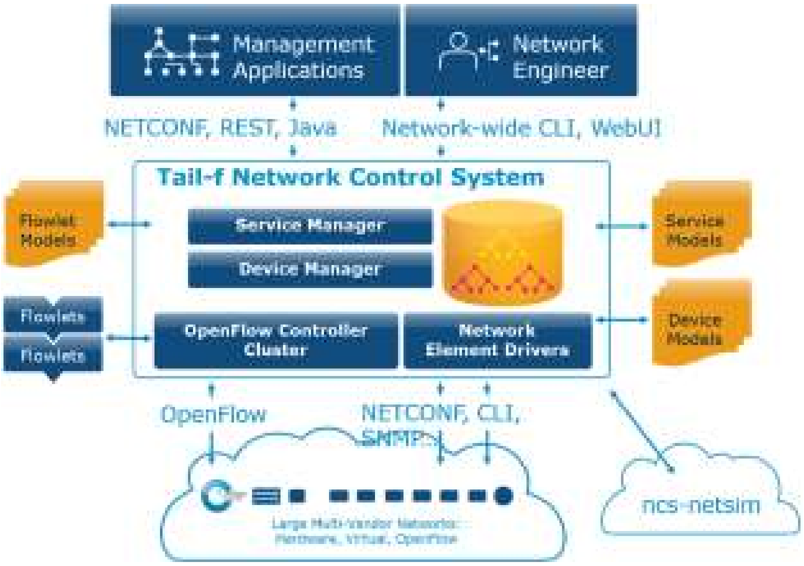

Key Features of NCS

- Model-driven Architecture: Network services and devices are modeled into YANG and their configuration states are centrally stored in NCS. Model changes are easy to adapt. NCS automatically renders the device interface from the vendor’s device models. NCS also renders the complete service management functionality from the service models that are specified by the Service Provider.

- Service-Aware: NCS provides a declarative way to specify how a network service shall be applied to the network infrastructure. This greatly facilitates the mapping of service configuration changes to device configuration commands. The entire service life cycle is supported including creating, modifying and deleting service instances.
- Programmability: NCS provides a rich set of transaction-safe and model-driven Application Programmatic Interfaces (API) such as REST, Java, JavaScript, and XML. All of them work in a publish-subscribe model so that other systems can be kept in sync in real-time.
- Transactions: NCS applies all service changes towards the network as an atomic change-set. This ensures that the network is always in a consistent state and can automatically recover from failed configuration changes.
- Consistent State: NCS represents the true current state of the network services, the device configurations and the mapping between services and devices through a central data-store in real-time.

C-Plane Networks (www.cplanenetworks.com)

C-Plane Networks **DVN** is fully integrated with OpenStack Icehouse. The product seamlessly replaces the standard Neutron VLAN networking with a high-performance VXLAN-based solution. DVN works with your existing network topology to quickly and reliably deliver network services for your OpenStack compute and storage projects.

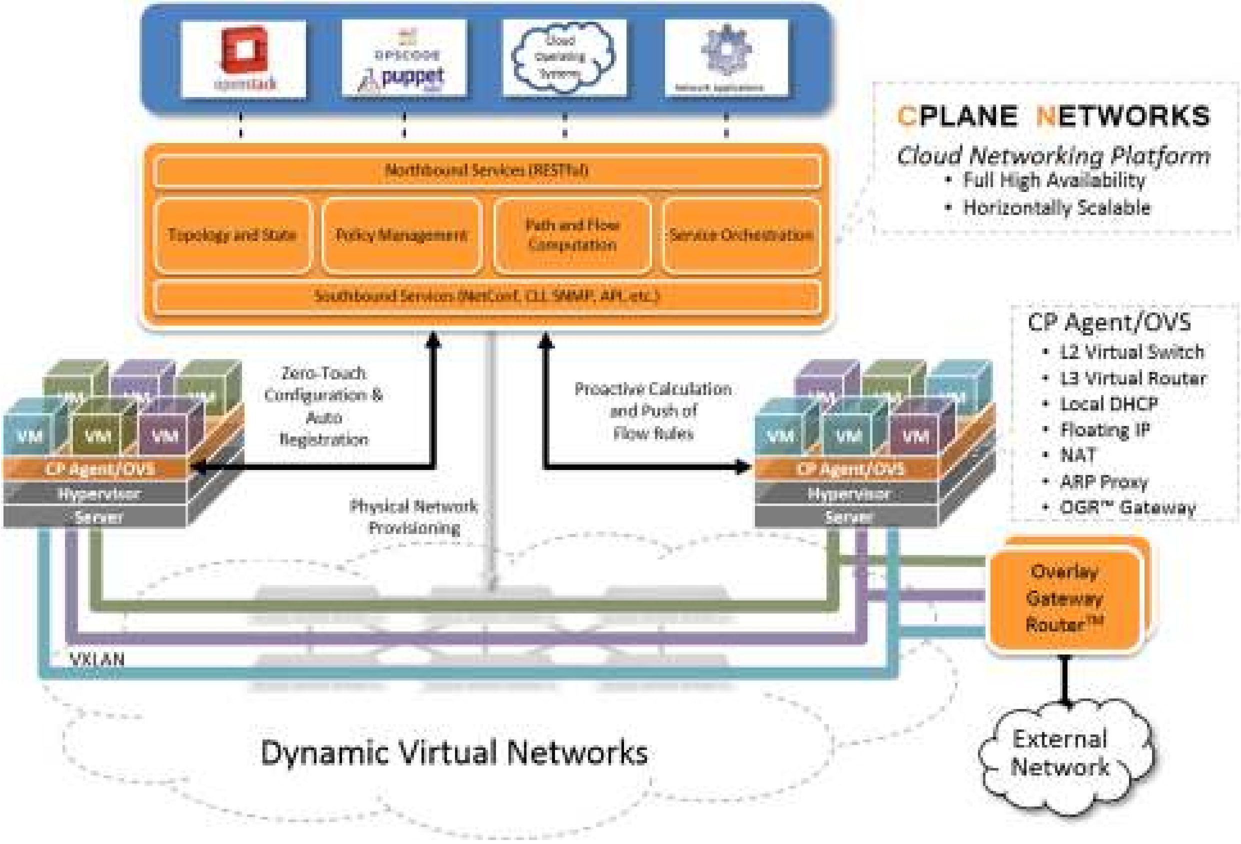

CPLANE NETWORKS **DVNi** lets you quickly and easily deploy global Layer 2 and Layer 3 VPNs over fully traffic-engineered MPLS networks with end-to-end bandwidth guarantees and quality of service that are defined, deployed and managed using powerful policy-based service class of service definitions. Now you can tune your network to the needs of your applications.

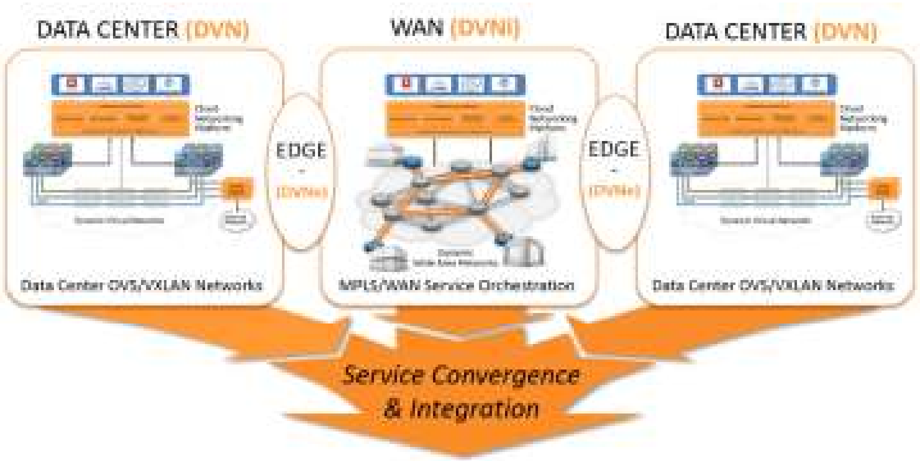
DATA CENTER (DVN)
WAN (DVNi)
DATA CENTER (DVN)
Cloud Networking Platform
EDGE
(DVNe)
EDGE
(DVNe)
Data Center OVS/VXLAN Networks
MPLS/WAN Service Orchestration
Data Center OVS/VXLAN Networks
Service Convergence
& Integration

End-to-End Networks

Scalable OpenStack network virtualization in the data center

Seamless, interconnected networks between data centers (WAN)

NFV Orchestration

Software-only Network Orchestration

- Dynamic Virtual Networks
- Policy and structure network orchestration
- Physical network integration and optimization
- Converged Virtual LAN and WAN
- Orchestrate NFV Services

OpenStack Networking

- Production ready Neutron Plugin

SDN Customization and Integration

- Sophisticated SDN platform allows customized solutions inside and outside the DC

Products:

Dynamic Virtual Network (DVN)

- Operational efficiency, reliability and secure multi-tenancy of OpenStack® networking
- Scale OpenStack network performance through elimination of OpenStack bottlenecks
- Close the gap between NetOps and DevOps through common Application-aware network

MPLS/WAN Network Provisioning and Orchestration (DVNi)

- Creates Multi-datacenter Wide Area Networks (WAN)
- Provides L2/L3 VPN with dynamic Class of Service/Quality of Service
- Provides optimal network utilization through patented bandwidth management

Plumgrid

PLUMgrid ONS for OpenStack is provides a fully distributed set of networking. Designed to scale from single-rack deployment to multi-rack OpenStack cells, PLUMgrid ONS for OpenStack gives you the extensibility you need to grow your offering as your needs grow. PLUMgrid ONS for OpenStack is tightly integrated with the building blocks in the OpenStack cloud.

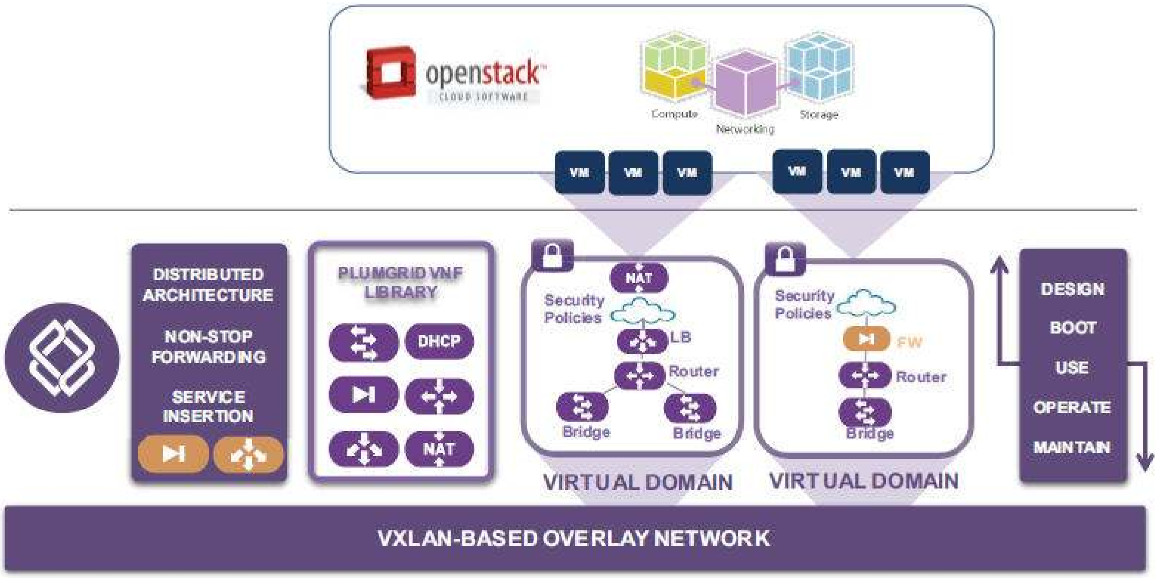

PLUMgrid OpenStack Networking Suite is based on a fully distributed architecture that is built for scale. Since forwarding decisions are distributedandmadeateachindividualserver,every newserveraddedto the cloud increases the cloud’s forwarding capacity. The product can deliver terabits of scale out performance by leveraging hardware offload capabilities within industry standard x86 servers andNetwork Interface Cards(NICs).

The PLUMgrid Director cluster provides a flexible and distributed control plane. This creates a fully distributed and high availability cloud network infrastructure. If a situation arises that causes the control plane tobecome unavailable, data forwarding continues without interruption because all forwarding decisions are distributed.

PLUMgrid OpenStack Networking Suite provides built-in network functions such as switch/bridge, router, NAT, IPAM, DHCP, security policies and end-to-end encryption.

PLUMgrid Virtual Domains are logical data centers that provide a comprehensive set of networking features and advanced description of policies for yourcloud.

PLUMgrid provides end-to-end data plane encryption for all network trafficwithin Virtual Domains. This unique capability provides cryptographic isolation across tenants, removes threats of cross contamination and security exposure.

PLUMgrid allows customers to easily add virtual network functions thatare not available in OpenStack Networking (Neutron), allowing them to immediately deploy OpenStack in production environments. Included are network functions (such as routers, switches and security policies, etc.)and services (such as NAT, IPAM and DHCP).

By providing ISSU capabilities, PLUMgrid is able to offer “always on”virtual network infrastructure by eliminating scheduled outages. ISSUreduces operational costs, eliminates cloud network downtime, and allows faster implementation of new features and mitigates security risks with timely fixes transparent to cloud users.

A library of application blueprints customized for specific applicationsstacks are included in PLUMgrid OpenStack Networking Suite. Theseblueprints accelerate deployment of these applications on OpenStackby taking the guesswork out of the network interconnect and services required to deploy these applications in a secure and scalable manner.

Cyan – Blue Planet (www.cyaninc.com/products/blue-planet-sdn-platform)

Cyan Blue Planet is a carrier-grade SDN and NFV orchestration platform built specifically for network operators and the WAN. Some major features are the following:

- Carrier-grade, multi-vendor SDN and NFV orchestration platform built specifically for network operators
- Includes a family of applications for simplifying end-to-end service deployment and virtual resource orchestration across telco data centers and the WAN
- Centralizes control and management of physical network elements and virtual SDN/NFV-based service resources
- Programmable architecture supports open northbound and southbound APIs to enable integration with existing systems, and accelerate service and application development
- Manages, automates and orchestrates services that leverage physical and/or virtual resources across the telco cloud and the WAN
- Flexible architecture integrates NFV, Cloud, and Multi-Domain Service Orchestration capabilities
- Open system ensures interoperability with different OSS platforms, cloud management systems, SDN controllers, network elements and VNFs
- NFV capabilities comply with ETSI’s NFV ISG Management and Orchestration (MANO) framework
- Powerful, intuitive html user interface provides a single-pane-of-glass for managing all service resources
- Support for multiple cloud management systems such as OpenStack and VMwar

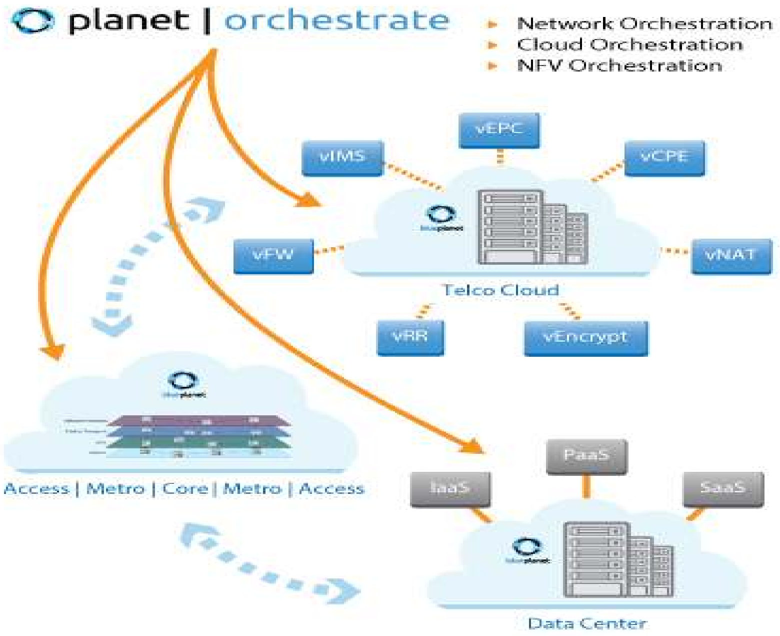

Overture Networks (www.overturenetworks.com)

The Ensemble Service Orchestrator/Ensemble Controller are open, extensible carrier-class NFV service lifecycle management and orchestration systems that coordinate virtual resources and physical network elements to create, activate and assure services using one or more virtual network functions.

ESO uses the OpenStack™ cloud controller – bundled with ESO – to manage the virtual compute environment, including virtual machines, virtual switches and top-of-rack data center switches. For management of the physical wide area network traffic flows, ESO leverages Overture's Ensemble Network Controller, but it can also be integrated with other third-party network controllers. Some major features include:

- Open northbound and southbound REST APIs
- Flexible, policy-driven workflow engine
- VNF agnostic management and orchestration
- Dynamic NFV service optimization
- Control multiple cloud environments

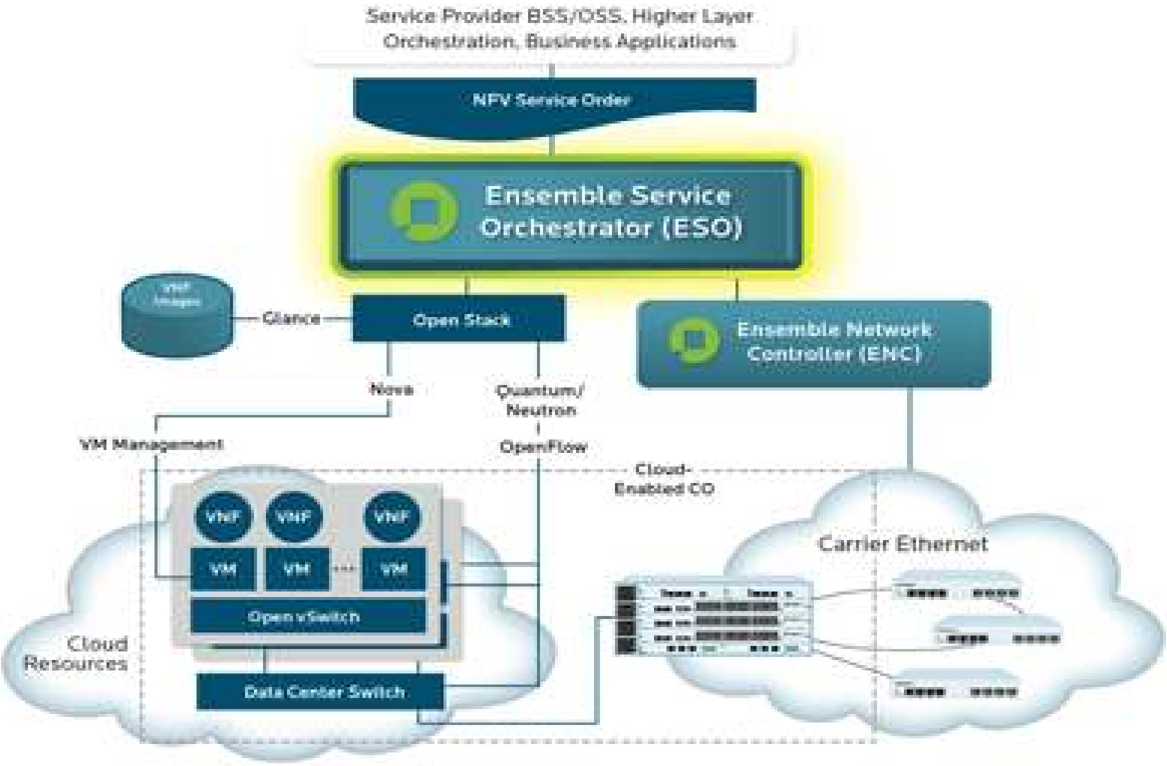

Nuage Networks VSP (www.nuagenetworks.net/products)

The Nuage Networks Virtualized Services Platform (VSP) is the foundation for open and dynamically controlled datacenter network fabrics aimed at cloud service providers and webscale operators.

Nuage VSP’s positioning is as follows:

- Making the datacenter network as dynamic & consumable as compute infrastructure through automated instantiation of network services.
- Eliminating cumbersome configuration-driven processes for datacenter networking.
- Separating and simplifying the definition of network service requirements and policies from the manner in which network services are established.
- Seamless connectivity across ANY datacenter network infrastructure (Layer2 through Layer4) incorporating both virtualized and non-virtualized compute environments
- Scaling to meet the demands of thousands of tenants with unique application requirements and enterprise policies.

The Nuage Networks VSP is comprised of three key software-based products:

1. Virtualized Services Controller (VSC): Datacenter network control plane
2. Virtualized Services Directory (VSD): serves as a policy, business logic & analytics engine for the abstract definition of network services, operating through RESTful APIs.
3. Virtual Routing & Switching (VRS): module that serves as a virtual endpoint for network services. Through the VRS, changes in the compute environment are immediately detected, triggering instantaneous policy-based responses in network connectivity to ensure that the needs of applications are met.

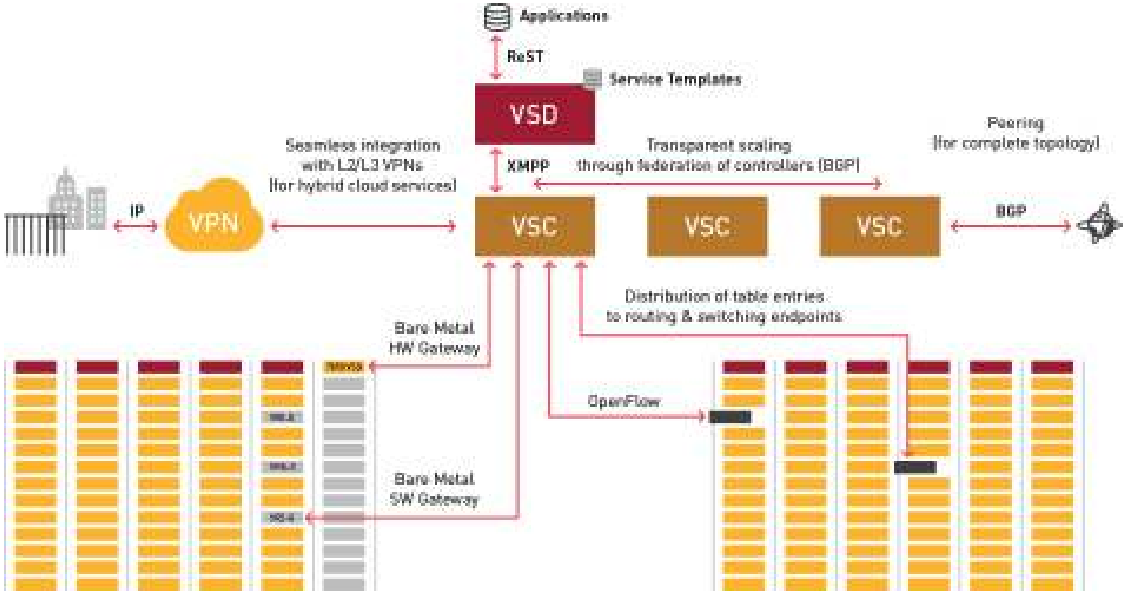

Juniper Contrail (www.juniper.net)

Contrail is Juniper's offer in the SDN orchestration and automation arena. Contrail is an open software defined networking solution that automates and orchestrates the creation of highly scalable virtual networks. Contrail is a scale-out virtual networking solution that integrates with physical routers and switches to eliminate the challenges of private and public cloud networking. Service providers can use Contrail to enable a range of innovative new services, including cloud-based offerings and virtualized managed services. Enterprises can use Contrail to increase business agility by enabling the migration of applications and IT resources to more flexible private or hybrid cloud environments.

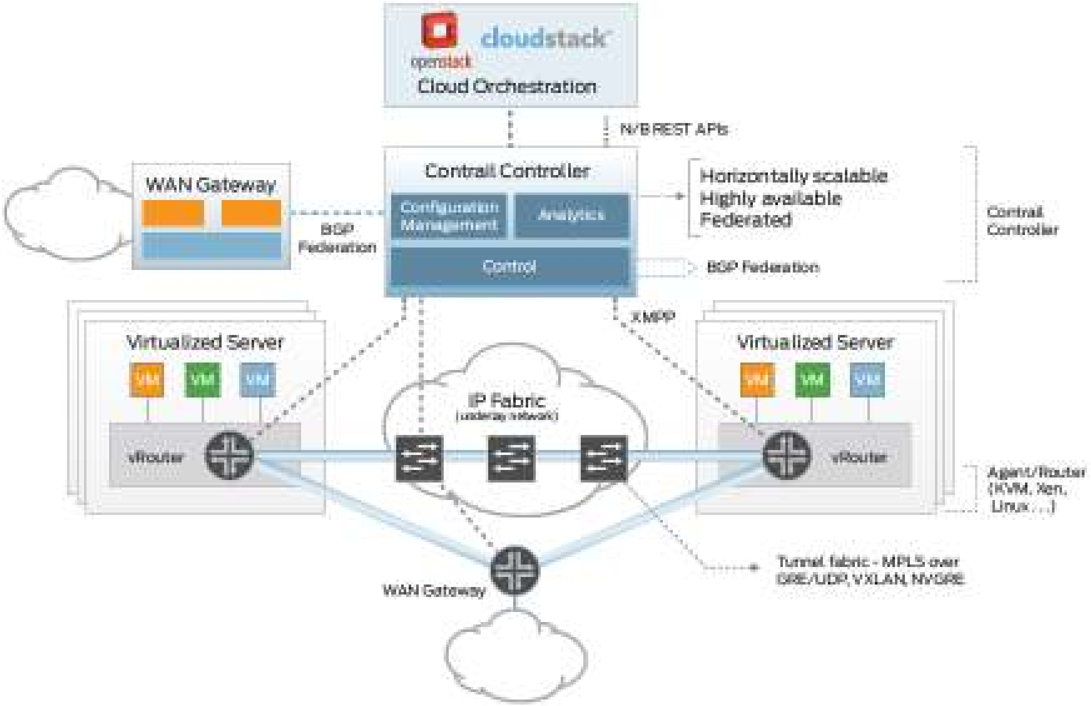

Open Daylight (www.opendaylight.org)

OpenDaylight is an open platform for network programmability to enable SDN and NFV for networks at any size and scale. The second release, "Helium", comes with a new user interface and a much simpler and customizable installation process thanks to the use of the Apache Karaf container.

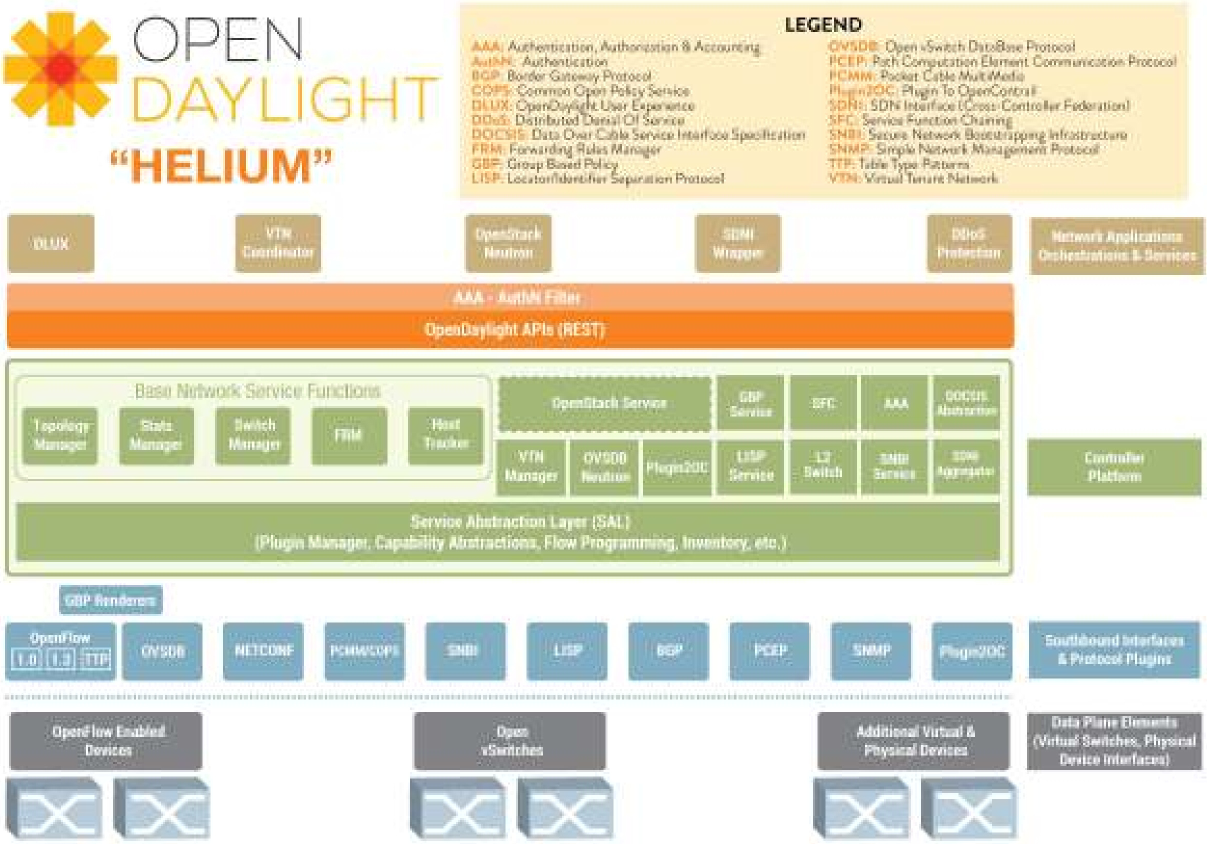

OPEN DAYLIGHT "HELIUM"

LEGEND

AAA: Authentication, Authorization & Accounting
AuthN: Authentication
BGP: Border Gateway Protocol
COPS: Common Open Policy Service
DLUX: OpenDaylight User Experience
DDoS: Distributed Denial Of Service
DOCSIS: Data Over Cable Service Interface Specification
FRM: Forwarding Rules Manager
GBP: Group Based Policy
LISP: Locator/Identifier Separation Protocol
OVSDB: Open vSwitch DataBase Protocol
PCEP: Path Computation Element Communication Protocol
PCMM: Packet Cable MultiMedia
Plugin2OC: Plugin To OpenContrail
SDNi: SDN Interface (Cross-Controller Federation)
SFC: Service Function Chaining
SNBI: Secure Network Bootstrapping Infrastructure
SNMP: Simple Network Management Protocol
TTP: Table Type Patterns
VTN: Virtual Tenant Network

For those looking to manage their networks using OpenDaylight, there is deeper integration with OpenStack, including significant improvements in the Open vSwitch Database Integration project.

OpenDaylight software is a combination of components including a fully pluggable controller, interfaces, protocol plug-ins and applications. This common platform allows both customers and vendors to innovate (see Brocade) and collaborate in order to commercialize SDN- and NFV-based solutions.

In the current release, the platform has evolved in other key areas including high availability, clustering and security, as well as strengthening and adding new protocols such as OpenFlow Table Type Patterns, PacketCable MultiMedia, an application policy framework and tools for Service Function Chaining – similar to what you would see in Overture's and Tail-f NCS' solutions.

Brocade (OpenDaylight-based)

The Brocade Vyatta Controller is a quality-assured edition of the OpenDaylight controller code, which includes the tools and services to quickly implement software-defined networks.

The Brocade Vyatta Controller is the first commercial controller built directly from OpenDaylight code, without any proprietary extensions or platform dependencies. Users can freely optimize their network infrastructure to match the needs of their workloads, and develop network applications that can be run on any OpenDaylight-based controller.

The Brocade Vyatta Controller provides:

- A common SDN domain for multi-vendor networks and virtual machines
- Single-source technical support for Brocade SDN Controller domains.
- A smooth on-ramp to SDN adoption with an easy-to-use GUI, installation tools and expert developer support
- Pre-tested packages and services optimized to the different needs of service providers and traditional network operations
- Complete portability for OpenDaylight applications developed on the Brocade Vyatta Controller
- Centralized Management function with new switching fabric

Cisco ACI

Cisco ACI is application aware, enables dynamic application instantiation and removal, and is capable of supporting physical, virtual, and cloud integration with full management visibility. The driver of its speed and efficiencies is the common policy-based operating model ACI employs across ACI-ready network and security elements.

ACI provides a common policy and management framework that enables automatic infrastructure provisioning based on application policy profiles. In the ACI model, networked infrastructure becomes a flexible and programmable pool of stateless resources ready to be provisioned for new applications and services.

A key architectural component of ACI is the Cisco Application Policy Infrastructure Controller (APIC), which provides a single touch point for all configuration, management, and operational tasks, including policy definition and health monitoring. By providing a common operational framework, it unifies applications, networking, cloud, and security teams in defining application requirements. These requirements are defined via an application network profile (ANP), which consists of the logical representation of all the application infrastructure requirements, connectivity, and network services that define their interdependencies. When an application is ready to be deployed, the APIC uses the profile to automatically provision the required infrastructure resources and services. This simplifies the operation and reduces infrastructure configuration and application deployment times.

Arista EOS (www.arista.com/en/products/eos)

Arista EOS is the core of Arista’s cloud networking solution for next-generation data centers and cloud networks. Cloud architectures built with Arista EOS scale to tens of thousands of compute and storage nodes with management and provisioning capabilities that work at scale. Through its programmability, EOS enables a set of software applications that deliver workflow automation, high availability, unprecedented network visibility and analytics and rapid integration with a wide range of third-party applications for virtualization, management, automation and orchestration services.

Arista Extensible Operating System (EOS) is a fully programmable and highly modular, Linux-based network operation system, using familiar industry standard CLI and runs a single binary software image across the Arista switching family. Architected for resiliency and programmability, EOS has a unique multi-process state sharing architecture that separates state information and packet forwarding from protocol processing and application logic.

Cloud Scale Architecture

- Scale to your needs, from 100 to 100,000+ compute and storage nodes
- Rich management and provisioning capabilities that work at scale
- Open, standards-based approach with MLAG at Layer 2, ECMP at Layer 3 with effective use of all available bandwidth in non-blocking modes while providing failover and resiliency
- Network virtualization using tunneling technologies such as VXLAN and NVGRE for seamless workload mobility
- Innovative Spline™ architecture for high density hosts in a single-tier or two-tier network

Open and Programmable

- Open integration with all application and infrastructure elements via eAPI and Advanced Event Manager (AEM)
- Programmable at all layers: Linux kernel, hardware forwarding tables, Virtual Machine orchestration, switch configuration, provisioning automation and advanced monitoring

- EOS Application Extensibility for the ability to run cloud infrastructure automation applications (such as Chef, Puppet or Ansible) and network analytics applications (such as Splunk)
- Easily adaptable to in-house network management systems
- Key enabler of Arista EOS software applications for automation and visibility, such as Zero Touch Provisioning, VM Tracer, and Latency Analyzer (LANZ).

High Availability

- Reduce maintenance windows with Arista Smart System Upgrade (SSU) reduces maintenance windows through intelligent insertion and removal of network elements
- In-service software upgrade (ISSU) for individual processes within EOS
- Self-healing resiliency for minimum downtime with fault containment to a single module and process restart without the need to rebuild state information
- Custom monitoring, failover and load balancing with third-party integration for custom monitoring, failover and load balancing

Visibility

- Unprecedented visibility into application performance and network-wide monitoring capabilities for both industry standards and customer specific dev/ops solutions
- Simplified Tap Aggregation with the Arista Data Analyzer (DANZ) feature set
- Rapid identification and troubleshooting of application and network performance problems through tracers such as VM Tracer, Latency Analyzer (LANZ), MapReduce Tracer, sFlow and Path Tracer

Automation

- Simplified provisioning for new and replacement switches with Zero Touch Provisioning (ZTP) and Zero Touch Replacement (ZTR)
- Advanced Event Management for automated responses to network and application events
- Automate complex IT workflows and simplify network operations to individual requirements through rich programmatic capabilities
- Automation Integration with partners enhance native capabilities with tools such as Puppet, Chef and Ansible and extends automation up the stack to include other network systems and applications including firewalls, load balancers and compute infrastructures with partners such as F5 and VMware.

As you can see, there are many possible ways to implement the control and orchestration functions across an SDN-enabled network. Which one you select depends entirely on the existing network infrastructure, current/future needs, application demands, budget, internal expertise and training levels.

Additionally, there are network security considerations that must be taken in to account – The specific methodologies and techniques used to secure the next gen data center are beyond the scope of this whitepaper. Hint: examine methods of distributed security solutions that are more attuned to East-West datacenter traffic rather than North-South flows.

VmWare NSX: (www.vmware.com/products/nsx)

NSX offers a distributed logical architecture for L2-7 services including, logical switch, router, firewall, load balancer and VPN. These logical network services are provisioned programmatically when virtual machines are deployed and move with virtual machines. Existing applications operate un-modified and see no difference between a virtual network and a physical network connection.

NSX exposes a RESTful API, allowing cloud management platforms to automate the delivery of network services. Because network services are now delivered to applications by the virtual network, no manual reconfiguration of physical network devices is necessary.

NSX Service Composer offers a way to automate the consumption of services and their mapping to virtual machines using logical policy. Customers can assign policies to groups of virtual machines and as more virtual machines are added to the group, the policy is automatically applied to the virtual machine. Customers can build advanced workflows automating security, compliance and network provisioning including load balancing and firewall rules.

Where NSX' platform sits in the stack:

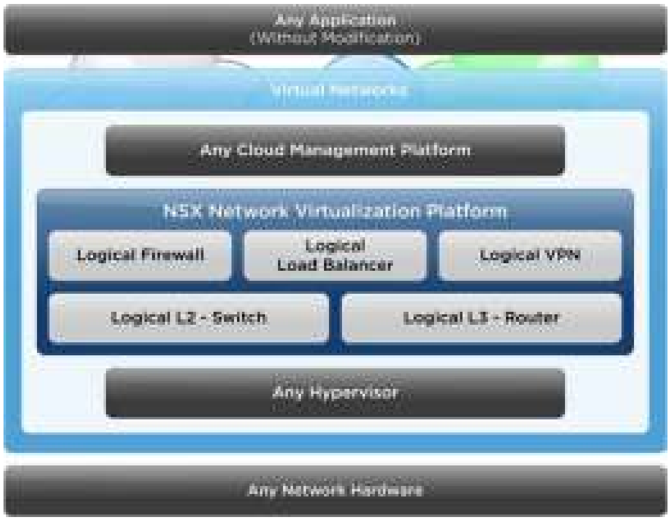

Key Features of NSX

- Logical Switching – Reproduce the complete L2 and L3 switching functionality in a virtual environment, decoupled from underlying hardware
- NSX Gateway – L2 gateway for seamless connection to physical workloads and legacy VLANs
- Logical Routing – Routing between logical switches, providing dynamic routing within different virtual networks.
- Logical Firewall – Distributed firewall, kernel enabled line rate performance, virtualization and identity aware, with activity monitoring
- Logical Load Balancer – Full featured load balancer with SSL termination.
- Logical VPN – Site-to-Site & Remote Access VPN in software
- NSX API – RESTful API for integration into any cloud management platform

Completely decoupled from physical network hardware, Network virtualization works as an overlay above any physical network hardware and works with any server hypervisor platform. NSX Gateway allows legacy VLANs and physical hosts to be mapped into virtual networks. Reproduce the physical network model in software. NSX reproduces the entire networking environment, L2, L3, L4-L7 network services, in software within each virtual network.

NSX offers a distributed logical architecture for L2-7 services including, logical switch, router, firewall, load balancer and VPN. These logical network services are provisioned programmatically when virtual machines are deployed and move with virtual machines. Existing applications operate un-modified and see no difference between a virtual network and a physical network connection.

Part 3: Virtualized Network Services Functions and Service Chaining

Network Functions Virtualization (NFV) and Service Chaining are interrelated concepts that drastically alter the cost and operational equation for both enterprises and service providers.

Network Functions Virtualization (NFV) is the instantiating of functionality found in routers, switches and other devices (Layer 4-7 appliances for example) into a virtual domain running on a x86-class computer, or "bare metal" Ethernet switch, depending on the product or solution

Service Chaining refers to the sequencing of Layer 4- Layer 7 network functions such as firewalls, load balancers, traffic shapers, DHCP and NAT. Traditionally, the functions implemented through service chains have necessitated the use of multiple appliances, with all of the attendant hardware, cabling and individual configuration that goes along with it. Additionally, when there are changes to be made, whether in an enterprise environment for a branch office for example, or managed services environment with a Service Provider, the configuration changes are prone to delay and error when handled in a primarily manual process.

For both enterprise and service provider clients, one approach that will provide security, agility and "speed to market" is to automate the Service Chaining functionality and process with the appropriate orchestration and Network Functions Virtualization Solution. Key capabilities to look for are:

- An ability to set up packet-forwarding rules that steer traffic through the appropriate network functions by flow. This also involves configurations of the switches and routers doing the traffic steering.
- Reconfiguring the specific L4-L7 functions in the chain with the appropriate processing rules for each flow.
- Providing life cycle operations for each chain; i.e. adding, upgrading or removing the :4-L7 network functions.

For a Service Provider, a solution set that provides NFV functionality for routing switching and the upper layer functions, along with a central orchestration controller can help drastically reduce the time to provision a new or changed Manage Service. This will enable a faster time to revenue for the Provider, and faster up-time for the end-Customer.

For an enterprise, managing these upper layer functions has always required additional configuration efforts and typically slow response times in setting up initial branch/remote office locations, along with the added expense and management overhead associated with multiple appliances for those functions.

Using SDN/NFV solutions abstracted to the physical network, and remotely orchestrated and controlled, offers a significant savings in both CapEx and OpEx across both an enterprise and Managed Service Provider's infrastructure.

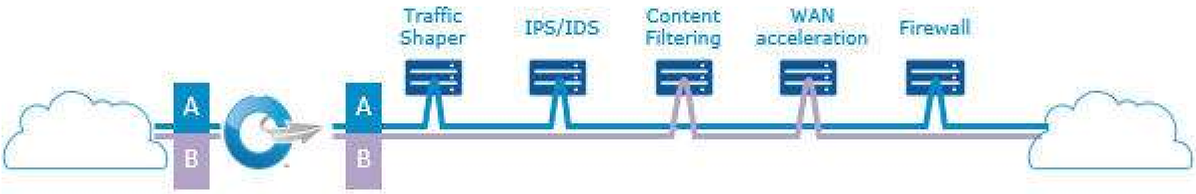

Some of the main advantages of NFV are as follows*:

- Improved hardware efficiencies compared to that of dedicated hardware implementations through the use of COTS hardware and commodity servers.
- Improved flexibility in assigning Virtual Network Functions (NFV's) to hardware. This adds to functionality and allows for the decoupling of the functionality form the physical location, whether at centralized Data Centers, POP's or customer premise.
- Drastically faster service implementation through the deployment of software.
- Improved Opex costs due to common automation and operating procedures
- Standard and open interfaces between VNF's and the infrastructure and associated management entities so that the decoupled components can be provided by multiple vendors.

**ETSI GS NFV 002*

Part 4: Network Orchestration

Since the key area of functionality in an SDN network is orchestration, the degree to which you can discover an existing “actual” or virtual, network element, manage its configuration and configure multiple devices to create your end-end solution, is one of the most important aspects of architecting a SDN-centric solution.

Up until the advent of controllers performing the orchestration function, the issues facing operators were: multiple EMS’s (”EMS Sprawl”), lack of service agility, hard coded service definitions in existing OSS’s, and, with the growth of VNF’s, translations of these services and configurations to Virtual Network Functions requiring a very time- and labor-intensive translation process.

Today’s Controller systems are literally transforming the back office with the ability to provide a high degree of automated service provisioning, orchestration and VNF control – without the army of developers. The controller helps to eliminate EMS sprawl, while simplifying the orchestration and OSS. Through providing vehicle to apply a standards and models-based approach (YANG, etc), the Orchestration platform allows dynamically definable network applications and automated translation to VNF operations for NFV “virtual devices”.

Through the application of NetConf and YANG modeling, services and devices are instantiated into an orchestration solution that enables operators to gain significant advantages in reducing OpEx and the time it takes to roll out new or updated services to customers. Operators using this type of controller solution will enjoy the following benefits:

- Faster Development and deployment of new services with a model-based approach
- Approach allows integration of multi-vendor environment in fractions of the time versus traditional methods
    - Services, device configurations, open flow apps defined in YANG, a standards-based modeling language.
    - Industry standard MEF CE2.0 Services are modeled and can be introduced to the network more quickly
- No disparate Element Management Systems
- Real time, dynamic capacity allocation
- Transition to and management of mixed environment of traditional hardware and software defined virtual devices and services.

Summary

In summary, there are a number of technologies and products available today that will help your organization migrate to the new network architecture. From decreasing the time to delivering a service within an enterprise to faster “time-to-revenue” for a network operator, the techniques and products we’ve described in this white paper are available and either in production or proof of concept environments.

From an economic standpoint, the reduction in operating costs to deliver the same amount of IT services, or the ability to grow at a certain rate without attendant head count increases is starting to manifest itself in what is becoming the widespread deployment of this technology. The promise of the ability to provide network agility and help operators make changes to the network and develop/deploy new services drastically faster than in the past is also coming to fruition.

For more information on these products, or engineering and consulting services to help you sift through the abundance of options available, please visit our Resources Center at http://datavision-inc.com/knowledge/. Here you will find Blog posts, White papers, with links to products and Industry web sites to expand your knowledge and awareness of what’s out there in the world of SDN and how it can help your business succeed.

The “science experiments” have ended, and, coupled with Openstack and other software defined systems, the Software Defined Data Center and Software Defined Network is a reality.

About Datavision:

Datavision is a network engineering consultancy focused on the emerging technologies and solutions around Software Defined Networking. Our suite of services around SDN helps or clients roll out SDN-centric infrastructure in an intelligent, structured way, helping to maximize the benefits and reduce the deployment risk of SDN:

***SDN Now* Lifecycle Management:**

- Datavision's SDN Now Lifecycle Management **reduces your capital and operational expenses while enhancing your network with agility, scalability and speed**. This approach covers the entirety from helping to define your business objectives to post-implementation support.
- SDN Now Lifecycle Management provides unprecedented programmability, automation and network control. By evolving beyond conventional hierarchical network structure, users get faster access and network support teams are free to make changes in significantly less time than current approaches.
- SDN Now Lifecycle Management provides Network Agility, ensuring maximum throughput and the ability to centrally make changes to network appliances to meet the user demands and achieve unprecedented scalability. Network Agility enables easy configuration changes without the need to manage on a device level, while also providing the agility to quickly and easily adjust security protocols and enforce universal policy guidelines.

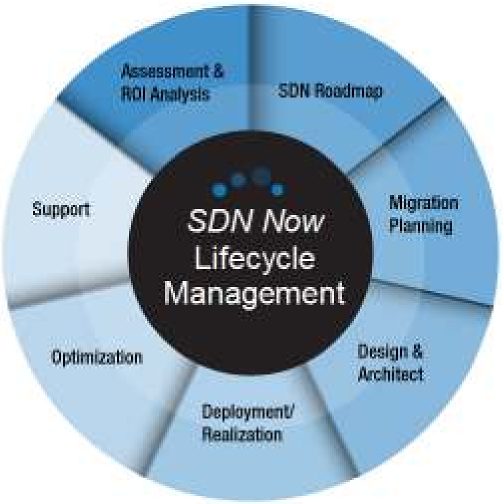

SDN Realization Roadmap

Experience a smooth transition from Proof-of-Concept (POC) to realization with Datavision’s expertise in full-cycle software-defined networking (SDN) implementation.

- Guides you through establishing a strategy and the implementation of an SDN, including a review of services, geographies and vendors.
- Creates a plan for POC, which allows you to test drive the SDN’s capabilities and ensures that it meets your requirements before rolling it out to the rest of the network.

Datavision’s future-proof network infrastructure approach meets the rapidly-changing needs for both bandwidth and application access while also ensuring user needs are met without draining resources or incurring major capital expenses.

SDN Quickstart Lab

- Implement Datavision’s SDN QuickStart Lab in your own testing and development environment. Solve your unique network puzzle with our specialized hardware and software configurations based specifically on your environment and user demands. Use the QuickStart Lab to test SDN capabilities and ensure KPIs are met in real-world simulations.

- Feel secure knowing that Datavision’s QuickStart Lab is being set up on your premises with a philosophy of vendor neutrality.

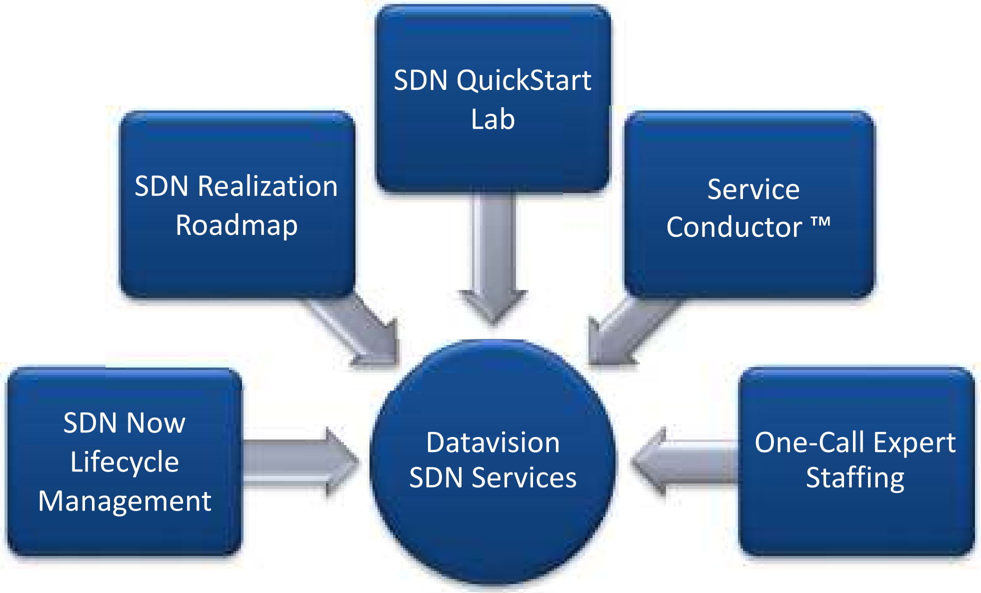

Datavision: Future-Proofing your networks: Helping you design and implement scalable and agile networks that are ready for your business - Today, and into tomorrow.